Report from Dagstuhl Seminar 14471

Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment

Edited by

Jon Crowcroft[1], Adam Wolisz[2], and Arjuna Sathiaseelan[3]

1 **University of Cambridge, GB,** Jon.Crowcroft@cl.cam.ac.uk
2 **TU Berlin, DE,** awo@ieee.org
3 **University of Cambridge, GB,** arjuna.sathiaseelan@cl.cam.ac.uk

Abstract

This report document the program and outcomes of the Dagstuhl Seminar 14471 *Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment*. At the seminar, about 27 invited researchers from academia and industry discussed the challenges and solutions for enabling universal and affordable Internet access. This report gives a general overview of the presentations and outcomes of discussions of the seminar.

Seminar November 16–21, 2014 – http://www.dagstuhl.de/14471
1998 ACM Subject Classification K.4 Computers and Society
Keywords and phrases Internet access, cost, wireless, regulation, social, deployments, protocols
Digital Object Identifier 10.4230/DagRep.4.11.78

1 Executive Summary

Jon Crowcroft
Adam Wolisz
Arjuna Sathiaseelan

Universal Internet Access is considered as one of the fundamental requirements in todays digital age as clean water, roads, schools etc. Enabling universal Internet access is one of the key issues that is being currently addressed at the European level via the Digital Agenda for Europe (DAE) as well as globally. Recognising the importance of broadband Internet, several developed countries have their own national broadband plan, such as the Broadband Development UK (BDUK) in the UK and the National Broadband Plan in the USA.

However a lack of access to the Internet and broadband is a global phenomenon that proportionately and negatively affects the poorest countries in world, where challenges to socio-economic development are most pronounced. It is estimated that only 41% of the worlds households are connected to the Internet. Half of them are in less developed countries, where household Internet penetration has reached 28%. This is in stark contrast to the 78% of households in more developed countries.

The disparity in access is even more worrying when one realises that the positive impact of increased Internet and broadband access is greater than any other ICT. In 2009, the World Bank found that in low and middle-income countries a 10% increase in broadband Internet penetration accelerated economic growth by 1.38%. Moreover, the positive effect of Internet and broadband on economic growth and social development are felt more in less

Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment, *Dagstuhl Reports*, Vol. 4, Issue 11, pp. 78–137
Editors: Jon Crowcroft, Adam Wolisz, and Arjuna Sathiaseelan
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

developed countries, like those in sub-Saharan Africa, than in more developed countries, creating opportunities for levelling up and greater equality.

The main barriers to the economic growth and social benefits identified by the World Bank include the cost of services and a lack of access to terrestrial and wireless networks. Indeed, there is general consensus upon the impact of these challenges, especially that of cost. Brahima Sanou, Director of the Telecommunication Development Bureau (BDT) at the ITU notes *Broadband is still too expensive in developing countries, where it costs on average more than 100 per cent of monthly income, compared with 1.5 per cent in developed countries*. There are indeed several challenges (political, regulatory, socio-economical, technological) to the realization of a Future Internet capability that will offer appropriate access to all parts of society.

The goal of our seminar was to bring together an interdisciplinary group of researchers from academia and research organisations as well as industry to understand the different challenges in enabling universal Internet access and to discuss potential solutions for solving some of the challenges.

This report provides an overview of the talks that were given during the seminar. We also had a dedicated breakout session with two groups specifically focussing on *Socio-Economic Models and Role of Community Networks* and *Internet in a box*. We also had longer informal discussions on specific focussed topics. The discussions and outcomes are summarised in this report.

We would like to thank all presenters, scribes and participants for their contributions and lively discussions. Particular thanks go to the team of Schloss Dagstuhl for their excellent organisation and support.

2 Table of Contents

3 Overview of Talks

3.1 The Internet is a series of tubes

Henning G. Schulzrinne (Columbia University – New York, US)

70% to 80% of the cost of deploying fiber networks is civil engineering. Thus, for most high-income countries, the challenge is to reduce the cost of fiber installation, not flying drones. The talk briefly discussed the experience with universal service funding and broadband build-out in the United States, and discussed the limitations of supposedly "cheap" approaches, as well as the wired- wireless trade-offs.

3.2 Monetization in the Internet

Falk Von Bornstaedt (Deutsche Telekom – Bonn, DE)

The Internet consists of about 50000 autonomous systems. Interconnections happen via mostly settlement free peerings or via IP-transit. If there are settlements, billing may happen on the basis of average traffic flow, on a flat rate, or in most cases on 95-percentile billing. This means that the buyer of IP transit cannot fulfil his wish of being billed on the average of his monthly basis. The seller would like to bill on the basis of maximum usage, because he has to provide and install this capacity, and even more to be prepared for singular events like big sports events. 95-percentile billing is somehow a compromise; it has strong elements of a peak load pricing.

Prices should lead to an optimal allocation of resources. In the Internet this means an even distribution of traffic over time. 95-percentile pricing for IP-Transit is sometimes giving false signals to the content providers. In 95-percentile billing the top 5% of all measurements are discarded, these are the highest values. As a concrete example, if Apple sends out an iOS upgrade, ISPs have an interest to spread substantial updates like iOS over several days. The current pricing model gives incentives to push all traffic into that period of 5%, so to make the peak even higher instead of levelling out traffic peaks. Any traffic put on top of the peak has a free ride with 95-percentile billing.

One possible consequence could be to move towards higher percentiles like 98-percentile. The option to charge for maximum values is problematic, since maximum values could be generated from denial of service attacks, and the IP transit buyer would not like to pay for, in worst case the IP transit seller might even have the incentive to order an attack.

3.3 Comparing Digital Exclusion with the Bottom of Pyramid (BoP): Economic Challenges and Strategies

Irene Ng (University of Warwick, GB)

Bottom of the Pyramid (BoP) was a phrase coined by U.S. President Franklin D. Roosevelt in 1932 when he referred to "the forgotten man at the bottom of the economic pyramid". In economics, the bottom of the pyramid is the 2.5 billion people who live on less than $2.50 per day. Bottom of the Pyramid is a phrase used by those interested to develop new business and economic models that deliberately target the poorest regions. When comparing the challenges of the bottom of the pyramid with those marginalised through digital exclusion, there are surprising parallels. Multi-national companies consider targeting BoP with great scepticism when it comes to doing business profitably often due to the perception of corruption, illiteracy, currency fluctuations and inappropriate infrastructures within the BoP. These factors and also the outdated representation of the poor hide the real potential of BoP markets. Similarly, those marginalised by digital exclusion are the vulnerable people, the poor, disabled, or elderly. Increased digital skills are of benefit to consumers (saving £254 a year), to public services, bring improvements in education; bringing older and isolated people to be more effectively engaging in their communities; helping adults back into work and improving health and social services. With digital skills having the greatest impact on the lives of the marginalised in society, it is alarming to see the latest figures from the Office for National Statistics that show that 6.4 million people have never been online. To face this challenge, digital exclusion is currently seen as a challenge for the state. This talk discussed the parallels between BoP challenges and that of digital exclusion and considers the lessons learnt and how new economic and business models could be applied to face the challenges of digital exclusion as a viable option for companies to serve and therefore enable public and private sector collaboration.

3.4 Business models for broadband rollout

Rüdiger Zarnekow (TU Berlin, DE)

The talk discussed a medium term deployment plan for fiber broadband rollout. The talk summarised the following key points:

1. FTTH should be deployed once it‘s capabilities are requested by the customers.
2. There is no market failure as long as the broadband supply meets the customer demand.
3. Business models can foster broadband deployment at various layers of a generic three-layer telecommunication value chain.
4. In rural areas typically only one fixed broadband infrastructure will be deployed.
5. Providers will compete at the Service layer.
6. The use of additional fiber capacities, which are deployed along alternative infrastructures, can decrease deployment costs for rural broadband provisioning.
7. To use these capacities, demand needs to be aggregated along the path of the alternative infrastructure.

3.5 Technical reasons for new data transmission charging models

Adam Wolisz (TU Berlin, DE)

Increasing share of Internet traffic is nowadays moving to wireless access. We read more and more about the growth of this traffic. And the operators are complaining that they have to increase massively the capacity of their networks without having a perspective of the proper revenues.

To give a more realistic perspective, it seems to be important to ask a set of basic questions: is the amount of bytes transported between the mobile device and the internet really closely related to the user needs for the content?

The analysis of recent research papers gives evidence of some unpleasant – if not entirely surprising – facts:

1. Many "free" applications generate side traffic (advertisements, analytics) in amounts significantly exceeding this overhead in comparable applications which have to be purchased.
2. Some of the most deployed software ecosystems use for the delivery of the same amount of content almost double the volume of transmitted data.
3. The potential for neither caching nor compression is utilized.

To improve this situation the change of economic model is necessary. The most radical way to go would be widely enforced charging the big content providers for traffic delivery. Charges for content from /to specific services (e. g., education, access to administrative services) could be covered by the local administration or sponsoring. This is only fair, as it can easily be calculated that the local administration saves a lot of money by moving to IT based interaction with the population. This is pretty easy to implement, as operators have very clear interfaces to the providers of large content volumes. In such model only the data form multiple small sources would be covered by a very low basic subscription.

An alternative model less controversial form the point of view of network neutrality purists might be at least a legal obligation to provide for any item of content a strictly standardized, unified label declaring the volume of the intended content and the volume of side content if any. In that case the consumer or the proper process running on his behalf – could make informed decision about selection of the content sources, form etc.

3.6 Avoiding building Hell 2.0 in a Paradise Built in Hell 1.0

Jon Crowcroft (University of Cambridge, GB)

As per ghostbusters, when unexpected things happen in society, it isn't obvious who to call. Indeed, BBN's report in to the events in NYC after 9/11 strongly advised against shutting off networks, based on data from emergency services that often, situational awareness was frequently based on data from citizens, as much as from specialists. This is common knowledge (see also Paradise built in hell, by Rebecca Sohit). One thing increasingly exercising interest is the shift from community to government as a social group grows beyond the level where trust via peer-pressure or just personal knowledge works – this is something that needs thought (one example wel documented is the self-help networks that emerge after disasters, and the very bad things that happen when the first responders arrive and try to take over without understanding the structures that have evolved – this transition applies in many arenas – even the Internet itself which didn't have very much spam or phishing or ddos till it hit a certain size – the emergence of anti-social or even byzantine behaviour is interesting and tools to combat it not well understood – think, also, Wikipedia edit wars, and non-terminating arguments in liquid democracy.

What could we learn and apply from those worlds in ours? Essentially, we require a set of tools to manage what looks at the human, power, and communications network layers, somewhat like a sequence of phase transitions from gas (opportunistic, DTN, dominated by altruism) through liquid (ad hoc, community mesh, peer-to-peer, peer production, mix of altruistic, rational, selfish, and Byzantine behaviours) to a solid (infrastructural, governance, service-oriented, customer-provider relationship managed, through economic/market) basis.

3.7 Towards Trustworthy Internet for All

Georg Carle (TU München, DE)

It should be considered a basic human right to have access to trustworthy Internet services, being compliant with personal rights and interests of individuals concerning their data and meta-data. There are a variety of threats against the trustworthiness of Internet services, including the threat of systematic espionage from intelligence agencies.

Network components from commercial vendors with closed-source software may have vulnerabilities due to software weaknesses and the possibility of backdoors. Apart from this, trustworthiness is affected by the country- dependent possibilities of security agencies to obtain legal access to personal data. As a consequence, service users do not know to which extent their data and meta-data is handled in a trustworthy manner by their service provider.

There are several possibilities to improve trustworthiness of Internet services. While it is hard to identify possible software weaknesses in closed-source products, components and services based on open-source software allow the community as a whole to systematically search for software weaknesses, and to assess trustworthiness of components and services in a transparent matter. This means that open-source software has the potential to lead to trustworthy Internet services. In order for this to happen, sufficient coordinated effort towards this goal is required. A promising approach for making such a coordinated effort happen is by fostering an ecosystem similar to the Linux ecosystem in which large and small commercial entities cooperate with a community of volunteers towards realizing trustworthy components and services.

There is also a potential to realize trustworthy Internet services for all by using decentralized authentication and authorisation, based on a Private Key Infrastructure with private X.509 Certificate Authorities associated with the homes of individuals, and using personal trust exchange. Further steps towards increased trustworthiness is realized by the Crossbear PKI notary service (cf. http://www.crossbear.org), which allows to detect and to localized TLS Man-in-the-Middle attacks.

As a concluding assessment, it appears impossible to ensure resistance against particular advanced targeted attacks today. At the same time, it appears to be a more realistic goal to raise the cost of large-scale attacks, and to make it difficult to have large-scale data gathering undetected.

3.8 Energy considerations for client-side Internet usage in developing regions

Saleem Bhatti (University of St. Andrews, GB)

As more users around the world come online, there is increased use of services and systems. More electricity is used world-wide, and this also increases the carbon emissions due to ICT usage: ICT already ICT produces 2% of worldwide CO2 emissions, second only (just) to the aviation industry, and those CO2 emissions are predicted to triple from 2008 to 2020.

In developing regions, there is far lower penetration of users: there are many more users yet to come online, so the potential increase in energy burden on those regions is greater than in developed regions. However, in those regions, the introduction of online services for application such as e-Health, e- Education and e-Government could also have the greatest benefits. For example, use of voice and video has high utility where there are a large number of languages in use and a high level of functional illiteracy. So, governments in such regions, e. g., India, are keen to bring these services to their citizens to improve overall standards of living. However, places like India also have power generation that does not meet normal daily needs, e. g., daily deficits of 6%-15%.

So, as we look at technology for deployment in such regions, to enable applications such as e-Health, we need to look more carefully for the energy usage of systems so that they can be introduced more easily without as a high an energy burden attached to their deployment.

There has been much research in how to improve energy efficiency of hardware components, server systems and datacentres, with applications of such research yielding great benefits in terms of energy usage and associated carbon emissions. However, there has been relatively little work on reducing energy-usage of client-side systems, especially by examining the energy usage of software systems and components.

Our position is that small improvements in energy usage of client systems can have a significant benefit when considering the multiplier of a large population of users [1],[2],[3],[4]. Additionally, improving software energy usage has the benefit of being applicable through software updates to legacy (non-energy efficient) hardware. Software improvements in energy efficiency work in complement to hardware efficiencies.

This talk looked at how changes in protocol usage and video codecs can have a measurable and potentially significant impact on energy usage of client systems, especially when considered for a large population of users.

References

1 O. Ejembi, S. N. Bhatti. Help Save The Planet: Please Do Adjust Your Picture. ACM MM 2014 – 22nd ACM Int'l Conf. Multimedia. Orlando, FL, USA. Nov 2014. DOI: 10.1145/2647868.2654897

2 Y. Yu, S. N. Bhatti. The cost of virtue: Reward as well as feedback are required to reduce user ICT power consumption. ACM e-Energy 2014 – 5th ACM Int'l Conf. Future Energy Systems. Cambridge, UK. Jun 2014. DOI: 10.1145/2602044.2602063

3 M. Abu-Tair, S. N. Bhatti. Energy Usage of UDP and DCCP over 802.11n IEEE. AINA 2014 - IEEE 28th Int'l Conf. Advanced Information Networking and Applications. Victoria, Canada. May 2014. DOI: 10.1109/AINA.2014.40

4 M. Tauber, S. N. Bhatti, Y. Yu. "Towards Energy-Awareness In Managing Wireless LAN Applications". IEEE NOMS 2012 – IEEE Network Operations and Management Symp.. Maui, Hawaii, USA. Apr 2012. DOI: 10.1109/NOMS.2012.6211930

3.9 The challenge of an Internet for all: can the grid be part of the solution?

Fernando M. V. Ramos (University of Lisboa, PT)

Around 80% of the worldwide population has access to electricity, a number that drops to 35% in the least developed countries. Nevertheless, in many developing countries the infrastructure exists, so the lines are already there. As such, in these countries power line communications (PLC) are arguably the only technology with a deployment cost comparable to wireless. This, we argue, makes it a technological solution worth exploring for Internet access.

Power grids today are being modernized through the introduction of communication networks dedicated to the management of energy transmission and distribution. Although the investment in smart grids today is concentrated in developed countries, smart grids may play an important role in the deployment of new electricity infrastructure in developing countries, by enabling more efficient operation and lower costs. For instance, in sparsely populated areas smart grids could enable a transition from one-off approaches to electrification (battery- or renewable energy-based electrification) to the development of community grids that then connect to the national grid.

Can the grid be part of the solution for the challenge of global access to the Internet? Some opportunities seem to exist. For example, exploring the existing power lines for Internet access using PLC; or leveraging on the distributed, ad hoc nature of community smart grid networks to explore novel low-cost access solutions. But many challenges remain: is the overall cost of these solutions reasonable, when compared to the alternatives? Can the existing infrastructure in these countries really be of some use? With many open questions and several unsolved issues, there's quite a big room for discussion and debate on the topic.

3.10 What rights should we have to the Internet?

Michael P. Fourman (University of Edinburgh, GB)

Internet access is increasingly a necessary prerequisite for participation in society. As those who are connected benefit, those who are excluded suffer economic, educational and social exclusion. Market forces do not result in Universal provision. Remote users value the Internet more highly, as it is more difficult for them to find substitutes for the services it enables; but Internet provision is most expensive in remote areas. Network effects mean that we all have an interest in getting more people online, so increasing Internet access brings many positive externalities. These positive externalities justify intervention to address the market failure to deliver universal provision.

However, a purely economic argument, based on such externalities has also failed to deliver anything close to universal provision. Those who remain unconnected, or with inferior service already isolated and/or disadvantaged, and their social exclusion is exacerbated by an increasing digital divide, as those who are well-connected enjoy increasing benefits from the digital economy, and existing non-digital services become uneconomic and are withdrawn.

We must either accept a growing digital divide, or provide stronger arguments for universal provision. In making these arguments we must balance Internet access against many other benefits of modern society, such as freedom, education, and health. Access to these benefits is often prioritised, because they are closely related to fundamental rights. Surely Internet access should be a secondary concern.

Of course the right to health does not entail a right to a hospital; nor does the right to education entail a right to a university. But rights such as these do affect policy. They entail a duty on governments to make best efforts for universal provision, and lead to levels of provision higher than would otherwise be justified. Information is power.

We argue that the changes enabled by modern information and communication technologies have fundamentally changed the balance of power, and that we must revisit established rights to freedom of speech and access to information and education if we are to avoid a fundamental digitally- enabled divide between those who can wield the digitally-enhanced power of information and those who cannot. This leads us to propose fundamental rights to store, process and communicate information. We argue that these are necessary for an inclusive digital society.

3.11 Life in the Slow Lane

Arjuna Sathiaseelan (University of Cambridge, GB)

Universal Internet access has become critical to modern life leading to many explorations of approaches to increase its availability. One means of addressing the cost of access is to share the WiFi of existing broadband connections utilising their unused capacity [2]. Systems that share home broadband in this way already exist (for e. g., FON), where subscribers' existing broadband connections are shared by making available a wireless network accessible to anyone with a Fon account. These networks are typically provided either as a paid service or incorporated as part of an existing broadband subscription, enabling subscribers to access their network's broadband service via other subscribers' access points while roaming. However, when considering whether to mandate such sharing, there are a number of questions that remain unanswered: (i) is there spare capacity on existing domestic broadband links to do this without impacting service? (ii) how would this capability be used? (iii) would such a service be useful to the target demographic? (iv) are existing subscribers willing to share their capacity?

The talk reports on a limited-scale feasibility study of a Public Access WiFi System (PAWS) [1] that answers these questions. The intent of PAWS is to explore provision of a restricted service that is free at the point of use, targeting demographics that want and need but cannot afford Internet access. As we are interested in the use of such a system in spatial and social context, we apply the "in the wild" method from HCI: lab-based studies fail to show appropriation and use in context; surveys/questionnaires reveal only stated preference rather than actual behaviour; and measurement studies fail to uncover the social challenges that are at least as important as the technical. The intent behind this study is not to make broad statements about the use of a system like PAWS, but to uncover a rich understanding of a small number of users so as to sensitise us to the underlying challenges in deploying such a system. Thus we do not collect the data necessary to comment on broader questions such as how to incentivise sharers or what the commercial or operational impact would be on ISPs. Instead, we make the following contributions: (i) there are a number of technical challenges, principally signal reach and the variation in both home router availability and ISP configuration, which suggest that the density of deployment required for a system like PAWS to provide free Internet access for the digitally excluded is quite high and probably requires regulation to be effective; (ii) perhaps surprisingly, free access to the Internet is not the instant success one might expect, although some citizens do find even relatively limited access of considerable use; (iii) many existing broadband subscribers appear quite willing to share their bandwidth locally for the common good without requiring significant financial incentive.

The experimental deployment demonstrated the potential utility and viability of a Public Access WiFi Service. Even in the relatively complex form it was presented, users signed up and used it for legitimate, socially beneficial purposes. The primary citizen was clearly sufficiently happy to use the service extensively, and all citizens used it for a wide range of purposes: this suggests that the overall performance of the system was acceptable once past the interactional overheads of signing up and connecting to the VPN. We also observed that several users, not just the primary users, made use of PAWS for the legitimate (foreground) uses it was intended, such as banking and retail. They also brought to light the value of other (background) uses that we did not originally envisage, e. g., software updates.

We believe we can now partially answer the questions we posed: (i) many existing domestic broadband deployments do have sufficient spare capacity that sharing some of it would not be a problem for the owner; (ii) such a capability will be used in a range of ways, some more acceptable than others, suggesting there may be a need to manage the uses to which the shared bandwidth can be put; (iii) the service was useful to at least some of the target demographic; and (iv) given that the compensation we offered failed to sway people who refused to participate, and some sharers did state a willingness to share bandwidth simply for the common good, at least some existing subscribers do appear willing to share their bandwidth without requiring any further incentive. Further research is required to better understand the incentive structure for both the sharers as well as the network operators to enable a service such as PAWS. One plausible incentive structure would be for enabling third party stakeholders such as grassroot user communities or local government who may have a socio-environmental objective rather than purely economical to manage the PAWS service (where sharers could get a small council tax rebate for sharing their connection) while for the network operator they get paid (again) for leasing out the unused capacity (which has already been paid for) [3].

Finally, the ISPs clearly have a role to play in this. Traditionally they have argued against giving free access based on the claim that traffic on their backbone networks costs them money. However, networks are (typically) provisioned for peak use, and citizens using PAWS would not significantly increase peak use (they will not increase the number of access links into ISP networks, and there would be little spare capacity for them to use at peak times anyway). Thus they would increase off-peak but not peak use; any slight increase in peak use could be further mitigated through deployment of less- than-best-effort protocols from the access points into ISP networks, enabling ISPs to further degrade the free service if they found it necessary. We might anticipate regulation would have a part to play in controlling such degradation. There is evidence that some ISPs (e. g., AT&T) are already applying such selective traffic grooming practices for commercial gain; PAWS would encourage them to do so for social benefit.

References

1 A. Sathiaseelan, R. Mortier, M. Goulden, C. Greiffenhagen, M. Radenkovic, J. Crowcroft, D. McAuley, A Feasibility Study of an In-the-Wild Experimental Public Access WiFi Network, ACM DEV, San Jose, December 2014.
2 A. Sathiaseelan, J. Crowcroft, LCD-Net: Lowest Cost Denominator Networking, ACM SIGCOMM Computer Communication Review, April 2013.
3 A. Sathiaseelan, C. Rotsos, C.S. Sriram, D. Trossen, P. Papadimitriou, J. Crowcroft, Virtual Public Networks, 2nd IEEE European Workshop on Software Defined Networking (EWSDN), Berlin, October 2013.

3.12 Community Networks: Access to the Internet or a Different Internet Model?

Renato Lo Cigno (University of Trento, IT)

Is a different (from the global Internet and Cellular Networks) model for urban communications and networking conceivable? Can Community Networks, now flourishing in many parts of Europe and the world, be the next "big thing" in networking, for once considering the needs of people and urban evolution as a key element, and not as a side effect of technology or business? This talk introduces some of the new trends on people-driven communication efforts, normally community networks, as alternatives to the business-based organization of the networks we use daily.

Let‘s start from a question: Is the Internet (as we know today, based on commercial services and most of all on “free” services paid by ads and loss of personal data, privacy and freedom) the only communication model possible? The follow up is a discussion of one possible alternative, or better complement, to the Internet and cellular networks: Urban Wireless Community Networks (UWCNs). We stress the urban dimension not because WCNs cannot grow in the countryside or cannot flourish in rural environments, but because we think that inside cities they can grow to have a role which is not only an access network (to the Internet) or an instrument to reduce the digital divide (whatever this means), but they can sustain entirely novel communication paradigms that not only break the Telco and ISPs oligopoly in communications, but also shatter the much more subtler trust of information mediators and aggregators (search engines, centralized social networks, etc.) that are today controlling the way information flows around the world, as the “Snowden Affair” has recently brought to the general public attentions.

Indeed, neither WCNs nor the challenging the Internet model are brand new topics [1],[2], and indeed also community networking go back at least 20 years [3], but today technology makes them more feasible than ever by making the technical design and management easier. However, what makes UWCNs different from the global Internet is their cooperative, bottom-up nature, and it is important to understand why this new model of social networking is emerging.

If we go back to the to the origins of Internet and TCP/IP we find out that the technical approach was distributed, cooperative and social event the web and most of all the “web 2.0” approach was conceptually peer-to-peer. Centralization and concentration are just market-driven economic trends, lead by the fact that the Internet is controlled, in the end, by a small oligopoly of very large players (content mediators and telcos). So, is economic sustainability the true reason for centralization? We claim this is not really true, and a more distributed, human and social-centric Internet is not only feasible but indeed desirable for the healthy growth of society and economy. What is lacking today is an holistic approach, a novel "science" that let scientists, engineers, economists, sociologists, but most of all politicians and decision makers understand the global context, the consequences of regulations and legal obligations, the far fetching results of apparently local decisions (as for instance implementing a UWCN supporting local services and independent – from the global Internet – communications for a community). Some of these issues have been discussed already [4], but much more remains to be done, specially from a interdisciplinary point of view, trying to make people from as different scientific fields as ICT, sociology, and law, cooperate, find a common language, establish a "new science" that will unleash the potential of UWCNs as means for a new perspective on global communications. We note that this approach will intrinsically reduce the digital divide both in developed countries and in the developing world, as a communal (not communist!) approach is far more sustainable, both economically and socially than a market-based one.

We have recently analyzed some technical challenges, as realizing distributed firewalls adapt for low-energy networking [5], or improving distributed filtering [6], but these are only timid approaches, in our specific field, to tackle a huge problem that requires many more efforts from the community and from society at large to be solved.

References

1. R. Battiti, R. Lo Cigno, M. Sabel, F. Orava, and B. Pehrson, Wireless LANs: From WarChalking to Open Access Networks, Mob. Netw. Appl. SpringerVerlag, June 2005, http://link.springer.com/article/10.1007Fs11036-005-6422-4
2. R. Lo Cigno, Untethered Local Communications: From Wireless Access to Social Glue, In IEEE WONS, Bardonecchia, Italy, Jan. 2011, DOI: 10.1109/WONS.2011.5720198
3. D. Schuler. New Community Networks: Wired For Change. AddisonWesley Publishing Company Inc., 1996, http://www.publicsphereproject.org/ncn/, last visited May 2014.
4. R. Lo Cigno, and L. Maccari. Urban wireless community networks: challenges and solutions for smart city communications. ACM international Workshop on Wireless and Mobile Technologies for Smart Cities (WiMobCity), Philadelphia, PA, US. 2014, DOI: 10.1145/2633661.2633669
5. L. Maccari and R. Lo Cigno, Waterwall: a cooperative, distributed firewall for wireless mesh networks, EURASIP Journal on Wireless Communications and Networking, 2013, http://jwcn.eurasipjournals.com/content/2013/1/225
6. L. Maccari and R. Lo Cigno. Betweenness estimation in OLSR-based multi-hop networks for distributed filtering. Journal of Computer and System Sciences, 2013, DOI: 10.1016/j.jcss.2013.06.018

3.13 Guifi.net experiences – Network infrastructure as a Common-Pool Resource

Roger Baig Vinas (Guifi.net – Barcelona, ES)

Bottom-up initiatives have brought connectivity to places that otherwise would have been left disconnected for a long time. Community Networks, a particular case of Bottom-up, have ported the concepts of the commons and non- speculative access to telecommunication infrastructure, proposing a new model for designing, deploying and operating networks and have developed the necessary tools to implement an entire economic ecosystem. The emergence of this telecommunication model opens the question how to integrate it in the, so far, unidimensional public-private debate, bringing new challenges in many fields, such as the regulatory framework, the internet governance, the internet management, etc.

Community Networks have proved that they can contribute in great manner in reducing costs of internet expansion, thus to be able to play a significant role in achieving and affordable internet access for everyone.

The underlying principle behind guifi.net is the firm conviction that the optimal way to design, deploy and operate a network is doing it as a common-pool resource (CPR), being the network infrastructure the core resource, which is nurtured by the network segments the participants deploy to reach the network or to improve it, and the fringe unit is the connectivity they get.

In our talk we discussed our fundamental ideas, presented some of the tools the guifi.net community have developed to put our ideas in practice and the results we have achieved. We are confident that most of these ideas and tools can be reused to bring connectivity to other places. All our materials, documents and software included, are publicly available.

The tools presented are:

1. The network license, the Compact for a Free, Open & Neutral Network (http://guifi.net/en/FONNC), which precisely defines the rights and the duties of the participants. It is mandatory to accept it to join the network. It is written to be enforceable under the Spanish legislation. Legal certainty is essential to stimulate participation and investment, which in turn, is at the base of any economic activity.
2. The network management and provisioning software tools, a set of software tools to ease the design, deployment, management and operation of the network in a self- provisioning style.
3. The conflicts resolution system, systematic and clear procedure for resolution of conflicts with a scale of graduated sanctions. It consists of three stages, conciliation, mediation and arbitration, all of them driven by a lawyer chosen from a set of volunteers. The cost of the procedures are charged to the responsible party or to both parties in case of a tie.
4. The collaboration agreements, 1) for professional activities and 2) for the rest of entities. Professionals can choose between three degrees of commitment, fully-committed (all the infrastructure deployed is in commons), partially-committed (sometimes they deploy in commons, but others not), or not-committed (they use the infrastructure available in commons, but they never contribute to it).
5. The economic compensations system, has been developed and implemented to compensate imbalances between investment in the commons infrastructure and use the professionals make of it. Expenditures (hardware and manpower) declared by the participants (the professionals but also the volunteers) are periodically cleared according to the network usage of the professionals and their degree of commitment. Fully-committed professionals are charged just to cover the administration costs, partially-committed are charged to meet the investments made by the fully-committed and non- committed are charged to meet the investments made by the fully-committed plus a an extra to ensure the CAPEX. At the moment the rates are 10%, 50%, and 100% of the cost of the network usage of each of them. The Foundation centralises and manages the billing system (each professional only makes or receives a single payment). The mechanism is applied at PoP level.

The results:

1. Over 26,500 working nodes
2. More than 50,000 kms of links

3. Multiple technologies used: OF, WiFi, etc.
4. Total estimated guifi.net users: over 80,000
5. Places we 90% of the households have guifi.net access
6. Total estimated OPEX: 6.5M€
7. Total estimated CAPEX: 2.75M€/year
8. CATNIX (Catalan exchange point) member
9. RIPE-NCC member
10. 2 Internet uplinks, 1Gbps each
11. 15 SMEs operating over guifi.net infrastructure
12. Hundreds of volunteers
13. A Foundation + number of local organisations

3.14 Experiences and Research in Community Networking: the Community-Lab.net testbed

Leandro Navarro (Polytechnical University of Catalunya – Barcelona, ES)

Experiences, research challenges, lessons learned, from researchers working together with community network activists to understand, model, engineer and experiment on sustainable ways to bootstrap connectivity and digital services at the scale of 7 billion people.

3.15 Some Challenges for Rural Broadband: Experience from Scotland

Mahesh Marina (University of Edinburgh, GB)

This talk gives a short overview of our work on the Tegola project that focused on bringing wireless Internet access to some of the remotest parts of Scotland. The talk also briefly presented a recent work on analysing mobile coverage in Scotland. Finally the talk reflects on the experience from these projects and outline some challenges for rural (mobile) broadband.

3.16 ICT4D and TVWS

Marco Zennaro (ICTP – Trieste, IT)

TV White Spaces (TVWS) technology and regulation has the potential to make connectivity both technically and economically feasible in Developing Countries where affordable access remains a challenge. The superior propagation characteristics of TVWS technology make it particularly well suited to connecting remote communities. How will this new technology affect ICT4D projects? What are the research challenges in TVWS?

3.17 Decentralized, multi-hop networks: Are They really different from the Internet?

Leonardo Maccari (Università di Trento, IT)

Community networks (CNs) represent a bottom up effort to build independent, community-owned network infrastructures. In light of the recent NSA scandal and the global trend towards centralization of networks and services, they have received large attention, since they represent a completely different approach.

CNs are changing from their initial model. In the past they were created by local hackers in regions where there was no broadband connections with the goal or replacing the last mile. Today, they are blooming in many countries, and in many cases they grow in densely inhabited cities that are already covered by broadband connections. So why do they grow? one reason is that they offer a different model of networking, that embeds social and political motivations that make them different from any other network. People build CNs because they want to own their networking infrastructure, neutral, uncensored, hard to disconnect. The last two years have witnessed the Datagate scandal, a large discussion (both in Europe and USA) about Network Neutrality, and governments that keep disconnecting, censoring and filtering the Internet to repress internal dissent. This produced a high media attention for CNs not because they are a cheap way to achieve Internet access but because they are a community-owned infrastructure that offers some kind of protection from privacy invasion, censorship and disconnection.

But is this effectively true? the talk presents some initial results on the study of some Wireless Community Networks showing that they are indeed more centralized than they seem, both topologically and socially. The topological features of CNs are indeed similar to other networks: a few fraction of the nodes (about 5% in the three CNs that were considered) are vital for the whole network. They are involved in approximately 90% of the shortest paths, and their failure would provoke the fragmentation of the network in many small disconnected islands. Moreover, the analysis of the ownership and of the interactions in the mailing lists in one of the community under analysis show that a few individuals own a large fraction of critical nodes and occupy the discussions in the mailing lists. Thus, in spite of the effort of the communities to maximise resilience, decentralization, and participation CNs are not much different to the Internet.

So the questions this presentation asks are: how can we help the communities to build networks that are truly decentralized? What are the tools they require to fulfil their goals?

3.18 Bandwidth for free?

Michael Welzl (University of Oslo, NO)

When we look at what is available in the academic literature as well as already deployed systems to support LBE communication, we can find quite an arsenal of possible mechanisms. These mechanisms differ not only in how they work but also in how they can be applied. For example, some methods involve using the network only in off-peak hours such as 3–5 AM.

Such usage requires scheduling of traffic, which may be fine even for large downloads, but limits usage to schedulable (i. e., non-interactive) applications. This is different with LBE congestion control such as LEDBAT [1], which may operate in parallel with other traffic. In addition to the benefits derived from the congestion control, LBE-marking traffic via DSCP in the IP header could yield more appropriate queuing in routers [2], and traffic engineering could then give such traffic a special treatment too.

LEDBAT is only one out of many congestion control possibilities. Depending on the requirements of the usage scenario, it may sometimes either be too aggressive or not aggressive enough. Alternatives include different LBE congestion controllers [3] and applying priorities in congestion control [4],[5]. Also, it is possible to allow users to transmit or receive traffic with "normal" congestion control whenever it is known that such traffic does not share a bottleneck with "foreground" flows (e. g., a user's downlink is not necessarily always the bottleneck, e. g., with P2P applications where it is more likely for the sender's uplink to be the bottleneck). There are methods to detect whether flows share bottlenecks, with one recent mechanism currently being proposed for standardization in the IETF [6]. These methods generally require both ends of a connection to be involved, but they could also be applied at the ends of just the link that is of interest.

LEDBAT-style mechanisms make sense when congestion controls compete. This is often not the case for a user's uplink, which usually mostly carries TCP acknowledgment (ACK) packets. These packets are small and usually do not reach a number where they begin to congest a link; this means that small data packets could be inserted in between pauses of ACKs, or LBE traffic could even be piggybacked on ACKs (which often do not carry data).

Here, we could only scratch the surface regarding all the LBE possibilities that are already known or can be imagined; it would probably make sense to create a taxonomy of such mechanisms, with their pro's, con's and limitations as well as their potential for application in a GAIA context.

References

1. S. Shalunov, G. Hazel, J. Iyengar, M. Kühlewind: "Low Extra Delay Background Transport (LEDBAT)", RFC 6817 (Experimental), December 2012.
2. R. Bless, K. Nichols, K. Wehrle: "A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services", RFC 3662 (Informational), December 2003.
3. David Ros, Michael Welzl: "Less-than-Best-Effort Service: A Survey of End-to-End Approaches", IEEE Communications Surveys and Tutorials 15(2), pp. 898–908, 2013.
4. J. Crowcroft and P. Oechslin. Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP. In ACM SIGCOMM Computer Communication Review archive, volume 28, pages 53–69, July 1998
5. Dragana Damjanovic, Michael Welzl: "An Extension of the TCP Steady-State Throughput Equation for Parallel Flows and its Application in MulTFRC", IEEE/ACM Transactions on Networking 19(6), December 2011.
6. David Hayes, Simone Ferlin, Michael Welzl: "Practical Passive Shared Bottleneck Detection Using Shape Summary Statistics", IEEE LCN 2014, 8–11 September 2014.

3.19 Can we fit the Internet in a box?

Gareth Tyson (Queen Mary University of London, GB)

What is the Internet? If you ask a network engineer they would explain the many details of TCP/IP. However, the everyday person might likely respond with "Facebook", "Twitter" or "Netflix". This talk will explore the feasibility of capturing these applications and services in a single locally usable "Internet Box". The box will operate independently of the rest of the Internet, allowing those without traditional connectivity to use the "Internet" in an simulated and disconnected manner. If this is proven feasible, many localities that possess no connectivity could hopefully start to use Internet services immediately. The talk will argue why this is a positive first step towards global access for all.

3.20 The Liberouter Neighborhood Networking Platform

Jörg Ott (Aalto University, FI)

Mobile opportunistic networking utilizes device-to-device communication to provide messaging and content sharing mechanisms between mobile users without the need for supporting infrastructure networks. However, enabling opportunistic networking in practice requires a sufficient number of users to download, install, and run the respective routing and application software to provide sufficient node density, and thus connectivity for the network to actually function. We provide a system called liberouter [1], an embedded Linux box based upon the Raspberry Pi platform, that allows bootstrapping mobile devices; serve a wireless access point and storage platform; and can mesh with each other. We explore reaching out to nodes that have not (yet) installed any dedicated software to: (1) allow them to access public content in an opportunistic network to possibly seed their interest and (2) instrument them to assist as (limited) message carriers to improve connectivity [2]. We present three applications to create and interact in temporal (instant) communities.

References

1 Teemu Karkkainen, Jörg Ott: Liberouter: Towards Autonomous Neighborhood Networking. Proceedings of IEEE/IFIP WONS, March 2014.
2 Marcin Nagy, Teemu Karkkainen, Jörg Ott: Enhancing Opportunistic Networks with Legacy Nodes. Proceedings of the ACM MobiCom CHANTS workshop, September 2014.

3.21 An interdisciplinary perspective on DIY networking: the case of the Internet Jukebox

Panayotis Antoniadis (ETH Zürich, CH)

The primary role of a local DIY network can be either the provision of cheap Internet access, as in the case of most community wireless networks, or the support of local applications, as in the case of the liberouter or the piratebox (and a fast growing number of related projects [1]). Depending on the choice of this primary role there is a wide variety of design options at different levels that will determine the successful deployment of such a network [2]. The talk reports on recent efforts to bring together researchers and practitioners from various fields around the (hybrid) design of such local networks: the recent Dagstuhl seminar on "DIY networking: an interdisciplinary perspective" [3] , the 3rd EINS summer school "From smart cities to engaged citizens" [4], and the LSE workshop on "The Alternative Internet(s): state of the art and possible futures" [5] . It then focuses on a novel application, the Internet Jukebox, that aims to combine the two possible roles of DIY networking, Internet access and local interactions. More specifically, the Internet Jukebox wishes to experiment with a different way for sharing a very thin Internet connection amongst a local community of users: enable them to decide all together on a single content item to be downloaded in every pre-defined period (according to the capacity of the link) and possibly projected in a public place. So, like in the case of a musical Jukebox the selected content will be the only choice for everyone. However, the selection will not be subject to payments but it will be the outcome of a democratic process, which we call "collaborative Internet consumption".

References

1 http://nethood.org/links.html
2 P. Antoniadis, B. Le Grand, A. Satsiou, L. Tassiulas, R. Aguiar, J.P. Barraca, and S. Sargento, Community building over Neighborhood Wireless Mesh Networks, IEEE Technology and Society, 27(1):48–56, 2008.
3 http://www.dagstuhl.de/14042
4 http://internet-science.eu/summer-school-2014
5 http://www.iscc.cnrs.fr/spip.php?article1912

3.22 Providing Connectivity or Avoiding Censorship by Mobile Peer-to-Peer Communication? Fiction or Practical Solution?

Karin Anna Hummel (ETH Zürich, CH)

Do-it-yourself delay tolerant networks that utilize user mobility to carry data are still an appealing participatory approach to provide multi-hop peer-to-peer connectivity when there is limited access to the Internet, or the access is unsecure as under repressive political regimes. Examples of limited Internet access are rural African and Indian regions. Another example recently covered in the media is Hongkong, where the basic-democratic movement makes indeed use of an ad-hoc communication technology to circumvent blocking of communication (cf. FireChat app). To serve these application use cases, the talk introduced two possible system architectures for mobile peer-to-peer networks: with and without stationary databoxes in addition to mobile carrier devices. The talk raised the question about the practicality of peer-to-peer communication in the envisioned scenarios. First, the talk discussed whether we have already a good understanding of mobility flows and contacts and whether these flows are sufficient to provide reasonable connectivity and network capacity. The talk showed how the cellular network can be utilized as a ubiquitous sensor of human mobility in principle, and exemplifies the approach for a developing country, namely Ivory Coast leveraging mobile phone data provided by Orange Ivory Coast. In the second part, the talk targeted technical obstacles for the deployment of mobile peer-to-peer communication with an emphasis on energy consumption.

3.23 Communication Issues and Challenges in the Amazon Region – A special focus at Itaituba, Brazil

Weverton Cordeiro (Federal Institute of Pará – Itaituba, BR)

Delivering affordable and effective Internet services in the Amazon has long been a challenge, mostly because of the hostile environment conditions (which imposes difficulties to wireless communications) and high cost of deploying wired infrastructures (such as optical fibers). As a consequence, available services are often overpriced and do not meet the needs of local population (e. g., very low bandwidth), not to mention the intermittent failures and service black-outs that occur very often. This situation is made worse by the apparent inability of existing protocols to perform adequately on narrowly constrained, congested links. In this talk, we provide a glimpse on the current panorama of Internet Services in the Amazon region, focusing on the case of Itaituba (a small town of about 100,000 inhabitants in the midst of the Amazon, Brazil). We also enumerate and discuss some challenges and opportunities for research, based on personal experiences of Internet usage in that region.

3.24 Which infrastructure for a better Internet in Africa?

Roderick Fanou (IMDEA Networks Institute- Madrid, ES)

Reasons for low penetration and low quality of Internet access in Africa are numerous: high Internet access costs inherent to energy instability, transit costs, network operation costs, lack of infrastructure in rural areas, lack of content in Africa, as well as the preference of users for popular Google, Facebook or Youtube content mostly hosted outside Africa which lead to a constant loop (no local content no peering, no peering no local content). We review approaches to tackle this problem. We suggest a better energy provision to the industry (by Africans governments or private companies) which could be boosted by competition in this sector, an orientation of electrification politics towards solar energy storage and furniture, renewable energy, gas and nuclear centrals, as well as the establishment of a climate of fairness and cooperation/partnership by the regulations in the telecoms market. Most importantly, ISPs could increase peering to save on transit costs. They could also implement traffic engineering and efficient routing to keep local traffic local. Meanwhile, ISPs should invest in terrestrial optical fiber networks deployment or alternative technologies for the middle mile in order to reach more customers. Creation, development of local content and web content hosting should also be encouraged. Considerable efforts are currently being done on the continent to achieve these objectives but they need to be multiplied, as Africa is massive.

4 Working Groups

4.1 Socio-Economic Models and Role of Community Networks

Roderick Fanou, Michael P. Fourman, Thomas Huhn, Renato Lo Cigno, Leonardo Maccari, Mahesh K Marina, Henning G. Schulzrinne, and Marco Zennaro

We identified four targeted groups when providing universal Internet access and separate issues given their types:

1. Individuals (with differences in availability and use by age, gender, education and to income, where the latter two are strongly correlated);
2. Schools and libraries: installing Internet at schools and libraries lead to the adoption at home;
3. Small businesses: Currently, many small businesses don‘t have any web or map presence. This reduces their ability to benefit from tourism, for example;
4. Health care facilities: hospitals, clinics, pharmacies, long-term care facilities.

Collectively, schools, libraries, public safety and hospitals are also called community anchor institutions.

We then focused on the following three dimensions:

1. Affordability: lack of income.
2. Availability: The high cost of deployment and maintenance (as well as lower income) typically causes broadband to be less available in rural areas. However, the proportion of served homes that subscribe to Internet access is often higher in rural areas than in urban areas. Availability is not binary, as Internet access may be available, but not reliable, or too expensive.
3. Relevance: Individuals and organizations may not see how Internet access meets their needs or furthers their objectives.

What do people use broadband for?

1. Content consumption
2. Communications: for coordination of personal and commercial affairs; for social connections to relatives and friends, often far away.
3. Applications: Increasingly, Internet access is no longer visible as web browsing. For example, people may use Internet-based mapping applications, but may not be aware that these require Internet access. Mapping applications, in particular, can be very useful for low-income countries.

We then answer to the following three questions:

1. How do you make broadband relevant? We need distinguish two aspects:
   a. Reasons why people find relevant or not.
   b. The factors that influence or motivate people to have it. In places where there may not be roads in time of disaster, for instance, talking to your relatives as you talk with your neighbours, communicating with the governments, getting services, participate to the development while your region is isolated due to a disaster, are those benefits that you can get from installing Internet. If we increase relevancy, we may also improve availability, by creating demand.
2. How to reduce the cost?
   Encouraging and enforcing sharing of critical facilities (from the middle to the last mile/connectivity) may help reduce cost. This could be improved by the regulations.
3. How do you make it widely available?
   Availability needs to include the notion of affordability.

We discussed ensuring that ducts or dark fiber are installed whenever a new road is built or an older road is resurfaced, or a railroad track is (re)built. Such fiber assets are needed for controlling traffic lights and train signals.

Similarly, all new homes should be equipped with fiber at time of construction, rather than adding it on later at great expense. The availability of fiber and Internet access should be part of the description of a home in real-estate web sites.

There are two models for increasing availability: In the bottom-up model, communities or NGOs build local networks, e. g., within a town or region, and link them later to the Internet. There is no empirical data, however, on how big such networks have to be so that they attract a sufficient number of users and a sustainable business model. In the top-down model, a national organization provides services to member institutions or communities. Examples include the national research and education networks (NRENs) such as Internet2 and JANET. For residential Internet access, the amount of capital needed may exceed the ability of private entities to finance. For example, the cost of running fiber networks to every household in Germany has been estimated at 80 billion Euros. While large, such expenses can be amortized over long time periods.

In many situations, building new networks requires a significant initial capital investment before revenue accrues to the ISP or community network. Thus, loans with reasonable terms can be very helpful. For example, the US Department of Agriculture provides low-interest telecom loans for rural development. It is not clear if larger banks understand telecom well enough to be a patient lender. Such loans may also facilitate the deployment of shared middle-mile infrastructure that allows small (wireless) ISPs to obtain Internet connectivity at reasonable rates.

Regulators can ensure access to backhaul and peering at fair and reasonable terms and rates. In UK, ISPs for instance are not interested in a network with less than a million of customers.

We discussed the use of the franchise business model for providing residential access. While common in other areas such as hospitality and services, we have not seen much of this for providing Internet access. Related models include the credit unions (cooperative banks) in many countries. Franchise models allow sharing of scalable resources, such as routing and security expertise, billing and network operations systems. Many common services, such as billing or network diagnostics, can be provided remotely.

4.2 Internet in a Box

Jon Crowcroft, Gareth Tyson, Jörg Ott, Arjuna Sathiaseelan, Georg Carle, Adam Wolisz, Michael Welzl, Karin Hummel, Panayotis Antoniadis, and Weverton Cordeiro

There are many areas of the world that have poor Internet connectivity, and extending high speed access is many years away. This section explores the potential of creating a locally available "Internet Box" that can be placed in a village. The Box will possess storage and computation capacity, alongside local network connectivity. By connecting to the Box, local residents will be able to get an Internet-like experience, accessing all content and services locally present. If a perfect prediction of content consumption can be achieved, the Box would truly be able to emulate a high speed "Internet" connection.

Initial Concepts. The Internet in a Box concept could be realised in a number of ways. This subsection details various "tiers" of functionality. Potentially, these tiers could be built iteratively one-by-one. Further, which tier is realised will vary heavily based on the target area of deployment.

Tier 1: Curated storage: The simplest model would be to provide a single box that possesses static storage and local network connectivity. The Box will contain a large block of storage that is filled with curated static content. The person selecting what is stored is left undefined. Local WiFi allows users to connect to the Box. Upon connection, users will be able to access a locally available portal that gives access to all content on the Box (e. g., via a web browser). In line with common usage, the most sensible portal might be a search engine interface but more structured access to content categories could be supported, too.

Tier 2: Curated and local user generated storage: The next tier would extend the Box to support user uploaded content. This would allow users to create and share their own static media (e. g., webpages, pictures, videos). The uploaded content would then be integrated into the portal to make it locally available to all other users.

Tier 3: Curated and local user generated storage and services: The next tier would introduce active services onto the Box. This would move beyond static content provision. Potential services could be online social networks and voice communications. This would be limited to local interactions.

Tier 4: Storage and services with one-way connectivity: The next tier would add one-way inbound communications to the Box. This is likely to be periodic and, potentially, unpredictable. This would allow the Box to receive one-way information from external parties.In particular, parts of the curated box contents could be updated on a regular basis.

Tier 5: Storage and services with time shifted two-way connectivity: The next tier introduced two-way communications. However, this will not necessarily be synchronous. Request/response intervals could be in the order of hours or days. Further, there may be extended periods without any connectivity whatsoever. This would allow usage data to be uploaded from the Box to a cloud service, which could then later send new content.

Tier 6: Storage and services with two-way connectivity: The last tier considered adds two-way synchronous connectivity. Unlike Tier 5, the Box will have on-demand connectivity.

Challenges. There are many issues and challenges with building an Internet in a Box. They can be categorised in various dimensions. This subsection explores some of the key problems.

Networking issues. There are a number of network issues. At a minimum, the Box must offer local network connectivity. Beyond this, it may be possible to provide some level of backhaul connectivity. This could be used for non-local interactions, as well as exchanging data for refreshing the Box‘s storage. These issues cut across multiple layers, as the role of network mediation might be performed at a content layer (i. e., selecting what gets sent over the network). Several key questions should be answered in this area:

1. How can the limited backhaul capacity be best utilised?
2. How can local network connectivity be provisioned?
3. How can content be prioritised for managing sequential delivery to the box (via a bottleneck path)?
4. How can content be transformed to address the needs of the network medium and box capacity?

It is envisioned that an “Internet in a box” could have three buttons: green, for schedulable traffic; red, for “direct Internet” (urgent communication); yellow, for lower-priority direct Internet (i. e., network traffic that does not get in the way of a potential “red” user). In an Internet context, where “normal” traffic is handled according to the “Best Effort” service model, lower-priority communication has often been called “Lower than Best Effort” (LBE).

Repository management and access issues. On deployment, the Box must be filled with important data in its content repository. An obvious challenge will be devising algorithms to define what content should be placed in the Box‘s storage. This raises several further research questions. Storage in the Box will likely be separated into curated and auto-selected content. The former is a key social issue that relates to who should have the power to curate what content gets stored. The latter is a greater technical challenge, as it will be necessary to automatically load in teresting content onto the box in a fair manner. This must address highly heterogeneous communications capacity, result in unpredictable intervals between backhaul connectivity. Key research questions to be explored include:

1. How can influence the content on the box?
2. How can they influence the content on the box?
3. How can fair sharing (of content capacity) be measured and ensured?
4. How can variations in communications intervals be managed?
5. What is eligible to be stored on the box?
6. How can the utility of content be measured, e. g., for replacement strategies?
7. How can the content of the box be made accessible for a user interactivity perspective?

Security and ethical issues. The introduction of any new technology into an area will create several security and ethical questions. It is particularly important to consider legal and cultural matters: What is acceptable in one region, might be unacceptable in another. Privacy relates closely to this, as it is vital to ensure that people cannot use the Box to monitor others. Important issues include:

1. How can the Box ensure that it adheres to legal constraints of the area?
2. How can the Box ensure privacy for users accessing it?
3. How can the Box reflect cultural aspects of its area?
4. How can the Box offer authenticity of content and services offered locally?

Deployment practicalities. A number of extremely challenging deployment issues are apparent when deploying a Box. Each box will need to be robust and sustainable. These factors must then, of course, be traded against price. Relevant requirements include:

1. The Box cheap must be cheap enough for wide-scale deployments.

2. The components used must be readily available on a commodity basis.
3. The Box must have a sustainable power source.
4. The Box must be physically robust relevant to the conditions of its location.
5. The Box must be appropriately and strategically placed to provide coverage to people.
6. The Box‘s storage and computation must be appropriately dimensioned relevant to price range and service needs.

Example: Collaborative Internet Consumption and the Internet Jukebox. The Internet Jukebox is a simple networking device, comprising of a Rasberry Pi connected to the Internet (e. g., through a 3G dongle), a WiFi dongle that leads those connected to a local captive portal, and an external hard disk. The assumption is that the Internet connection is very "thin" and expensive (e. g., a 3G connection with 500MB limit per month). The main difference of the Internet Jukebox, compared to other "Internet-in-a-box"-like technologies is that the content to be downloaded from its thin Internet connection is decided collaboratively between those that will consume it, unlike for example the case of the Outernet (https://www.outernet.is/en/).

The selection of the content happens in two separate (but possible parallel) phases: search and download. In the search phase, users search for content in the Internet (through a specific set of supported search engines, e. g., the Youtube API) and choose amongst the results the content that they would like to download, forming a list of desirable content to vote upon. In the download phase, users can place their votes until the "download deadline" (it could be one vote per person or a total "budget" per person in MB that can be distributed in different items). When the deadline arrives, the most popular content is downloaded and possibly projected in place, in the (typical) case when most users do not have a personal device. So, one could imagine for example the "movie of the month" application in a small village where users can only search for youtube videos through the Youtube API and perhaps using only a small provided set of shared devices. Every Sunday the most popular movie is downloaded and projected while previously downloaded movies can be consumed in private by those that do have a device.

5 Overall Summary, Conclusions, and Recommendations

This seminar was successful in creating awareness on the different barriers to digital inclusion, understanding the requirements, potential and the limits of solutions that have been proposed to address in this space. While there will continue to be a need for diversity in research approaches, we could identify some areas of synergy.

5.1 Infrastructure Deployment

It was clear when hearing about African and South American deployments in detail, that we could make valid comparisons with the experiences of European and North American Internet deployment in the past (going back over twenty to thirty years). There are regulatory barriers, which are highly reminiscent of problems in creating really widespread connectivity in the early 1990s between European countries, between regional networks in the US, and between Europe and the US. The barriers are not the same, but require similar thinking to alter the landscape more favourably. On the other hand, there were new problems to overcome, but also these can be seen as new opportunities to develop technology. In particular, very long range wireless communications techniques need much more work to deal with the massive distances and somewhat adversarial deployment environments in these areas (e. g., Amazon or Congo rainfall, foliage). At the same time, massive reduction in optical equipment costs, and increases in capacity available affordably on long haul fiber will also help (e. g., coming from some heroic astro-physics groups deployment for VLBI, just for one example). This technology would also address precisely the problems we see in sparsely populated fringes of the EU and North America, however, there are few market-led reasons to see technology developed for people there, whereas in Emerging economic regions such as Brazil, we could see leadership if the political obstacles are reduced.

The main opportunities here are in regulatory/policy work in designing solutions (not something we had more than modest expertise present in the seminar) and application of technical operational experience to fix problems, so that we can move other regions along the s-curve (logistical) faster, if possible (something we had some good examples of during the seminar).

5.2 Protocol Work

It is very clear that the Internet is suffering from ossification and obesity in the protocol world. The heterogeneity of available transport protocols and shims is somewhat overwhelming, as well as the prevalence of middle boxes within the net to try to improve performance in the presence of so-called impedance mis-matches between legs of an end-to-end path (e. g., at boundaries between cellular data net and backhaul).

The interactions between new protocol combinations are not well understood. (Number of TCP connections for a web download, adverts from 3rd parties, ECN deployment, SPDY, QUIC, Minion, HTTPS + MPTCP, or video streaming and DCCP or SCTP and some middle box functions, the list goes on and on). Luckily, there are pressures on many companies producing good engineering of protocols to improve things and the IETF has groups (e. g., TAPS, TCPM and others) with strong membership from people actually trying to fix the eco-system.

There is a good academic opportunity for much good empirical work to evaluate the solutions and document best practice.

5.3 Socio-Economic Research

It is entirely unclear what the social and economic benefits of Internet access actually are, especially outside of the main deployment areas. It is entirely clear that alternatives have had major successes socially and economically at far lower cost (use of cellular voice and SMS in many places), and that we suggest caution in making naive assumptions that the Internet automatically brings benefits that exceed the costs. Estimates have been made in highly developed regions (oft quoted UK number is 3 Billion pounds per year is saved by the government if the last 15% of the population get access), but even this figure is not really robust.

Empirical social and economic research is needed to understand the cost/benefit equation in each and every area. We had little expertise in the seminar, but there was consensus on this, that technogeek led heroic deployments of novel solutions often failed to sustain due to (for example) higher operational effort than their value justified. A study of failures would probably make a good academic research topic for groups that work on this (e. g., the Oxford Internet Institute, just for one).

5.4 Community Networks

As can be seen in the body of the report, we had very good representation from groups that had highly successful community network deployments. It is clear that these islands are growing, and starting to gain sufficient momentum to hit sustainability. Part of this due to their local, grass roots engagement, but also technologically, they have started to build management tools, and this should certainly help other groups. The main recommendation here is that this is a space to watch! To paraphrase William Gibson, the Street creates its own uses for the tech.

5.5 Internet-in-a-box

We also had a very good set of research groups working on these "extreme DTN" solutions to some forms of information access. Crucially, (despite the somewhat misleading name) many of the ideas seemed to offer somewhat different services than a vanilla Internet. Again, the recommendation here is that this is an area to watch. There are, as with community networks, novel technologies, but the bulk of the future work will entail optimising the sustainability of any deployment. Both Community Networks and Internet-in-a-box also require the Protocol Optimisation work that is ongoing in the IETF in any case.

6 Appendix

6.1 Individual communications tariffs and links to social services as regulated social and economic inclusion measures

Louis-Francois Pau (Erasmus University-Rotterdam, NL)

While network and access technologies offer diverse performances and operating costs, the question will always remain of how less favoured users can pay, and for society also how to use communications to embed them back into society with incomes over UN defined poverty line. A key issue has been to render such initiatives sustainable for operators or OTT operators as well, and to have technology-neutral solutions as required by policy makers. Our past and on-going work on mass customized communications tariffs (esp. mobile and Internet), regulated or corporate "good citizenship" driven individual tariff bundles, links between the above and public social measures for less favoured citizens have led to the adoption of social tariffs in some EU countries, with more to follow. A formal methodology has been proposed and adopted by regulators on affordable tariffs. We will also report of live trials of the same in China (agricultural communities), India, South Africa and Laos, and policy initiatives. New research questions arising from these interactions with users, NGOs, operators and regulators, and from the trials or deployments, will also be summarized.

References

1 L.-F. Pau, Affordability, Social eInclusion, and Addiction in Telecommunications, IEEE Technology and Society Magazine, Vol 31, no 3, Fall 20, 20–26+33
2 L.-F. Pau, Telecommunications in South Africa: who is actually regulating?, Journal of Governance and Regulation, Vol 2, no 3, 2013, 128–138 (ISSN: 2220-9352)

3 L.-F. Pau, Enabling mobile communications for the needy: affordability methodology, and approaches to requalify universal service measures, Informatica Economic? Journal, Vol 13, no 2, 2009, pp. 128–135
4 H. Chen, L.-F. Pau, with H. Chen, Mass Customization in Wireless Communication Services: Individual Services and Tariffs, in F.T. Piller & M.M. Tseng (Eds) , Handbook of research in mass customization and personalization (in 2 Volumes: Volume 1: Strategies and Concepts Volume 2: Applications and Cases), World Scientific, 2010
5 L.-F. Pau, J. Mottiwalla, J. Motiwalla, India: a case of fragile wireless service and technology adoption? (short version), Int. J. Mobile Communications, Vol. 6, No. 3, 2008, 376–389
6 H. Chen, L.-F. Pau, Individual telecommunications tariffs in Chinese communities, in: Huang, Wang and Day (Eds), Global mobile commerce: strategies, implementations and cases, IGI Global (formerly IDEA Group), September 2007, ISBN 978-1-59904-558-0

6.2 Challenges in developing regions

Weverton Cordeiro and Roderick Fanou

Communication Issues and Challenges in the Amazon Region. In certain places, local communities do have infrastructure for communication (fiber optics for instance), but it is often overpriced and lacks quality. An example is the Brazilian Amazon, a sparsely populated region with low GDP per capita. A fraction of the fiber optic network deployed there was initially meant for monitoring power transmission lines. After some time and political arrangements, that network was leased for commercial and public use.

In the city of Itaituba (PA), a small town (approx. 100,000 inhabitants) in the midst of the Amazon, fiber and satellite are available for data communication. However, the infrastructure in place cannot fully serve the needs of the local community. The current social situation is also not favourable. Approximately 40% of individuals (according to Brazilian IBGE) live below poverty line; they participate on social welfare programs, which improves their monthly income. In spite of financial constraints, in general they have smart phones (which are easier to afford due to instalment buying).

Because of the non-favourable social situation, a large fraction of the community cannot afford (good) connectivity services. For example, a satellite link costs $95 dollars for 2 Mbps, plus $380 dollars installation fee. Local providers are cheaper, e. g., $40 for 512 kbps, and $46 for 1 Mbps, plus installation fees (these costs may vary depending on the customer‘s location). The region experiences significant amounts of precipitation (rain) and cloudy weather yearly, and also very high air humidity; this scenario is often hostile for wireless communications there. The city also has four nationwide mobile operators offering data services. The dominant technology is 2G, but two of them also run 3G (deployed more recently); both services are offered with very poor coverage, though. There is no 4G service in the area. The mobile data charges are also expensive. As an example, one of the operators offer a $0.30 per day pay as you go plan, 10MB allowance, which is a bit expensive for local standards. In addition, voice and data services often black-out. Some mobile operators connect their antennas directly to satellites, possibly to avoid fiber as it fails quite often. Finally, there are public funded connectivity initiatives: Internet for free, for all, through public wireless hotspots. However, they have very low capacity (most of places have just one off-the-shelf, Access Point deployed) and fail quite often too (because of intense rain or fiber failure).

The local university campus is not an exception to this reality. The Federal Institute of Education, Science, and Technology of Para at Itaituba has two radio links of 4Mbps each. These links connect the campus to a state-run point of presence, and then to the fiber network. Those links could be of higher capacity, and in fact there are some arrangements under way for expansion. However, expansion plans are severely delayed due to financial constraints, and possibly because of lack of a proper political mindset. The radio links frequently disconnect (e. g., during rain, or when the fiber itself fails). The students used to have access to the local network in the past, but because of the high demand and severe resource constraints, they had to be excluded for the sake of professors and staff. In order to improve this situation, the local administration is deploying a backup link to satellite.

In addition to the aforementioned technical problems, the local community also experience political challenges too. For example, there is an ongoing project for building a high speed, optical network interconnecting several public and private institutes in the town. However, the project is now one year delayed because of political affairs. It basically means you now have to increase the OSI layer to 8 (financial) and 9 (political), and the local community as a whole seem to be stuck in 9.

In summary, having the necessary infrastructure for guaranteeing digital inclusion is often not sufficient to guarantee affordable connectivity services to poor communities. The deployment of technology appropriate for the local context, and a political mindset towards ensuring effective digital inclusion, are paramount to deliver good Internet services for the other billion.

We consider the town of Itaituba (Para) in Brazil as an example. It has a population of around 100,000, a Federal Institute for Education, Science and Technology with approximately 600 students. The Institute has an 8Mbps connection to unreliable fiber. Students have no access to the Internet in the campus, as the bandwidth is required for teachers and administrators. A wireless LAN with an "internet on a box" could provide many of the benefits of a campus network, without Internet access and, if designed with this in mind, could be easily connected to the Internet if and when a good fiber connection becomes available. The benefits could include web-based communication between staff and students, social networking tools, and access to content, educational and recreational, selected for its value to this community.

Access to this LAN could also be extended to the local population and businesses, and eventually to neighbouring communities. This would allow the students to become familiar with Internet technologies. We can also view this as a pilot. This same model could also be used in other communities where there is no Internet connection at present. Students from the Institute could be trained to contribute to such future projects.

Interconnecting Africa: Challenges. According to the world stats, Africa and Asia have the lowest penetration rates (21% and 32%) with about the half of the world population. More specifically, African countries have some common specificity, as far as Internet access is concerned:

- High cost for customers: due to low income, some people cannot afford access to Internet. For others living in rural areas, Internet is still a luxury, as these areas are not yet reached by technology, ISPs infrastructures or power.
- Poor QoS of Internet.
- Lack of Content: accessing to Internet is for many users, only for browsing, for accessing to Google, Youtube or Facebook. Their blogs or even their mails (private or not) are

hosted outside the continent on Yahoo, Gmail, Hotmail servers mostly for high availability purposes. Besides, people are not used to create local content even though they are knowledgeable: for instance, the lack of will and hardware prevent universities from hosting thesis, works, products or applications developed by their students while similar works are available online.
- Lack of economic incentives for ISPs to peer.
- It is difficult for ISPs to move content closer to people, as alone they do not meet the middleman‘s requirements.
- Education (from primary school to most universities) do not include familiarisation to computers.
- Most universities are connected to ISPs networks instead of NRENS.

In such context, the key milestones for a better Internet access in Africa can be listed as follows:
- The energy furniture.
- A climate of Fairness and cooperation established by regulators to secure and pave the way for huge investments
- Traffic Engineering, efficient routing by ISPs and Mobile operators, both aiming at keeping local traffic local to the continent. Besides, peering as much as possible and adding services (DNS root servers, cctlds, etc) to IXPs will definitively make ISPs save on transit costs.
- ISPs could then use these saved costs or interests to invest in building the physical infrastructure.
- Developing local content and content hosting to boost local economies. We need to stress that content developed in each country should have to be attractive enough and potential to be exported (education, culture, music, videos, activities specific to the country but well appreciated elsewhere) at least to other countries in its region.

As for the power, the continent is characterised by an abundant solar radiation through the year, abundant wind energy resources along coastal countries, an abundance of coal resources in Southern Africa, over 117 billion of proven oil reserves, overs 14.6 trillion cubic meters of proven gas reserves, etc. However, Energy is still unstable. Since, power is essential for industry and therefore for Internet access, Governments have to do their best to make it stable and sustainable. To this end, they could privilege and focus their efforts from short to long term on solar energy, renewable energy, gas and nuclear centrals.

The African terrestrial optical fiber map looks like the railroad one: very dense in South Africa, less in Eastern Africa and even less in the West. This reflects the level of interconnection among countries in which operate 1,159 ASes as of June 2014.

One way for its improvement is the bottom-up model. It consists for universities and NRENs (Ex: WACREN) to build per country an academic network linking primary schools, universities and hospitals. The existence of such networks would incentivize ISPs and Governments to invest in cross borders connections and Internet provision.

Another way is the up-bottom model which would need the regulations to facilitate fiber deployment, suppress the dominance by incumbents or the discrimination of the new companies as well as ensure price control of essential facilities. That is, regulators must ensure that all ISPs (private and incumbent) have the same rights on telecoms market. Mainly, they must get rid of monopoly (Djibouti, Ethiopia), enable both competition and partnership, encourage infrastructure sharing for the welfare of the citizens, make “crossing borders” easy for the enterprises, make license for ISPs or hosting companies declarative, etc.

Overall, the followings are key points for a better Internet in Africa:

- An affordable (cheap) International connectivity.
- Cross-borders interconnections and regional transit networks.
- Terrestrial optical fiber within regions and cities of countries.
- A Wireless access for the users.
- Content produced by these users (especially students in universities) available online.
- Datacenters connected to IXPs to host servers or government services (content produced locally, caches, CDNs).

Considerable efforts are being done throughout the continent to achieve these objectives. Among others, the AXIS (African Internet Exchange System) set up by the African Union (AU) and conducted by the Internet Society aimed at establishing IXPs per African country after organizing best-practises workshops as well as national and regional capacity building workshops. Liquid Telecom has deployed over the last 5 years its own optical fiber network covering Eastern and Southern Africa. Similarly, the Government of Rwanda has built optical fiber network covering all the 30 districts of its country, even in areas where there are no business reason yet. We can also list the GamersNight project of Liquid Telecom in collaboration with Uganda IX. Such efforts have to be encouraged and multiplied all over Africa to extend the network coverage and empower the digital development of the continent.

6.3 Solving the Geographical Challenges: Access Technologies for Rural and Remote Access

Arjuna Sathiaseelan, Mahesh Marina, Marco Zennaro, Fernando Ramos, Panayotis Antoniadis, and Karin Anna Hummel

Internet access in rural and remote areas constitutes a significant part of the problem of enabling Internet availability for all, in both developed and developing countries. In the developed world, unavailability of Internet access in rural areas is largely a situation of digital exclusion by default and not by choice, caused by low population densities in rural areas and higher per person cost of deployment of common broadband access technologies (ADSL, cable, fiber) per person. On the other hand, the problem is a lot more complex in developing regions involving socio-economic factors.

Broadly speaking, the nature of the rural access problem in a given region depends on several factors, including population density and distribution, remoteness, nearness to infrastructure that connects into the global Internet, income and education levels, cost and reliability of power infrastructure. Technologies employed for providing rural Internet access vary widely depending on the specific situation. They range all the way from fiber and copper based wired solutions to satellite and wireless technologies. Some of the newer types of wireless technologies (e. g., mesh, long distance WiFi, TV white spaces) are discussed later in this report. Also, in many cases where conventional wired Internet access solutions are not economically viable, the gaps are often filled by network infrastructures involving participation of communities. Therefore, various types of community networks (e. g., Guifi in Spain, Tegola in Scotland, Cybermoor in Cumbria, UK) and their characteristics are also discussed separately in this report.

In many developing countries, significant a percentage of population live in rural areas. For example, two-thirds of people in India live in rural areas. These environments are typically economically deprived and also suffer from poor infrastructure (power and transport).

In Europe and North America, the scale of the rural access problem varies anywhere from 5%–30%. In the UK, a tenth of the population do not have access to basic 2Mbps connections and a third lack access to connections with speeds greater than 30Mbps (what is now being considered superfast broadband). Situation in some parts of the UK is worse. For example, only 9% of the households in rural areas have access to faster connections although studies on data use over broadband connections by Ofcom suggest that rural users would equally benefit from high speed Internet access.

Some lessons and challenges for enabling broadband for all in rural parts of developed countries are outlined here. While these are largely based on the experience from the Tegola rural wireless access project in Scotland , they may hold for other developed country situations as well. Two main policy related challenges for addressing the rural access issue can be identified.

1. Firstly, the models for government interventions to extend broadband reach in rural areas are as important as interventions themselves. In particular, for rural areas, promoting and enabling bottom-up community driven efforts are more likely to succeed as opposed to top-down approach via procurement.
2. Secondly, accessible and affordable backhaul (and middle mile) remains a challenge. Where there is publicly subsidized fiber infrastructure being rolled out, it should be made openly accessible for local ISPs and communities. Also wider use of the public services network infrastructure (e. g., serving schools and hospitals) would be an effective method to intervene in areas where there is no alternative infrastructure.
3. From a technical perspective, more detailed understanding of wireless technologies suitable for addressing the middle mile / backhaul problem (e. g., 24GHz) would be useful for DIY efforts.

Having mobile voice and data access widely available is important in view of growing adoption of smartphones and especially in areas where alternative forms of communication are non-existent or unreliable (e. g., for emergency services). A key challenge in this context is quantifying the gaps and quality of mobile services. While modelling based 2G/3G/4G coverage maps from regulators and operators might suggest that mobile services are ubiquitously available, user experience differs greatly from modelling based coverage estimations. For example, Ofcom‘s modelling based coverage analysis focusing on outdoors indicates that there is less than 4% of UK premises lacking in 2G/3G coverage whereas consumer experience shows that the actual scale of the problem is much larger around 20%. Therefore it is crucial to assess mobile coverage and quality via measurements, which however is a technically challenging problem. For instance, context needs to be accounted for in user-side measurement of mobile performance. As another example, rural areas may need a combination of network-side management information and user side measurements to identify mobile not-spots. Measurement of mobile performance is also crucial from the perspective of enforcing the coverage obligations associated with spectrum licenses held by mobile operators. Current enforcement mechanisms rely solely on modelling and simulations, which have the same problem, mentioned above. Also in some countries like the UK, newer coverage obligations for 4G require only one operator to provide near universal coverage, which leads to lack of market for mobile services in rural and underserved areas.

In addition to the above, there is a significant value in documenting (or at least highlight the importance of documenting) the current experience and lessons on several important aspects like sustainability models for community efforts and technology adoption models in developing and underserved areas.

Role of long distance Wireless. The long distance Wi-FI (easily 30 km) are a low cost opportunity for "low bit rate" connectivity of villages (or parts thereof) needing the some connectivity. 802.11 Wi-Fi technology is commonly used for creating wireless networks with a range of about one hundred meters. With careful planning and proper antennas, this same equipment can be used to make point-to-point links of dozens of kilometers.

The two major limitations for using Wi-Fi over long distances are the requirement for line of sight between the endpoints and the vulnerability to interference in the unlicensed band. The first limitation can often be addressed by taking advantage of the terrain elevations, or by using towers to overcome obstacles and provide Fresnel zone clearance. As the distance between sites increases, the curvature of the Earth becomes a serious obstacle, requiring higher elevation at both ends. The second limitation is less pronounced in rural areas, and can be alleviated by migrating to the less crowded 5 GHz band. Low cost equipment supporting TDMA and equipped with high gain directional antennas allow medium to long distance links to be deployed with a very small investment.

Utilising higher frequencies. Low cost equipment of higher frequency bands are now becoming available and can be used on unlicensed bands. For example, the 17 GHz frequency is now unlicensed in Europe, with a maximum Equivalent Isotropic Radiated Power (EIRP) of 20 dBm. 2 4GHz is an unlicensed frequency that can be used for microwave communication for point-to-point wireless backhaul. This is subject to local country regulations. These frequency ranges are subject to rain fade, meaning that atmospheric rain absorbs microwave radio frequency (RF) signal, thus causing a signal loss.

High frequency point-to-point links on these frequency bands may be installed to service locations many kilometers farther than could be served with a single link requiring 99.99% uptime over the course of one year. A secondary lower bandwidth link such as a 5.8 GHz may be installed parallel to the primary link, with routers on both ends controlling automatic failover when the primary link is down due to rain fade. Recent equipment on 24 GHz can deliver wireless throughput of up to 1.4+ Gbps, surpassing conventional wired backhauls, over real-world, 10+ km links.

A lot of effort is devoted to the development of Transmission on 60 GHz. The 802.11 ad is only the beginning of the "wave of development". Using multiple antennas (24+) on a very small surface point-to-point wireless links for distances in the range of 300 Meters can be expected in a few years. With price and reliability comparable to the today's WLAN access points. This might be a game changing technology as for "connecting WLANs from individual houses" into local structures.

TV White Space as a solution. The growing demand for wireless data transmission drives the search for alternatives to the current spectrum management schemes. In the long term, the only viable solution seems to be dynamic spectrum access once the technical details for its implementation are solved. In the short term, the use of currently vacant spectrum allocated to TV broadcast (so-called TV white spaces or TVWS) can alleviate the spectrum crunch while opening the path for dynamic spectrum access. Several measurements campaigns have shown that the TV broadcasting spectrum is mostly unused in sparsely populated areas, especially in developing countries, as there is not enough return on investment for broadcasters to provide many simultaneous TV channels. For the same reasons, these are precisely the areas in which Internet access is frequently lacking.

TV White Spaces technology can take advantage of the improved propagation capabilities of these frequencies to provide affordable Internet access in rural areas. White spaces are also present in densely populated areas as a consequence of the transition from analog to digital TV, and these can be harnessed for wireless Internet access as well as other wireless communication services. The lower frequencies as compared with the ones used for WiFi (which in some places is becoming too crowded), are less attenuated by walls and offer an interesting alternative also for indoor Internet access, as well as for multimedia distribution. Finally, for machine-to-machine (M2M) applications and the "Internet of Things" paradigm TVWS technology has significant advantages both for developed and developing economies.

A quest towards understanding the potential of TVWS has been initiated by way of experimental measurement campaigns that strive to establish the nature and extent of spectrum usage and the resulting TVWS availability. It is imperative that long term occupancy measurement campaigns and analysis studies be carried out so as to ensure that government bodies, research and development agencies, and other interested parties target the real and evolving spectrum situation.

Satellite Broadband. Satellite systems have traditionally used Ku-Band operating at 11-14 GHz, and were often limited to rates less than needed for broadband services. Satellite systems can provide access to large geographical areas and are not constrained by the need to provide backhaul capacity for each location.

A new generation of Ka-Band systems, operating at 20-30 GHz have largely removed the previous constraints, enabling much higher speeds and reducing the size of the antenna needed by a customer. This new generation is now seen as a key enabling technology to deploy broadband access to locations that cannot be cost-effectively reached by other technologies.

Broadband satellite support forward link rates of up to 80 Mbps per forward link with typical user rates offered 2-20 Mbps per terminal and return link rates of 256 kbps – 4 Mbps currently common.

Powerline Communications. According to public data (e. g., from the World Bank), there are more electricity lines than telephone lines worldwide. This difference is particularly significant (an order of magnitude at least) in low-income countries. As such, in these countries power line communications (PLC) are arguably the only technology with a deployment cost comparable to wireless. This makes it a technological solution that may be worth exploring for Internet access.

Power grids today are being modernized through the introduction of communication networks dedicated to the management of energy transmission and distribution. Although the investment in smart grids today is concentrated in developed countries, smart grids may play an important role in the deployment of new electricity infrastructure in developing countries, by enabling more efficient operation and lower costs. For instance, in sparsely populated areas smart grids could enable a transition from one-off approaches to electrification (battery- or renewable energy-based electrification) to the development of community grids that then connect to the national grid.

Can the grid be part of the solution for the challenge of global access to the Internet? Some opportunities seem to exist. For example, exploring the existing power lines for Internet access using PLC; or leveraging on the distributed, ad hoc nature of community smart grid networks to explore novel low-cost access solutions. A peer-to-peer network called Gridmates already lets those with excess energy share it with neighbours in need, so one can imagine similar sharing of Internet access using the power line.

But many challenges exist, and indeed the discussion that ensued has led to the conclusion that this solution entails enormous challenges that may or may not preclude its adoption.

These challenges include:
1. problems of noise and interference that are difficult to solve. For example, PLC turns electrical wires into antennas effectively, and this can be problematic. PLC could hamper the reception of broadcast shortwave radio, for instance. More modern protocols, however, seem to include techniques to mitigate this problem;
2. the power infrastructure in low income countries may be really poor, making it extremely difficult (if not impossible) to use the higher frequency bands using the existing power lines;
3. it is important to make realistic calculations of the bandwidth per household, since the last mile is shared, and the backhaul connection can be very limited in these countries;
4. another point worth investigating is the energy cost of providing internet access;
5. also, is the overall monetary cost of this type of solutions reasonable, when compared to the alternatives?

DTN/Store-carry-forward for developing regions. In the last decade, opportunistic and delay tolerant networking (DTN) research has resulted in a set of advanced concepts and technologies that are ready to be employed in particular use cases, such as Twimight or WLAN-Opp. Among the most attractive application fields of opportunistic networks are empowering people when disasters strike, circumventing censorship in repressive regimes, and connecting regions where no network infrastructure is available in developing countries. One example of employment of DTN technology is the recent democratic movement in Hongkong, 2014 that made use of a simple DTN-enabled app, namely FireChat.

Although community networks demonstrate that volunteers provide and maintain resources deliberately, it is yet unclear how a solution incorporating mobile carriers will be actually accepted and which capacities one can expect from such networks. A potential system architecture that makes also use of opportunistic store-carry-forward transmission may consist of stationary boxes with storage facilities typically installed at popular places and mobile data carriers. Based on CDR (call detail record) data of cellular network operators, mobility flows, e. g., from city centres to suburban areas can be estimated (e. g., making use of the D4D challenge data of Orange in Ivory Coast) and used to characterize the capacity of a hypothetical DTN network. Although cellular data can provide countrywide mobility flows, these data include only the subset of commuters (those, that are calling) of one operator, which comes at a bias of unclear importance. The capacity formulation considers also the storage on devices, success of transmission, the user's willingness to cooperate, and a scaling factor. An open research issue is how to find good models and configuration ranges of theses factors based on social, behaviour, and demographic observation. Questionnaires used to capture the opinion of inhabitants in developing region can provide insights, though the way of asking is important. A general question about whether people would volunteer carrying data of others resulted in a fraction of 8% agreeing if rewarded, 15% agreeing to carry for friends and relatives, and a majority of 77% who would not do it at all. Yet, putting this question into concrete context might result in higher acceptance of the people as they would see concrete application dependent benefits for themselves and their social community.

6.4 Community Networks

Leandro Navarro, Renato Lo Cigno, Leonardo Maccari, and Roger Baigl

Traditionally networking infrastructures and services were provided by the ecosystem of telecom monopolies, mobile and fixed network operators, telecom service providers, complemented with manufacturers of carrier-grade equipment, terminals, etc. However these are driven by business interest and cannot provide networking services everywhere, where there's no market yet, and to everyone in the world, who may not cover the required investment with service fees or generate enough economic profit.

Community networks (CNs) are a bottom up effort to build independent, community-owned network infrastructures. They serve two roles, first they are used to replace the last mile and distribute Internet connectivity, second, they provide services to the local community with a governance model centred on the community and the stakeholders, on an infrastructure owned and managed by the community. They are created starting from principles of neutrality, non-discriminatory access, collective ownership and participated management, and are growing in many places around the world.

Cheaper and simpler to setup network components, combined with open spectrum for WiFi, open knowledge, and open source software, has enabled individuals and small communities to create their own network infrastructures and services without the need of large companies with large investments and complex structures. This model expands the opportunities and empowers citizens to be in the digital society.

This opportunity has enabled new forms of networking service from a local perspective. Individuals, small organisations, and their neighbours can cooperate to create and operate a network infrastructure and services locally. This typically results in creating a local ecosystem open for participation.

Probably for the first time, the design of a communication system can start from the needs of people, from urban evolution and planning, and not be the outcome of technology or business analysis (that failed so many times in the last decades)

Some CNs (one for all: Guifi counts 26000 nodes and about 80,000 users. Originally based in Catalonia is now expanding to all Spain), reached a size that is not anymore governable with only peer-to-peer interaction and independent node management. CNs are actually proposing new models of cooperative, bottom-up organization to face the costs and the challenges of becoming almost a nation-wide network, and are developing the required tools to put them in practice.

CNs should not be considered only as "cheap internet access". A challenge they take is how to reshape the communications infrastructure in order to be more fair, more decentralized and based on the principles of neutrality and community empowerment. Nevertheless, the networks that are effectively in place do not really reflect this organization, they are topologically fragile [1], with few nodes that can be routing the majority of the traffic, and socially the communities risk to be dominated by few individuals. There is a need for instruments (not only technical but also social, economical and legal) to help the community grow in a sustainable way [2]. Guifi is in the process of creating some of these instruments, that are freely accessible to anyone that want to replicate the approach, among these are: software, licenses, conflict resolution procedures, etc.

Community Networks have developed around the world in very diverse environments to enable local networking infrastructures that allow every citizen to join in and participate in the digital society. However, while being quite successful in diverse locations in providing local digital network infrastructures and services, there are several challenges, obstacles and lessons learned in its way.

The Community-Lab testbed. Community-Lab is an open, distributed experimental infrastructure or testbed with around 200 small computer nodes. It is embedded in several community networks with more than 40,000 routers and users spread in diverse regions of Europe, using mostly wireless but also a growing usage of fibre links. Researchers can deploy experimental services, perform technical and social experiments or access open data traces in Community-Lab. The testbed follows the slice architecture that PlanetLab started and thus allows multiple experiments to share experimental nodes together with production traffic and community network users. Community-Lab is designed and has been used in experiments ranging from the link layer to routing, transport and application layer, or even social experiments.

Community-Lab consists of at least one portal or controller and a set of nodes that are embedded in different community networks. Each node consists of several slivers (Linux containers) which are grouped into slices that represent an experiment or service. The purpose of the Community-Lab controller is to manage and control the testbed, i. e., manage its users, nodes, slices and slivers. Currently there are about 25 research and community organisations involved in the testbed in Europe, Africa and America. Research groups already working with local community networks are encouraged to join Community-Lab http://community-lab.net.

Challenges, Research and Development. While community networks are widespread, numerous challenges still lie ahead. Some closely resemble "classical" networking, software systems or collective social organisational challenges, while others are very specific to community networks.

The free and open source software development community, working with community networks and the networking and software systems research has developed and enhanced key enabling technologies like the OpenWRT Operating System for embedded devices or the OLSR mesh routing protocol among many others.

Among of the most relevant research results in the last 3 years have been the development of the OLSRv2 mesh routing protocol, the development of hybrid nodes with network attached radio interfaces (DLEP), the optimization of the IPv6-based BMX6 routing protocol with multi-topology extensions, comparative analysis of the topologic properties of diverse community network graphs and its implications.

Collaboration between the FOSS community and community networks has resulted in the development of generic firmware for bootstrapping a community network such as the Quick Mesh Project (qMp) http://qmp.cat/, Libre-Mesh http://libre-mesh.org/ router software distributions (including BMX6) and the NodeDB to manage a set of nodes involved in a community network.

The FOSS software developed for the Community-Lab testbed has been extended to be a resource infrastructure manager for a set of hosts embedded in a community network that can provide distributed resources for experiments or for community services. These nodes include a remote configuration, management and monitoring daemon that, combined with central components, allow managing all of them in unison. Together with the Cloudy distribution developed in the Clommunity project by Guifi.net, and combined with experimental peer-to-peer platform services for video distribution and distributed storage, Community Cloud services have been demonstrated in Guifi.net.

Several open challenges: Scalable and dynamic mechanisms for the interrelated, dynamic and global problem of the allocation of resources such as spectrum; Economic and sustainability models; Collaborative sensing; Local community cloud services and uses to provide digital services to citizens, digital tools for collective awareness, collective action, and social innovation; Socio-economic aspects; Multidisciplinary methods of evaluating community networks as well as their "performance" and sustainability; Self-organising systems; Production of open data sets.

These experiences are being shared in informal and formal interactions in events such as the International Summit for Community Wireless Networks (IS4CWN), the International Workshop on Community Networks and Bottom-up-Broadband (CNBuB), the BattleMesh workshop, or the IRTF Global Access to the Internet for All (GAIA) group.

References

1 L. Maccari and R. Lo Cigno, A week in the life of three large wireless community networks, To appear in Elsevier Ad-Hoc Networks, 2014. Online: http://www.sciencedirect.com/science/article/pii/S1570870514001474

2 R. Lo Cigno, and L. Maccari. "Urban wireless community networks: challenges and solutions for smart city communications." ACM international Workshop on Wireless and Mobile Technologies for Smart Cities (WiMobCity), Philadelphia, PA, US. 2014, DOI: 10.1145/2633661.2633669

6.5 Key Parameters for Universal Internet Access: Availability, Affordability and Relevance

Henning G. Schulzrinne

Internet adoption tends to progress in phases: early rapid adoption in urban areas and by relatively well-educated and higher-income households, then transitioning to a slower pace once between 70% and 85% of households subscribe. For example, in a 2013 survey, the Pew Research Internet Project found that approximately 15% of the adult population does not use the Internet; those least likely to use the Internet include "senior citizens, adults with less than a high-school education and those living in households earning less than $30,000 per year [1]."

Non-adoption has multiple reasons that vary in each country, and, within each country, by demographics. These reasons are often grouped, somewhat crudely, as availability, affordability and relevance. For the United States, a 2014 NTIA survey[1] (Figure 1) identified 48% as stating "Don‘t need. Not interested", often summarized as "relevance", as the main reason for not using the Internet at home, followed by "too expensive" (29%), "no computer or computer inadequate" (11%). Only 1% each claimed that there was no Internet access available or that privacy or security concerns dominated. A small percentage (3%) use the Internet somewhere else, e. g., at work, in coffee shops or a library. It is likely that many individuals will have multiple reasons and some non-users increasingly use functions on

[1] http://www.ntia.doc.gov/files/ntia/publications/exploring_the_digital_nation_embracing_the_mobile_internet_10162014.pdf

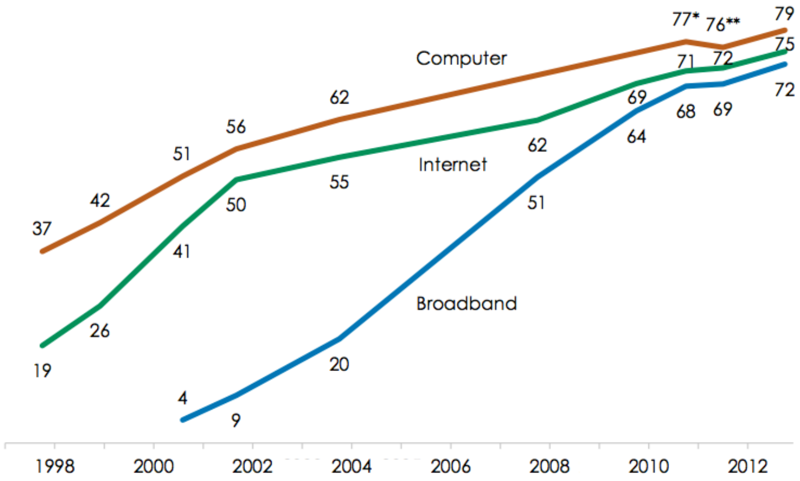

Figure 1 Overview of household adoption rates by technology, Percent of US households, 1997–2012 (Source: NTIA 2014).

smartphones that may require Internet access, e. g., for maps or games, that may not be appear to the user as such.

While the emphasis is often on private use, either at home or mobile, availability, affordability and relevance play a role for use by community anchor institutions such as schools, libraries, health care facilities and local government offices, as well as for small businesses, whether retail, hospitality or small manufacturing. For example, a recent report "finds significant gaps in the availability of high speed broadband among traditionally underserved students and their more affluent, suburban, and White peers' [2]'.

Availability, affordability and relevance are in turn influenced by a number of external factors, such as industry structure, the regulatory environment, consumer behaviour overall, and available technologies and speed tiers. For example, the level of competition is likely to influence both the price consumers' pay for access, i. e., affordability, and the investment made in new high-speed access and middle-mile technologies. In turn, the regulatory environment strongly influences industry structure. Consumer behaviour, such as the shift from landline to mobile phones or from linear video to streaming video-on-demand, creates opportunities for existing and new entrants.

What Kind of Internet? When we discuss Internet access, it is helpful to distinguish multiple tiers, with very different capabilities (Figure 2). For example, basic messaging, whether via SMS or low-bandwidth IP-based technologies, can already satisfy key information and transaction needs, whether for electronic payments (M-PESA), person-to-person communication or for information retrieval about weather or medical information. Such services may only require access bandwidths of the order of 10 kb/s, such as those enabled by 2G wireless or dial-up, and are likely to consume at most a few MB of data a month. The next step up enables text-heavy web content as well as text email, and is easily achievable through (reliable) 3G coverage or basic ADSL. With bandwidths of a few Mb/s and monthly bandwidth-allowances of around 1 GB, users can access most main-stream web sites, even those not optimized for low-bandwidth devices. At speeds of roughly 10 Mb/s and monthly data usage around 10 GB, the home Internet connection can offer short-form video (e. g., YouTube) and replace FM radio. Between 10 and 100 Mb/s, with data budgets of at least 100 GB, streaming video-on-demand becomes plausible, while, longer term, a full TV replacement

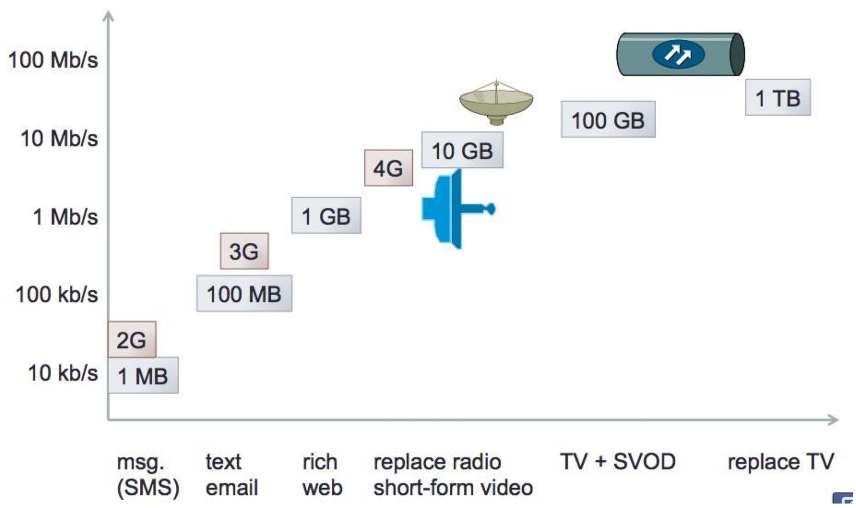

Figure 2 What kind of Internet do we want?

at 4K quality is likely to push the monthly data volume per household to 1 TB and above.

Beyond speed and monthly data volume, there are at least three other dimensions along which one might characterize Internet access: latency, symmetry and uniformity. Latency, typically expressed as round trip time, may vary between 20 ms to the edge of the access provider for fiber and LTE to more than one second for satellite links. Dog-leg routing of traffic through far-away interexchange points may further increase the effective latency. Latency reduces the usability of interactive applications such as phone calls, telemedicine, distance learning and some games, increases the load time of web pages referencing network-based APIs, and may reduce the achievable throughput of bulk data applications. Traditionally, most residential access links, such as ADSL and HFC, have been highly asymmetric, with download speeds ten times higher or more than upload speeds. Creating content requires higher upload speeds, and thus more symmetric networks. Finally, some networks for challenging environments may offer non-uniform access to content[2], where some content, cached locally, is available immediately and at high bandwidth, while other content may only be available after significant delay measured in minutes or hours, e. g., after fetching it during low-usage overnight hours. Similar ideas are also implemented by modern entertainment systems on commercial aircraft, which host high-bandwidth content locally, supplementing a lower-bandwidth Internet access channel. Some of the "Internet in the box" projects discussed in the workshop report occupy spots on this continuum, with the ideas exploring aspects of related to delay-tolerant and disruption-tolerant networks, or more classical web proxy caching. (For example, the ICOW work [3] places caches on public transit vehicles. Earlier, we had explored local applications, such as bulletin-board systems [4].)

Centralized vs. Distributed Infrastructures. All societies above a certain level of development depend on four key infrastructures: potable water and waste water, energy such as electricity and natural gas, transportation and communication (Figure 3). In each case, albeit to differing degrees, these infrastructure services can be provided with various degrees of centralization or local and individual autonomy. Recently, the delivery of electrical energy

[2] The allusion to NUMA concepts in computer architecture is intentional.

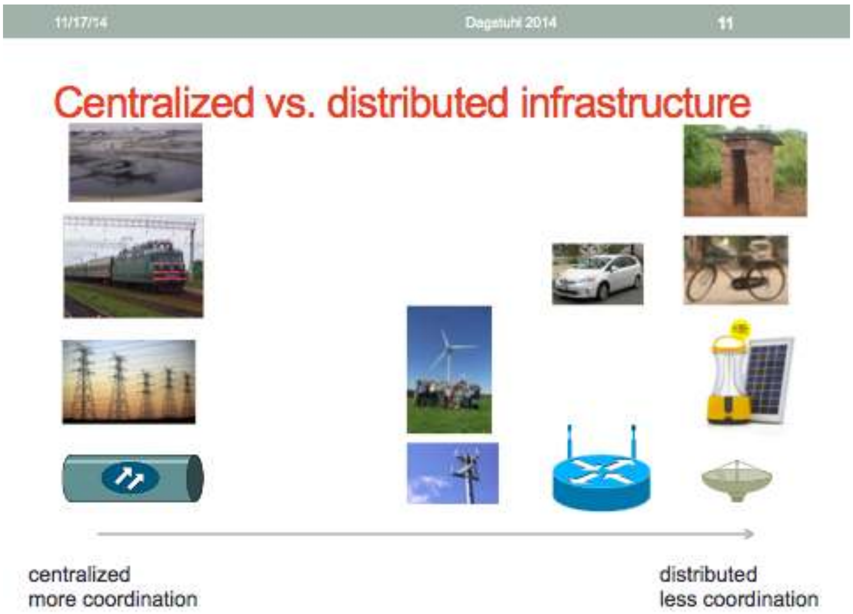

Figure 3 Centralised vs. Distributed Infrastructures.

has probably been most visible in exploring these trade-offs, whether through solar powered lanterns and cell towers on one end of the spectrum, or huge wind farms at the other. For communication, a purely local infrastructure seems to have inherently lower value than connecting to the global Internet, even if such purely local infrastructure can clearly be "better than nothing". Even for communication, there are trade-offs of having infrastructure such as satellite that requires very little coordination among participants and very little on-the-ground investment by local governments vs. a fiber infrastructure that requires the consent and participation of thousands of property owners and local governments, for example.

What's Expensive About Networks? The network engineering and research community has spent most of its efforts on making networking equipment more efficient and cost-effective (Figure 4). However, for most network build-outs, the electronics, optical components and fiber contribute only a small fraction, often less than 30%, of the total cost. For example, a fiber cable with 48 strands might cost $5,000 per mile, while the finished cost of a middle-mile deployment may range between $50,000 and $70,000 per mile. Cables with more fiber strand are proportionally cheaper; for example, a 144-strand cable costs roughly $10,000 per mile.

The cost of deploying fiber is driven first by the number of homes passed, regardless of how many of those homes subscribe, and then the actual number of subscribers. Depending on how networks are built out, the second cost component is only incurred when the carrier or community has a paying customer. For typical densities in urban and sub-urban environments, the cost seems to be around $1,000 per home passed for recent deployments. The per-home optical network termination units (ONUs) used to be relatively expensive, at $300-$400 per home, but have been dropping to around $100 to $200. The costs are generally higher if ducts have to be placed underground, compared to stringing fiber on utility poles, but duct or other underground burial tends to yield a more reliable network. Actual historic FTTH costs range between $1,000 and $1,300 per home, for outside plant only.

Independent 2" Conduit Run for Three User Co-Location

LABOR						
Category	Quantity	Unit	Low Cost/Unit	High Cost/Unit	Low Cost	High Cost
Design	5,280	FT.	$0.08	$0.10	$422	$528
Engineering and Permits	0	FT.	$0.25	$0.25	$0	$0
Railroad Crossing	0	LOT	$5,000.00	$15,000.00	$0	$0
Directional Boring for 2" Conduit	0	FT.	$8.00	$20.00	$0	$0
Directional Boring for 4" Conduit	0	FT.	$11.00	$25.00	$0	$0
Trenching for 24" - 36" Depth	5,280	FT.	$5.00	$12.00	$26,400	$63,360
Place Conduit	15,840	FT.	$1.00	$1.75	$15,840	$27,720
Place Inner Duct	0	FT.	$0.50	$1.50	$0	$0
Place Vault	33	EACH	$500.00	$750.00	$16,500	$24,750
Place Fiber in Conduit	15,840	FT.	$1.25	$2.50	$19,800	$39,600
Install Splice Enclosure	3	EACH	$300.00	$500.00	$900	$1,500
Splice Fiber	648	EACH	$12.00	$30.00	$7,776	$19,440
TOTAL LABOR					**$87,638**	**$176,898**
MATERIALS						
Category	Quantity	Unit	Low Cost/Unit	High Cost/Unit	Low Cost	High Cost
216 Count Fiber	18,216	FT.	$1.80	$2.50	$32,789	$45,540
Splice Kit	3	EACH	$500.00	$750.00	$1,500	$2,250
4" Conduit and Materials	0	FT.	$2.98	$3.50	$0	$0
2" Conduit and Materials	15,840	FT.	$0.88	$1.50	$13,939	$23,760
1" Inner Duct	0	FT.	$0.30	$45.00	$0	$0
Vault	33	EACH	$450.00	$600.00	$14,850	$19,800
Tax and Freight	1	LOT	$6,307.80	$9,135.00	$6,308	$9,135
TOTAL MATERIAL					**$69,386**	**$100,485**

Figure 4 Brief Engineering Assessment: Efficiencies available through simultaneous construction and co-location of communications conduit and fiber, (Source CTC, http://www.ctcnet.us/, 2009).

To first approximation, it appears that the cost per home is closely proportional to the distance between homes within one network deployment. Thus, for the same population density, serving homes spaced one-by-one along rural highways may be much more expensive than serving small villages where homes are clustered.

The figure below shows an example calculation of a fiber build-out [5]. Other cost examples include a build-out with 96% aerial and 4% buried construction, in the Northeastern United states, with a cost of $30,000 per mile. Aerial overlash can be cheaper at $15,000 per mile, while buried fiber can cost $89,000 per mile.

It is currently unknown whether combinations of fiber to the neighborhood (FTTN) and VDSL, such as G.Fast, or wireless, whether 5.8 or possibly 60 GHz, can offer a long-term attractive alternative. They may trade-off an initially lower investment with higher operational costs or lower reliability, as well as practical difficulties such as where to place DSLAM and radio components. A more systematic exploration of this trade-off, with real-world cost figures, would be enlightening.

Current field engineering wisdom seems to have converged on a number of recommendations for reducing the cost of building new “greenfield” or overbuilt networks:

- To reduce the cost and difficulty of deployment, active network elements are placed in as few locations as possible. This avoids negotiating for “fiber huts” with local communities or finding in-building space that is accessible for maintenance. Recently, even passive optical networks place the splitter in the head-end and create fiber home runs.
- To reduce the cost of dispatching technicians, all per-home fiber drops and other installa-

Company	Revenue	Capital expenditures	%
Comcast (US), 3Q14	$11.04B	$1.644B	14.9
Telekom (DE), 3Q14	€15.6B	€2.58B	16.5
Safarikom (KE), H1FY15	Ksh 79.34B	KSh 12.37	15.5

Figure 5 Expenditures.

tion is performed at once, e. g., in one Google "fiberhood".

- To reduce the cost of in-home installation, each home is connected through an all-in-one optical network unit that includes, for example, a high-speed 802.11ac wireless interface that can directly provide network connectivity to the whole home. (Early Verizon FiOS deployments had the ONU in the basement, which is then connected via coaxial cable and MoCa data transmission to the set top box. The set top box is placed near the family TV and also contains a Wi-Fi interface. This requires in-home cabling.)
- Any active in-network components are reverse-powered through the home, to avoid having to connect to utility power.
- To the extent possible, homeowners are asked to self-install, as had happened earlier with DSL and cable modems. If ONUs become standardized, it may also be possible for the carrier to ask the consumer to purchase the device, possibly using the Equipment Installation Plan (EIP) model now popular in the United States for smartphones, further reducing capital outlays for the carrier.
- Communities and other government entities responsible for roads, railways and pipelines should, by default, install at least fiber ducts or dark fiber whenever a road, railroad or pipeline is built or major improvements are made.

Given the large cost differential between the cost of fiber and the total construction cost, there seems to be room for significant improvement on the civil engineering side. One can imagine autonomous underground boring vehicles, possibly leveraging advances in horizontal drilling pioneered by oil shale fracking, that place fiber or ducts with minimal human intervention.

While the capital investment for fiber or other network access technologies is substantial, it is a small fraction of the total cost of operating networks. Below, we show three examples that illustrate that, across three continents, the capital costs are roughly 15% to 16% of the revenue (Figure 5).

These capital expenditures include costs other than for building networks. For example, Figure 6 shows the Q2 2014 capital spending for Comcast, the largest US cable communications company. About half the spending is for customer premises equipment, primarily video set top boxes.

This observation indicates that attempts to reduce the cost of network operations may be more productive than reducing capital expenditures. For example, self-configuring and self-diagnosing networks could reduce the number of field service calls and customer support incidents. Simpler billing models may reduce the need for extensive billing-related expenses. In general, fiber networks have significantly lower maintenance costs than copper or coax

Comcast's Q2 2014 Capital Spending Trends

Category	Growth CapEx ($ mil.)	% of Total (%)	Maintenance CapEx ($ mil.)	% of Total (%)	Total CapEx* ($ mil.)	% of Total (%)
Consumer Premises Equipment	668	65	72	16	740	50
Network Infrastructure	107	10	287	64	394	27
Support Capital	48	5	89	20	137	9
Commercial	209	20	0	0	209	14
Total*	**1,032**	-	**448**	-	**1,480**	-

As of Aug. 2014. * Total excludes $13 million in discretionary capital. Total including discretionary spending was $1,493 mil. Source: Comcast. © 2014 SNL Kagan, a division of SNL Financial LC estimates. All rights reserved.

Figure 6 Comcast spending trends (Source: SNL Kagan, "Comcast, TWC drive spending as Q2 cable CapEx spikes 27.7% YoY", August 2014).

networks. One of the gaping holes for engineering solutions that address these problems is the apparent lack of any public data on the cost of operations in real, particularly large-scale, networks. Community networks and networks run by public organizations may offer opportunities to gain insight into real-world operational costs. Without precise cost models, it is impossible to say, for example, whether and where replacing FTTP networks with FTTC plus wireless models is cost effective.

The Cost of Carrying Bits. In the figure below (Figure 7), we try to capture, grossly simplified, the three principal components of network costs, namely the Internet backbone, middle mile and last mile, i. e., access. The relative impact of each cost factor depends on the size and geographic scope of the network operator. For example, a small rural network operator may have very limited competitive choices among middle mile providers that can carry its traffic to the nearest interexchange point (IXP). This motivated the United States BTOP program to fund middle mile access, e. g., in Vermont. On the other hand, larger operators with a significant footprint can probably justify either leasing dark fiber or building their own, so that the cost of the middle mile infrastructure is more modest as it can be amortized over large traffic volumes.

As the figure tries to illustrate, the volume-dependent cost also differs significantly among the three components. For example, the middle mile may only require additional electronic components to increase bandwidth, while HFC networks may need nodes to be split. If the last mile is based on fiber, there is essentially no incremental cost of carrying additional bits on that last mile. (This does not necessarily argue for flat pricing; it has been argued that charging based on data consumption reflects the value users place on network connectivity.)

As video has come to dominate Internet traffic, much more of the content can be cached close to the edge, either at the IXP or even at the fiber head end. Thus, additional video volume may only impact the volume-insensitive components of the network.

In many parts of North America and Europe, the cost of Internet transit has decreased steadily, as shown in the figure below (Figure 8)[3]. The cost is shown in dollars per month and Mb/s. (Transit capacity is typically sold at the 95th percentile of 5-minute traffic intervals.) In addition, many large "eye ball" networks, i. e., networks serving consumers,

[3] http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php

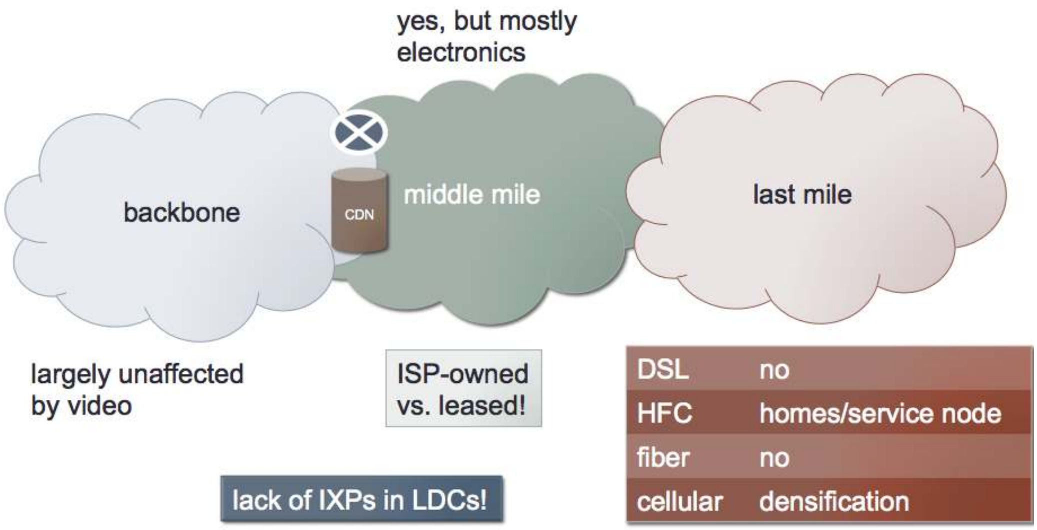

Figure 7 Network costs.

peer settlement-free, i. e., without charge, with other networks, or, in some cases, even charge content distribution networks (CDNs).

Overall, the cost of delivering bits across networks has now, in the United States at least, dropped below the physical cost of delivering bits on disk or DVD. As an example, the postage cost of shipping a Netflix DVD round-trip is approximately $0.70, translating into $100 per TB of data. For user-created data, the DVD-R cost itself is about $0.25, or roughly $30 per TB. Typical CDN pricing is $7 to $20 per TB.[4]

How to Finance Networks? As noted earlier, building new network infrastructure requires an initial substantial outlay of capital. Since the take rate is often not known ahead of time, an investor, whether for-profit or community, incurs substantial uncertainty how much revenue they can expect. The uncertainty grows if there is a lower-performance competitor in the same area that might be able to undercut the new entrant on price or upgrade their own facilities. There are a number of approaches that can be tried to provide "patient" capital or reduce the need for up-front investment:

- Wireless mesh networks essentially allow a network to grow with each new user. The only, albeit substantial, investment is middle-mile backhaul. However, such networks may require significant expertise to set up and are probably best suited for DSL-like per-user bandwidths.
- Demand aggregation can take many forms, with Internet cafes, community centers, libraries and schools offering well-known examples. Cooperatives can aggregate both last-mile and middle-mile demand.
- In some countries, governmental and non-governmental organizations provide loans at low interest rates to eligible entities to fund the construction of networks. Similarly, rural electric cooperatives may have longer investment horizons than public companies.
- Vendor financing may be another option, where vendors of equipment and construction

[4] See http://www.cdnpricing.com for examples and history.

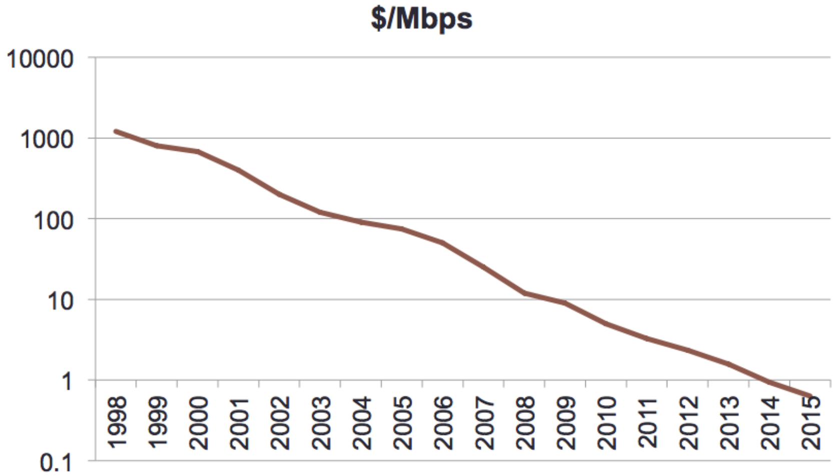

Figure 8 Transit prices

services retain a portion of the revenue. This is attractive if the vendor can obtain better financial terms than the local network operator[5].

- Franchise models have been successful in areas ranging from take-out food to dog kennels, but appear to be uncommon for providing Internet access. Franchising provides two key benefits: It allows to centralize functions that require significant technical expertise and that enjoy economies of scale, such as, in our case, network management, billing and other OSS functions, as well as bulk purchasing. Franchisees contribute their own capital, labor and knowledge of the local market. However, the investment needed for each local area may well exceed the typical small-business franchise investment of up to, say, $500,000 in the United States. (At $1,000 per home passed, this would suffice only for small rural communities. Indeed, the emergence of wireless Internet service providers, without the benefit of franchising, seems to support this scale estimate.)
- Broadband Internet access is unique among residential infrastructure, as it is often added long after a home is built. For gas, water, and electricity, home owners typically finance the last few meters as part of the cost of construction, through a long-term low-interest mortgage. For FTTH, the per-home connection costs are as large as half the total. With standardization of interfaces, it is conceivable that at least new homes and developments make fiber part of the initial construction. Indeed, this idea has been proposed several years ago [6] under the moniker "homes with tails", but seems to have achieved limited traction in practice.
- In theory, it might be possible to have content providers, whether commercial or not-for-profit, pre-fund the construction of networks, allowing them to reach new consumers or users. However, aggregating diverse content creators and avoiding the free-rider problem appear to be challenging for an infrastructure that is inherently shared. (This is a very old idea – after all, radio and TV stations fund their own transmission systems.)

[5] However, this model can endanger the vendor, as large telecom equipment vendors found out after the 2000 telecom bubble burst in the United States and many competitive local exchange carriers (CLECs) were unable to meet their debt obligations.

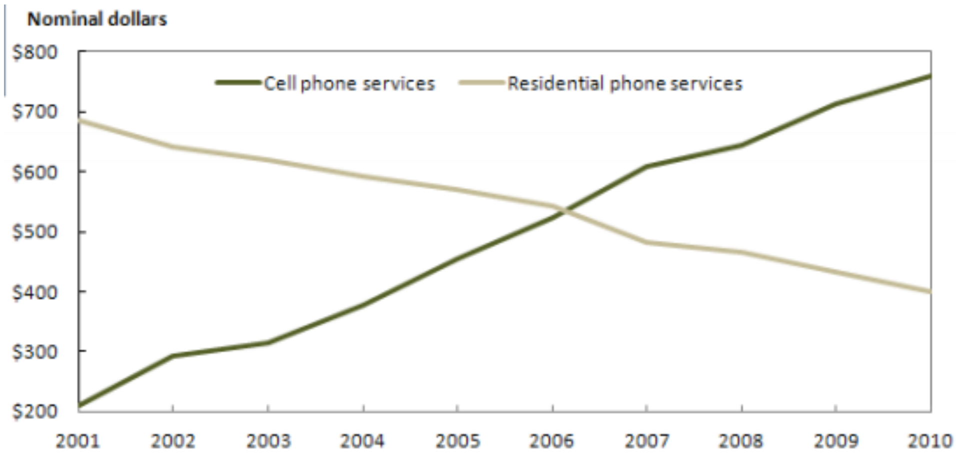

Figure 9 Average annual expenditures on cellphone and residential phone service, 2001–2010 (Source: U.S. Bureau of Labor Statistics December 2011 Volume 2, Number 12).

How to Pay for Networks? American households have significantly increased their spending on telecommunication services, as shown below (Figure 9)[6] [7]. Given that overall disposable household income (Figure 10), and in particular income not committed to health care, housing and food, has not been increasing for most households in Europe and the United States, there may be limited opportunity for additional growth. In many countries, expenses for cellular and Internet services have displaced paying for landline phone service.

It is well-known that the value of bits to the user differs dramatically, with the lowest-bandwidth applications having the largest value per bit. For example, at 10c per minute, a GSM voice call at 13 kb/s costs roughly $1,020 per GB, while 4G data is sold at approximately one hundredth of that price, at rough $10 per GB. Consumers have been willing to pay 10c for each SMS, which translates to $625,000 per GB, even if the short message fills the 160 byte capacity.

While a detailed analysis is beyond the scope of this review, experience seems to indicate that consumers place significant value on predictable charging models [8]. Even simple usage-based charges generates complaints about discrepancies between volume measured by the end system and by the provider or surprises when better connectivity increases the cost of viewing videos. Application-based charging encourages traffic masking, i. e., converting expensive bits into cheaper ones, which then requires deep packet inspection or other countermeasures[7]. For bandwidth-constrained networks, time-of-day charging appears to have the advantage of being easy to understand, and probably approximating more sophisticated congestion-based charging.

Given the roughly factor-ten difference in bandwidth costs between 4G and landline networks (and arguably the largely zero incremental cost for the user), mobile operating systems have already adapted by automatically delaying high-bandwidth activities such as software updates to times when Wi-Fi is available. (Unfortunately, side-loading of video does not appear to be automated yet. Operating systems, including those for tablets and laptops, could provide significantly better policy support to applications.)

[6] http://www.bls.gov/opub/btn/archive/consumer-spending-in-2010-pdf.pdf
[7] For example, Skype masqueraded as web traffic at some point, both to bypass restrictive firewalls and to avoid VoIP limitations.

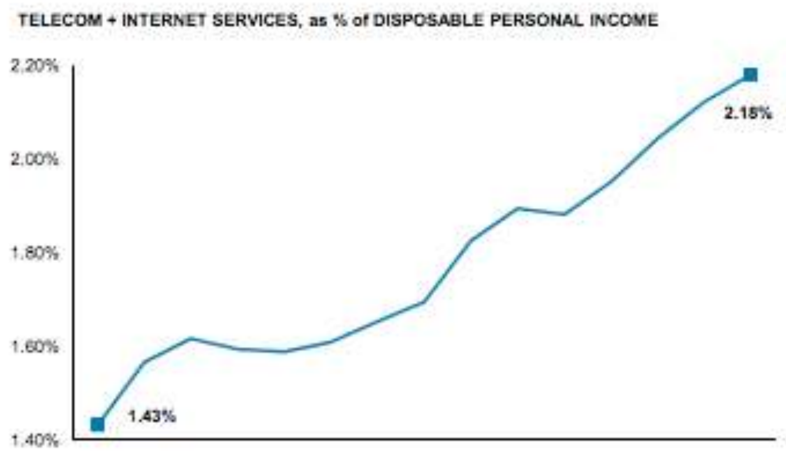

Figure 10 Internet services as % of disposable income (Source: US Equity Strategy: "Macro Meets Micro: A Five-Year View of the Telecom Services Industry", Morgan Stanley, July 21, 2014).

Zero-rating Content – An Alternative? The idea of having content providers or advertisers pay for network access, thus reducing the cost burden particularly on low-income consumers, has received a fair amount of discussion. The economics appear to be challenging, however. For example, at a CPM (cost per thousand) for a video pre-roll advertisement of roughly $10, this translates into 1 cent per video. Since the content creator, e. g., the creator of a YouTube video, receives some compensation and the video distribution service also incurs costs, significantly less than that amount will be available to zero-rate other content. A 720p H.264 video consumes approximately 20 MB per minute, i. e., the typical 30 second video ad itself generates 10 MB. At the incremental rate of $10/GB, the cost is $0.01 per MB. Thus, at current rates, the advertising revenue could not even cover the cost of the ad itself, let alone other content that the user might want to watch.

Universal Service. The concept of universal service, i. e., providing telecommunications services to all residents regardless of geography or income, has a long tradition, dating back in the United States to at least the 1934 Telecommunication Act. There have been three traditional mechanisms to provide service to low income households and high-cost areas, by cross subsidy, by explicit subsidies by rate payers in general and finally through general tax revenue. Both Europe and the United States have had implicit subsidies within the dominant telecommunication provider, i. e., AT&T and the Bell Operating Companies, as well as explicit subsidies with the United States choosing universal service fees imposed interstate telecommunication charges. Implicit subsidies include obligations to serve all reasonable requests within defined geographic areas, such as "carrier of last resort" obligations in many US states. They also include higher-than-cost long-distance rates, intercarrier compensation for terminating calls in rural areas, higher business line rates and above-cost rates for services such as caller ID.

The United States has a universal service program that consists of four parts, with the most relevant for purposes of this discussion being the Connect America Fund to support high-cost (largely rural) areas and Lifeline, for low-income households. In 2013, the program

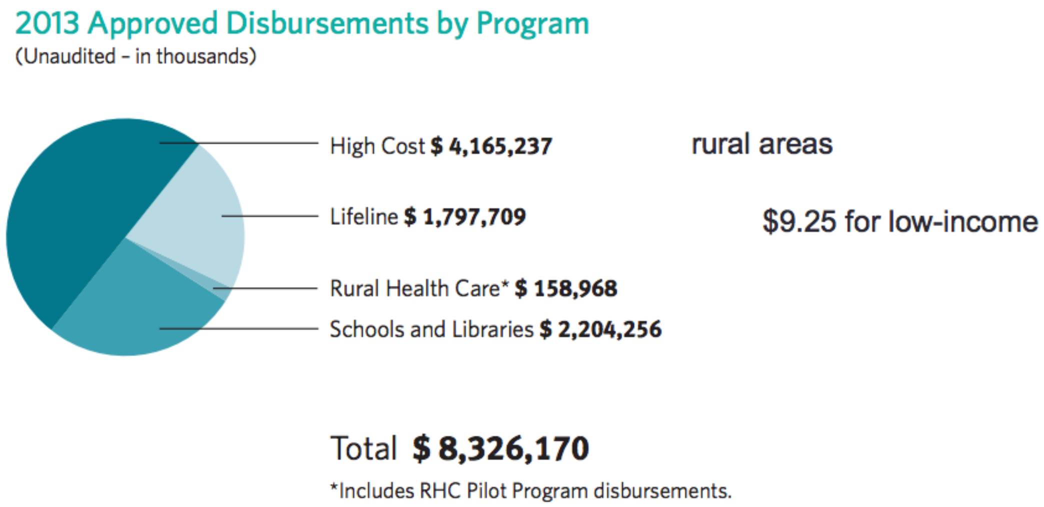

Figure 11 2013 Approved disbursements by program (Source: Universal Service Administrative Company (USAC), Annual Report 2013).

disbursed more than \$8.3 billion, i. e., roughly \$75 per household and year (Figure 11)[8]. High-cost and LifeLine support is paid to carriers, while the rural health care and school and library (e-rate) fund is paid to organizations that then purchase telecommunication services.

A detailed treatise of universal service is well beyond the scope of this summary. The programs have needed significant reforms to address inefficient and fraudulent expenditures. Also, the current contribution mechanism relies on a decreasing portion of consumer and business telecom expenditures, and there is no clear path on whether and how to include Internet access, for example, in the contribution base. Currently, broadband with speeds of 4 Mb/s down and 1 Mb/s up is eligible for support. An increase to 10/1 Mb/s has been proposed.

Currently, the FCC is conducting a small-scale (\$100 million) experiment on allocating high-cost support for higher-speed broadband based on a reverse auction, i. e., the entities with the lowest relative cost per subscriber served receive support, thus encouraging an efficient use of funds.

The Lifeline fund [9] largely supports a \$9.25 subsidy for mobile phone service, which is offered competitively by a number of MVNOs (Figure 12)[9]. A typical plan includes 250 minutes of voice calling per month, along with text messaging. There has been discussion, but no action, on extending Lifeline funding to data services. Subscribers have to be on a social service plan such as SNAP ("food stamps") or be within 135% or 150% of the poverty limit. Each household is eligible for only one phone.

The private EveryoneOn [10] effort offers discounted Internet access, either through some HFC companies (5/1 Mb/s) or via 4G/3G wireless to low-income families with children eligible for free school lunches. For example, 1.2 GB of 4G wireless data is available for a monthly fee of \$10 for families with children that are participating in the free school lunch program. The Comcast Internet Essentials program serves approximately 350,000 households

---

8 http://apps.fcc.gov/ecfs/document/view?id=7521096190
9 http://mercatus.org/sites/default/files/Ellig_FCC-phone-subsidy_MOP_101813.pdf

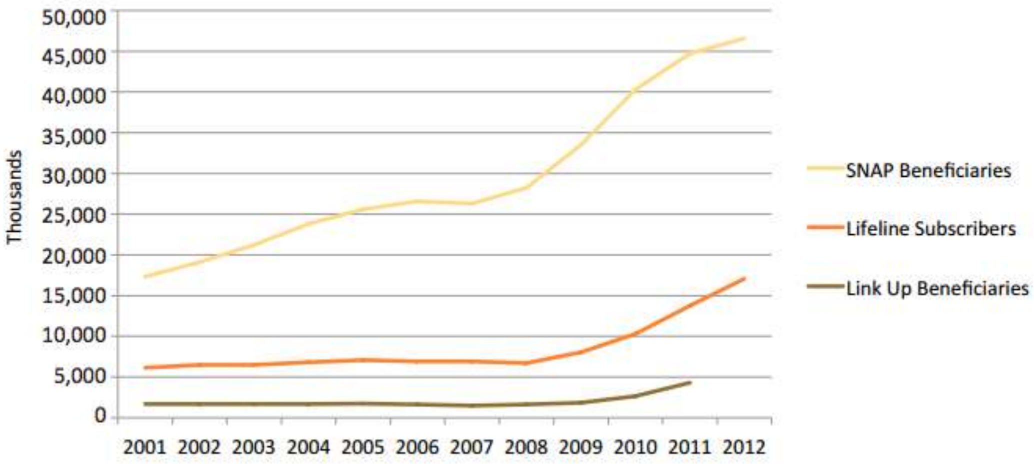

Figure 12 Lifeline subscribers and beneficiaries (Source: Mercatus.org).

in 2014 . The latter program was instituted as part of a merger condition.

These programs illustrate that there are at least two ways to offer affordable Internet access within the existing industry structure, often in combination. From a provider perspective, the carrier wants to avoid losing price-sensitive customers that currently pay significantly more to low-cost plans. Providers can limit the performance of such plans, or make them otherwise inconvenient, so that they are only attractive to households that cannot afford regular plans. Thus, such plans typically only offer limited speed or data budgets. However, reducing the performance too much will likely limit the take rate or may cause users to abandon a service that is unreliable or not capable of supporting common Internet experiences such as watching video content. Secondly, and probably more effectively, such programs can restrict eligibility by income. Providers may thus primarily reach customers that would otherwise not subscribe at all. Since many families move in and out of poverty, they may also help to attract new customers that stay even after they are no longer eligible for the discounted rate.

In summary, addressing availability, affordability and relevance requires both engineering and economic approaches. To make programs sustainable, they need to offer services that meet modern Internet expectations, if logistically possible. Predictable limitations are likely to be more successful than highly variable quality, but there appears to be little quantitative research on many aspects of universal service, including the cost of building and operating networks.

One shoe salesman reported back that there was no market because no one wore shoes. His companion reported back that there was a fantastic market because no one wore shoes.
Edward de Bono, Textbook of Wisdom

References

1 http://www.pewinternet.org/fact-sheets/broadband-technology-fact-sheet/
2 http://99in5.org/resources-news/equity/
3 Se Gi Hong, SungHoon Seo, Henning Schulzrinne, and Prabhakar Chitrapu, ICOW: Internet Access in Public Transit Systems, Under review at IEEE Communications Magazine.

4 Arezu Moghadam, Suman Srinivasan, Henning Schulzrinne, 7DS – A Modular Platform to Develop Mobile Disruption-tolerant Applications, IEEE NGMAST 2008, Wales, UK, September 2008.
5 CTC, 2009, Brief Engineering Assessment: Efficiencies available through simultaneous construction and co-location of communications conduit and fiber.
6 http://www.newamerica.net/publications/policy/homes_tails, (2008)
7 Morgan Stanley, Macro Meets Micro: A Five-Year View of the Telecom Services Industry, July 2014.
8 Andrew Odlyzko, Will smart pricing finally take off?, http://www.dtc.umn.edu/~odlyzko/doc/smart.pricing.pdf, Oct. 2013.
9 http://www.fcc.gov/guides/lifeline-and-link-affordable-telephone-service-income-eligible-consumers
10 https://everyoneon.org/
11 http://corporate.comcast.com/images/MB-10-56-Comcast-Internet-Essentials-Annual-Report-2014-07-31.pdf

6.6 Affordability? Or Willingness-to-pay? The Perspective of Future Markets

Irene Ng

Another perspective of affordability is to look at it not from what the market can afford, but rather what the market is willing to pay. The paradox of understanding a future that has no data is that we can only talk about the future in present terms, using present words, and through the lens (and the data) of the present and formulate some extrapolation based on some assumptions. Yet, if we do so, we run the risk of using the past to predict the future, which, as past innovation have gone on to demonstrate, may not be effective. Affordability is a case in point. We may be able to afford a dictionary, a flashlight or a calculator but when they become freely available on our smartphone, our willingness to pay for the phone increases (though probably not proportionately). That means the current boundaries of products requiring connectivity may not hold and the market is not a zero-sum game where one assumes a market size based on current prices and then analyse how total revenues may be distributed amongst all players within the network. Instead, there may be greater demand for connectivity, quality and speed but this commensurate with an expansion of markets and revenues for new products hereinto not yet known.

To understand a future where markets could expand resulting in greater demand for connectivity and yet generating marginally greater revenues, one would need to know how markets are re-forming in the digital economy and the drivers for such expansions.

Commodification: Commodification (a term which dates back to 1975, believed to have originated from Marxist political theory) is used to describe the process by which something with no economic worth is assigned some worth. This is different from commoditisation which is when an offering moves towards undifferentiated competition. The Italian economist Piero Sraffa [1] once said that firms are involved in "the production of commodities by means of commodities", which means that firms make offerings based on what is offered to firms. This line of thinking of course has led to the commodification of many aspects of life that some consider excessive. Our dreams, our fears, our stem cells have been commodified to become offerings of status, security and hope. Dallas Smythe, back in 1994 before the proliferation of the Internet, considered TV audiences as commodities. By watching advertisements, the audiences "work" for the media by letting themselves be marketed to and by doing so, media firms mass produce audiences and "sell" them to advertisers [2]. Audiences therefore "labour" for distribution and consumption of things produced. The moral limits of commodification have also been widely discussed [3].

The digital economy has seen a proliferation of commodification practices since many of our value-creating activities have become digitally visible. We click on ads suggested by Google, write information on Wikipedia, update our status on Facebook, tweet about our lives and our opinions, and upload videos on YouTube. With our lives becoming more digitally visible, we produce more data about ourselves even as we consume more digital offerings. We become "prosumers", a term coined by Alvin Toffler in 1980 to describe a consumer who produces as well as consumes [4]. This thinking has been extended to the digital realm where the digital self becomes "digital labour" through digital value-creating practices [5]. The technical, functional and social capabilities of mobile media mean that there are more varied forms of commodifying our digital labour practices [6]. We contribute "creative", "intellectual" and "emotional" labour through the way we interact [7],[8]. These digital labour practices are made ever more possible through smartphones. Musicovery is an app that categorises your music collection on your smartphone and plays music to you according to your mood, in degrees of "dark", "positive", "energetic" or "calm" modes. While doing so, your mood becomes part of digital labour. Indeed, our digital labour spawns even more digital labour through the need to maintain our digital identities and social networks. Firms learn about their customers through the collection of their personal data, and repackage them to sell on to other firms. There are many out there who have been campaigning to protect the rights of the increasingly perilous, alienated and exploited, digital "labourer", arguing that commodification of digital labour dehumanises us.

Commodification of labour or digital labour does not automatically assume that worth is created or that market expands. Rather, it is the step towards which human activity contributes to the firm‘s business model and the firm‘s capability to create new connected products to compete in digital markets. Commodification is therefore a potential resource acquisition for a market expansion strategy in future markets.

Digitisation and market expansion: On the 19th of January 2012, Eastman Kodak US filed for bankruptcy protection. An iconic firm formed in 1889 by George Eastman, Kodak dominated the photographic film business for more than 100 years. During its heyday in 1976, it commanded 90 percent of film sales and 85 percent of camera sales in the U.S [9]. Hundreds of business school students and numerous newspapers, magazines and journals have described and analysed the Kodak failure, with many concluding that Kodak did not move quickly enough into the digital era even though the digital camera was invented by Steven Sasson, a Kodak engineer. Kodak‘s dallying for fear of cannibalising its film business resulted in competitor Canon seizing a sizeable market share. The subsequent ubiquity of camera phones created further problems. Urban legend has it that a well-known business consultancy commissioned by Kodak to research into the future of photographs around the turn of the century, when approximately 86 billion analogue photographs were taken per annum, reported that only one in 10 photographs printed were shared. The rest were left gathering dust in photo albums. So it was concluded that the future of photographs would probably not be in the direction of shareability. Imagine the surprise when, within 10 years by 2010, the number of photos shared online is 11 billion and rising, with experts predicting that the numbers will double by 2015. The number of photographs taken have also soared, with some studies estimating the total number to be around 380 billion. Printing photos at home and in store has dropped to a third at 27 billion. The digital revolution of photography is a case in point on the power of digitisation to change both the message (what photos are taken) as well as the medium (the channels for taking the photos and for communication and sharing). Where we used to take photographs for memories, the ability now to generate (and erase) a photo on demand means that photographs could be taken for many reasons without incurring any further economic costs. Today, photographs are no longer just for keeping memories but for collecting evidence after an accident, or photographing a notice or telephone number when there is no pen or paper. The consultancy engaged by Kodak over 10 years ago didn‘t get it wrong. It was probably true that only one in 10 photographs were shared before photos were digitised; it‘s just that the sharing mechanism was through a printed medium. Today, freed from the need to be printed, photographs and other images have become social resources on steroids. The photograph as an offering has completely changed in terms of its benefits, its content, and its channels, from creating it to communicating it, labelled as a digital backwash. Digitisation therefore creates an expansionary effect on markets that would drive changes to the product, its content, the medium and the message.

Consumption and Experiential Contexts: The New Focus For Markets. Digital connectivity is also driving the location of markets. Traditionally, market exchanges happened at retail locations. We bought what we needed from shops, whether they were just around the corner or required a shopping trip into town. As Internet access became ubiquitous, online shopping through the World Wide Web became commonplace. Yet, this still meant visiting websites in the same way we visited physical shops. Buying can therefore be viewed as an interruption to our lives (unless you are seeking retail therapy or window shopping). If we wanted something, buying it would be a “cost” in terms of effort, to get what we really want-which is to use or experience it. Google became popular through its ability to match our search needs with websites, reducing such costs. Even so, it still meant that we had to buy in advance of use, whether from a shop or online.

With greater Internet connectivity, exchanges are slowly becoming closer to the contexts of use experience, even allowing us to choose how we wish to be served. We can now do our banking, watch the latest movie on demand and read our newspapers without leaving the house and without waiting too long to do it.

As more digitised offerings become available, we have switched our purchases from single-function items to a “platform” that allows us to buy more digitised offerings to serve us in context. Many widgets in our homes will want to be the platform to serve up more services on demand, as firms begin to realise the strategic advantage of being a potential channel to serve individuals in context and on demand, rather than in traditional market places.

What this means is that the separation between the location of where exchanges happen (i. e., the market space) and the consumption and experiential space is collapsing into a similar time and space. For products that can be fully digitised, such as music, this spatial collapse has had a profound impact on the business and economic model of an entire industry. Digitisation, as Normann [10] puts it, enables individuals to create what he calls “density”. Density seeks the mobilisation of the best combination of resources for a particular situation. Ultimately, density means that the customer has the world of specialist knowledge available when and where they like. Creating density is the fundamental driver to the digital economy. Being able to digitally enable resources to be used in context and on demand has an expansionary effect on the market as goods and services that serve contexts also means that firms can tap into a market that is more willing to pay, since the service or product acquisition is at the point of need. It also helps firms create new business models of both ownership and access. For example, the market for the ownership of luxury cars currently stands at 70 million cars sold in 2012. That is, however, the market for the ownership of these cars. Yet, there may be many out there who do not wish to own, but would like to drive a Porsche on Mondays and a Rolls-Royce on Saturdays. If digital services in cars can help us detect where they are and what condition they are in, there could be a vast appetite for such services. Technology that makes this viable will create an expansionary effect on the entire automotive industry.

The above outline only 3 drivers of market expansion that could increase the total market affordability for greater infrastructural investments in connectivity. There are many more such as lowering coordination costs, reducing overall exchange risks, emerging latent needs and creating increasing returns of scale through both scalability and personalisation, etc [11]. While market expansion could justify greater infrastructural investment, it doesn‘t however, solve the issue of digital exclusion. However, a more tightly connected ubiquitous market of digital connectivity can create positive externalities that could benefit those who are marginalised by digital exclusion.

Society has benefited greatly from network effects of digital offerings, directly and indirectly. Amongst many other benefits, consumer welfare is increased because of greater choice, lower coordination and reduced search and information costs. Digital connectivity has made us share and collaborate more. This suggests that the digital economy, with lower coordination costs and high scalability and low marginal could reallocate quickly, with the right market incentives to do so. This is especially so when the strategy is that of market expansion.

References

1. Sraffa P (1960) Production of Commodities by Means of Commodities: Prelude to A Critique Of Economic Theory. Vora and Co., Publishers PVT Ltd
2. Smythe DW (1981) Dependency Road: Communications, Capitalism, Consciousness And Canada. Norwood, N.J.: Ablex
3. Sandel MJ (2000) What Money Can‘t Buy: The Moral Limits Of Markets. Tanner Lectures on Human Values 21:87–122
4. Toffler A (1984). Future Shock. Bantam
5. Manzerolle V (2010) Mobilizing The Audience Commodity: Digital Labour In A Wireless World. Ephemera: Theory of Politics in Organization 10(4):455–469
6. Hearn A (2008) Meat, Mask, Burden: Probing The Contours Of The Branded “Self”. Journal of Consumer Culture 8(2):197–217
7. Fuchs C (2009) Class, Knowledge And New Media. Media, Culture & Society, 32(1):141–150
8. Cohen NS (2008) The Valorization Of Surveillance: Towards A Political Economy Of Facebook. Democratic Communique 22(1):5–22.
9. Shapiro C & Varian HR (1999) Information Rules: A Strategic Guide To The Network Economy. Harvard Business Press
10. Normann R (2001) Reframing Business: When The Map Changes The Landscape. John Wiley & Sons Inc
11. Ng, Irene C L (2014) Creating New Markets in the Digital Economy: Value and Worth, Cambridge University Press, UK

6.7 Rights on the Internet

Michael P. Fourman

Information is power. We argue that the changes enabled by modern information and communication technologies have fundamentally changed the balance of power, and that we must revisit established rights to freedom of speech and access to information and education if we are to avoid a fundamental digitally-enabled divide between those who can wield the digitally-enhanced power of information and those who cannot.

The per-subscriber cost of providing Internet access depends on both geography and uptake. Uptake depends on a combination of perceived relevance and affordability, while per-subscriber cost of provision depends on terrain and population density. These factors mean that, in most countries, market-driven provision will not be universal. In the UK provision of modern Internet access to the 'final third' of the population was deemed uneconomic.

Increasing Internet access has social value. It enables significant cost reductions and efficiencies in operations and service provision for all sectors of society: public sector, commercial, and third-sector organisations. It provides local businesses with extended visibility and access to global markets. It provides new tools for communication that enhance social cohesion. Network effects make all of these factors increasingly valuable for existing participants once new individuals and organisations come online.

These positive externalities justify intervention to address the market failure to deliver universal provision. However, a purely economic argument, based on such externalities has also failed to deliver anything close to universal provision. Those who remain unconnected, or with inferior service already isolated and/or disadvantaged, and their social exclusion is exacerbated by an increasing digital divide, as those who are well-connected enjoy increasing benefits from the digital economy, and existing non-digital services become uneconomic and are withdrawn.

This social argument has led some to suggest that Internet access should be seen as a human right. Others, with whom we side, disagree. For example, Vint Cerf has said, "There is a high bar for something to be considered a human right. Loosely put, it must be among the things we as humans need in order to lead healthy, meaningful lives, like freedom from torture or freedom of conscience."

Nevertheless, it is easy to argue that Internet access plays a significant role in enabling several established rights, such as the rights to freedom of expression, to work, to access information. As Vint Cerf says, "...technology is an enabler of rights, not a right itself. The Internet has introduced an enormously accessible and egalitarian platform for creating, sharing and obtaining information on a global scale. As a result, we have new ways to allow people to exercise their human and civil rights."

We claim that existing rights are insufficient to ensure that everyone has an equal opportunity to exercise these opportunities. These technologies amplify our human abilities to store, process, and communicate information. The amplification is extreme. Modern machines can have the power of hundreds or thousands of horses. Nuclear explosions can be many thousands (or millions) of times more powerful than the largest conventional detonations. Our information technologies enable the collection, storage, processing and communication of information at a scale that is millions (or trillions) of times that of a naked brain.

Technology is an enabler of rights, not a right itself. We propose fundamental rights to store, process, and communicate information. Clearly internet access is neither necessary nor sufficient for the exercise of such rights; but it can certainly contribute to enabling them.

The rationale for such rights is that equality of access to the power of information is required to ensure equality of opportunity to choose and pursue a worthwhile life. Vulnerable and marginalized groups in societies tend to be excluded from the opportunities afforded by information technologies.

Rights have value only in their consequences. Our exploration of the consequences of our proposed information rights is an adaptation of an analysis, The right to health. Fact sheet No. 323, published by the World Health Organisation.

The right to health means that governments must generate conditions in which everyone can be as healthy as possible. Such conditions include ensuring the availability of

- health services,
- healthy and safe working conditions,
- adequate housing and nutritious food.

The right to health contains four elements: Availability, Accessibility, Acceptability, and Quality. It imposes on States three types of obligation: to respect, protect, and fulfil the right.

Our proposed information rights would have analogous consequences. For example, that governments must generate conditions in which everyone can share in the benefits of an information society. Such conditions include ensuring the availability of

- information infrastructure,
- a safe information environment,
- adequate access to open data and information.

More generally, information rights include four elements:

- Availability: A sufficient quantity of reliable infrastructure, services, tools and data.
- Accessibility: Infrastructure and services accessible to everyone. Accessibility has four overlapping dimensions:
  1. Non-discrimination
  2. Physical accessibility
  3. Economical accessibility (affordability)
  4. Information accessibility.
- Acceptability: All infrastructure, goods and data must be respectful of ethics and culturally appropriate as well as sensitive to gender and life-cycle requirements.
- Quality: Infrastructure, services and data must be sound and adhere to open standards.

Information rights would impose, on States, obligations to respect, protect, and fulfil the rights. Respect: means simply not to interfere with the enjoyment of the right ("do no harm"). Protect: means ensuring that third parties (non-state actors) do not infringe upon the enjoyment of the right (e. g., by regulating non-state actors). Fulfil: means taking positive steps to realize the right (e. g., by adopting appropriate legislation, regulations, policies or budgetary measures).

Recognising such information rights would provide a framework for the development of policy. In particular, such information rights would justify the commitment to universal inclusion that is necessary to avoid a growing digital divide.

The requirement to respect these rights would inform discussions of privacy, surveillance, censorship, and some aspects of net neutrality. The requirement to protect would entail obligations on governments to act against abuses and cybercrime. The requirement to fulfil would include actions to ensure the availability and accessibility of both infrastructure and education to ensure that all citizens are able to share in the benefits of a digital society.

Another core obligation would be the adoption and implementation of a national strategy and plan of action. This must address the information and communication concerns of the whole population; be devised, and periodically reviewed, on the basis of a participatory and transparent process; contain indicators and benchmarks by which progress can be closely monitored; and give particular attention to all vulnerable or marginalized groups.

Just as for health, each country must move forward in line with a principle of progressive realization. Elsewhere in this document others outline some deliberate, concrete and targeted steps towards universal Internet access that should be taken using the available resources. These resources will include those allocated within a State as well as resources available through international assistance and cooperation.

Participants

- Panayotis Antoniadis
ETH Zürich, CH
- Roger Baig Vinas
Guifi.net-Barcelona, ES
- Saleem Bhatti
University of St. Andrews, GB
- Georg Carle
TU München, DE
- Weverton Cordeiro
Federal Institute of Para-Itaituba, BR
- Jon Crowcroft
University of Cambridge, GB
- Roderick Fanou
IMDEA Networks Institute-Madrid, ES
- Michael P. Fourman
University of Edinburgh, GB
- Thomas Hühn
TU Berlin, DE
- Karin Anna Hummel
ETH Zürich, CH
- Renato Lo Cigno
University of Trento, IT
- Leonardo Maccari
University of Trento, IT
- Mahesh Marina
University of Edinburgh, GB
- Leandro Navarro
Polytechnical University of Catalunya-Barcelona, ES
- Irene Ng
University of Warwick, GB
- Jörg Ott
Aalto University, FI
- Fernando M. V. Ramos
University of Lisboa, PT
- Arjuna Sathiaseelan
University of Cambridge, GB
- Henning G. Schulzrinne
Columbia Univ. – New York, US
- Gareth Tyson
Queen Mary University of London, GB
- Falk Von Bornstaedt
Deutsche Telekom-Bonn, DE
- Klaus Wehrle
RWTH Aachen, DE
- Michael Welzl
University of Oslo, NO
- Adam Wolisz
TU Berlin, DE
- Rüdiger Zarnekow
TU Berlin, DE
- Marco Zennaro
ICTP-Trieste, IT
- Anatolij Zubow
TU Berlin, DE