Exploring the Role of Smart Data Pricing in Enabling Affordable Internet Access

Arjuna Sathiaseelan[§] Gareth Tyson[‡] Soumya Sen[†]
[§]University of Cambridge, UK, [‡]Queen Mary University of London, UK, [†]University of Minnesota, USA

***Abstract*—There are several challenges for enabling an Internet that is universally available and accessible to all. These challenges can be overcome by innovations in technology, economics and regulation/policy. Over the last few years the smart data pricing community have proposed several innovative pricing solutions which have the potential for reducing the cost of Internet access. However, it is imperative to understand the network operator's perspective regarding the applicability and wider deployability of these proposed models. In this paper, we discuss the benefits and challenges for the successful adoption of these models backed by a small scale survey filled out by several network operators across the globe.**

I. Introduction

Universal Internet Access is considered one of the fundamental requirements of today's digital age. As such, enabling universal Internet access is one of the key issues that is being currently addressed both nationally and internationally. It is estimated that only 41% of the world's households are connected to the Internet. Half of them are in less developed countries, where household Internet penetration has only reached 28%. This is in stark contrast to the 78% of households in more developed countries [1]. The reasons behind poor Internet penetration in developing regions are diverse, including regulatory, economic, political, geographical, educational and sociological factors. Importantly, broadband is still too expensive in developing countries [2], where it costs, on average, more than 100 per cent of average monthly income (compared with just 1.5 per cent in developed countries).

Unaffordability is not just a developing regions issue: most of the challenges are also applicable to the developed world as well. For example, access to standard broadband is still not possible for many individuals in developed countries who find themselves unable to pass a necessary credit check while others are simply living in circumstances that are too unstable to commit to lengthy broadband contracts [3]. Consequently, we argue that *pricing* is one of the fundamental obstacles that must be overcome. Without suitable pricing models, customers will not be able to afford access, while organisations will not be able to gain sufficient revenue streams to support universal deployment.

Although there is no silver bullet to remove socio-economic barriers, we believe that developing and deploying new targeted *economic models* is a vital initial step. We further posit that these should be combined with new and innovative *access models* that change the ways that people interact with the Internet (*e.g.,* introducing delay tolerance). Flexibility is a key requirement for ensuring that the diverse sociological, economical and geographical make-ups of different regions can be effectively handled. This flexibility could be embodied in a wide range of pricing and access models. These include opportunities for multi-sided markets; third parties to become virtual network operators [4] (*e.g.,* charities); new models for revenue creation from currently underutilised infrastructures [3]; allowing time-shifted services to utilise offpeak capacity; the introduction of micropayments for fixed broadband networks (not just for mobile as currently seen [1]) and reverse payment models.

This paper surveys a range of recently proposed smart data pricing (SDP) and access models, exploring their applicability for enabling global Internet access. As of yet, few have been widely deployed, particularly in regions that have poor levels of connectivity. As such, we focus on *near future* deployment of these techniques. To help fuel our discussion, we issued a questionnaire to network operators asking their opinions of a variety of models we consider to be well suited for improving access. We use this to explore key issues and challenges that remain in the domain. We make a number of observations that will help both academics and industrial researchers better focus their efforts on driving forward the deployment of these new models.

II. Survey Methodology

Introducing new economic and access models requires the direct cooperation of network operators. As such, to fuel our discussions, we circulated a questionnaire amongst various operators (network operators, virtual network operators and community network operators). This asked questions about various recently proposed pricing and access models aimed at enabling low cost access for the disadvantaged. It was publicised on several related mailing lists (for e.g NANOG, RIPE, AfNOG etc), as well as through social media.

We received 21 responses to our survey, out of which 12 were network operators, 2 were virtual network operators and 7 were community network operators. As information is sensitive, we did not ask operators to identify themselves. However, we did collect data on where the operator is based, as well as which form of demographics they primarily serve. Table I presents on overview of this data[1].

Before continuing, it is important to highlight one point. We asked operators how they perceived the Internet; Figure 1

[1]GDP numbers refers to the approximate World Bank ranking of the network operators country of operation which the respondent chose not to disclose.

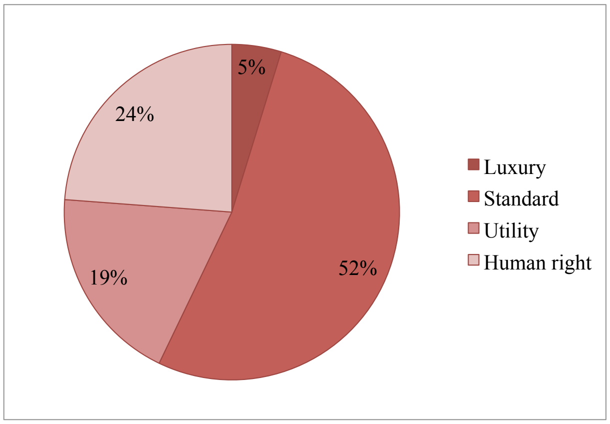

Fig. 1. Perception on the Internet service they delivered.

presents the results. We found an interesting dichotomy, with community operators *all* viewing Internet access as a human right (compared to commercial operators who tended to label it as a standard product). We communicated directly with several community operators to understand their perspectives, as several were happy to de-anonymise their responses. We found that community operators were *not* interested in SDP. This was not because they disagreed with the benefits but, rather, because most felt that network access should be entirely *free* and that SDP models were not applicable to their goals. Although this is an interesting finding in itself, this clearly creates bias when exploring the potential of pricing in the wider domain. Consequently, unless otherwise stated, our results relate solely to the other operators (network operators and virtual network operators).

III. A Review of Pricing Models for Consumers

We begin by exploring recently proposed pricing models, and their applicability to enabling widening access to the Internet. Here we focus on pricing models for consumers, whereas in the next Section (IV), we look at diversifying this (*e.g.,* sharing costs with third parties). Broadly speaking, these pricing mechanisms can be classified either as static or dynamic on whether the price points adapt to the congestion level in the network. We provide a brief discussion of these pricing schemes below; a detailed survey of various pricing plans can be found in [5].

Popular static pricing mechanisms include flat-rate pricing, usage-based pricing, time-of-day pricing and app-based pricing. *Flat-rate pricing* plans charge users a flat monthly fee for access, irrespective of the actual time spent on the network or data usage. While such plans are simple and convenient for billing purposes, such pricing plans are quite inefficient in extracting revenue for elastic traffic if all users run identical applications but with different valuations for them. Such plans have also been criticised for being unfair to light users. Flat-rate pricing is often accompanied with negotiated contracts under which the rate plan is adjusted from period to period depending on the user's resource requirements. *Usage-based pricing* attempts to solve the latter problem by charging users according to the amount of data used. Although simple, they do not offer approachable pricing models for low income users. For example, usage-based pricing will simply force such users into using the Internet less. In essence, they do not allow low income users to devise ways to intelligently maximise their usage for a low cost. In our survey, we asked the operators to specify which of these pricing plans have they been using in their network. Figure 2 presents the number of operators who partly, exclusively or never supported the different pricing schemes (both static and dynamic). From the survey it was evident that all the operators apart from one South American operator have either exclusively or party deployed the flat-rate pricing model (including unlimited data plans), suggesting that the flat-rate pricing model is the most popular due to the lower administrative overhead and ease of billing. Usage-based pricing is also popular with 71% of the operators having partly used this scheme. We argue that these more simplistic schemes are often not helpful for low income individuals. There are, however, several other more flexible smart pricing schemes currently proposed:

Time-of-day pricing charges different hourly price points that are pre-determined based on the expected temporal demand pattern. The most basic form of a two-period time-of-day pricing uses different rates for day time and night time usage. A dynamic version of time-of-day pricing adjusts these prices in near real time based on the network congestion measurements. From our survey, time-of-day pricing seems to be less popular with only 28% of the operators (from Africa and Latin America) partly using them. However, this could be attractive to low income users who could limit their usage to the cheapest time period.

Application-based pricing uses differential pricing of applications to create personalised service offerings. A common realisation of such pricing is through the creation of toll-free apps or zero-rated plans in which application providers (or advertisers) subsidise the users' costs. Such pricing plans, however, have faced some concerns about user lock-in, deep packet inspection and net neutrality. A related pricing plan is that of *Sponsored content* in which third-party advertisers reward users with incentives (*e.g.,* additional quota over their data cap) in exchange for the user's engagement with their content and advertisements [6]. Arguably, such plans are more network neutral in the sense that these do not tie the incentives to any specific application(s) that the user has to use. In our survey, 43% of the operators (from North America, South America and Africa) have partly used these models. This could be extremely beneficial for initial Internet uptake as it offers the opportunity for zero cost access. However, this should be considered carefully, as long-term implications could be serious, with only a subset of Internet services available to those who have not been able to move beyond application-based pricing.

Congestion-based pricing adapts the price to reflect network congestion. Such ideas include *Real-time Responsive pricing* (which sets prices so as to keep user demand under a certain

Continent	# of Operators	TIER	Urban/Rural	Countries
Africa	4	1, 2 &3	Both	South Africa
Asia	1	not disclosed	Rural	India
Australasia	2	2 & 3	Urban, Rural and Remote	New Zealand, Vanuatu
Europe	7	2 & 3	Both	Czech Republic, Italy, Austria, Germany, Catalonia, Spain, Italy, not disclosed (GDP 51-75)
South America	4	1, 2 &3	Both	Ecuador, Venezuela
North America	3	1, 2 &3	Both	USA, not disclosed (GDP 1-25)

TABLE I
SURVEY RESPONDENTS: DEMOGRAPHY AND GEOGRAPHY.

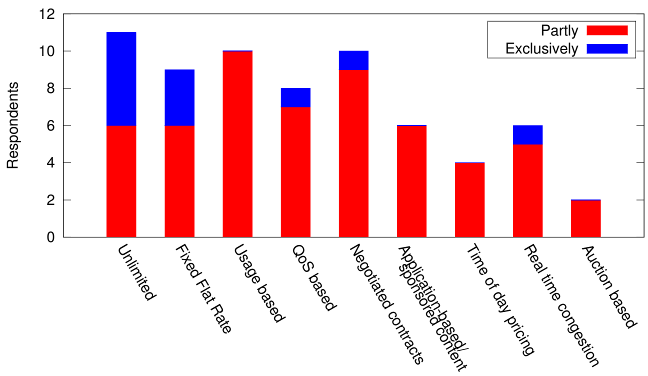

Fig. 2. Number of operators who support the different pricing schemes (partly or exclusively).

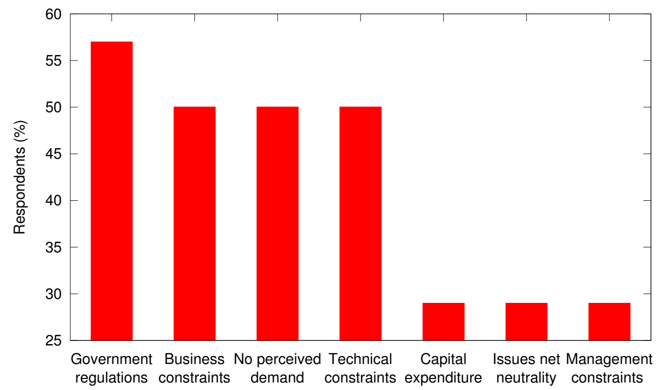

Fig. 3. Main barriers to deploying new pricing models for affordable communications.

threshold), *QoS or Priority pricing* (which explicitly accounts for QoS by allowing users to pay less by accepting a longer delay at congested times), and *Smart Market Auction pricing* (which decides whether to admit a packet into the network at congested times based on the user-specified bid attached to that packet) [7]. This, again, could be beneficial to low income users who may have the flexibility to wait before using the Internet. Alternatively, they may be willing to accept a lower quality of service in return for reductions in cost. It is important to note, however, that although these dynamic pricing plans have the theoretical potential to make resource allocation more efficient, they have largely remained unrealised in the market. Our survey points out that only 14% of the operators (from Latin America) have partly used auction based pricing. However, 57% of the operators have partly or exclusively used QoS or Priority pricing models and 43% of the operators (from Africa, Australasia and Latin America) have used real-time responsive pricing.

Given that these flexible pricing schemes have the potential to enable low cost access, we next explore the barriers blocking their deployment. We asked what operators perceived the key challenges to be. Figure 3 summarises the results. Regulation was the key point raised, with 57% saying that government regulations limited their flexibility on pricing. Interestingly, many also mentioned that there was a lack of customer demand. Clearly, hurdles do not solely exist on the operator side with the necessity of mobilising consumer interest as well. Without improving awareness among customers, it is unlikely operators will engage. A subset of operators also raised specific concerns over net neutrality, which was somewhat surprising.

A further important observation we made from our survey was that operators from developing economies such as Africa, Latin America and Australasia have either partly or exclusively used all of the different pricing schemes. Considering the technical, operational and cognitive overhead associated with such plans we found this surprising. We hope that there will be a wider adoption of these different pricing schemes in the near future.

IV. NEW ACCESS AND STAKEHOLDER MODELS

In the previous section we discussed several smart pricing models for consumers. Whereas these have great potential, we argue that, in many cases, they will not be sufficient to lower the cost of Internet access alone. In this section, we discuss a set of access models that propose alternative ways of using and financing the Internet beyond customers simply transferring money to their ISPs in return for always-on connectivity. When these models are tied up with SDP mechanisms, we posit that they would be an effective solution to reduce the cost of Internet access.

A. Access Models Utilising Spare Capacity

In many cases, network infrastructures are largely under utilised. This is mainly due to over-provisioning and statistical multiplexing of user traffic (with different user access patterns). A network provisioned for a daytime peak, will naturally be under utilised during a nighttime trough. These observations are precisely what led to the commonly used 95-percentile pricing scheme used by transit ISPs to charge their customers according to peak demand [8]. This pricing model, however, means that often customer networks can generate traffic free of charge during troughs (because they are getting charged for it anyway).

We argue that this observation could be exploited to enable new access models for the economically disadvantaged. Network capacity that has already been paid for by network operators could simply be given (either free or at a low cost) to people who are prepared to be flexible in when (and how) they use the Internet. This would be a dynamic pricing strategy, which might be highly attractive to ISPs, as this is currently an untouched consumer group (and extra costs would only be administrative). In fact, simple versions of this already exist, whereby "nighttime" bundles are included in contracts, *i.e.,* allowances that can be only used during the night.

To ensure that this low cost traffic does not interfere with "traditional" customers, it could be allocated to a basic less-than-best-effort service (scavenger class) while providing opportunities for enhanced QoS through micropayments (discussed in Section IVC) [3]. Although, potentially controversial from a net neutrality perspective, we argue that this practical solution is a viable first step towards enabling wider access as demonstrated in [3]. Reselling the unused capacity could be done using one of many different SDP techniques.

To explore the feasibility, we turn to our survey. We asked if the operators were willing to monetise their unused capacity by selling it at a low cost. 71% of the respondents said yes while the remaining 29% said no. There is some resistance to such models. To push these principles further, we asked if operators were prepared to contribute to the running costs themselves (as a societal service). A surprisingly large number were quite engaged with this idea. More than half said that they would contribute some level. More specifically, 43% of the respondents said they were willing to cover 25% of the costs, while 14% said they were happy to cover 50% of the costs. No one was ready to subsidise 100% of the costs though. The remainder demanded that costs were covered entirely by another (either the customer or a third party, *e.g.,* charity). We surmise that willingness to engage with such pricing models is still extremely divisive. Whereas some see business (and societal) benefit, others do not. It is therefore likely that, initially, it will be necessary to have external forces driving this (*e.g.,* governments). In all cases, it will be vital that third parties are integrated into the process, as no operator offered to pay *all* costs; instead, they always require a revenue stream.

To inspect the barriers to deployment, we then asked what operators perceived the key challenges to be. Figure 4 summarises the results. Interestingly, the key reason listed was a lack of customer demand. Clearly, hurdles do not solely exist on the operator side with the necessity of mobilising consumer interest as well. Without improving awareness and

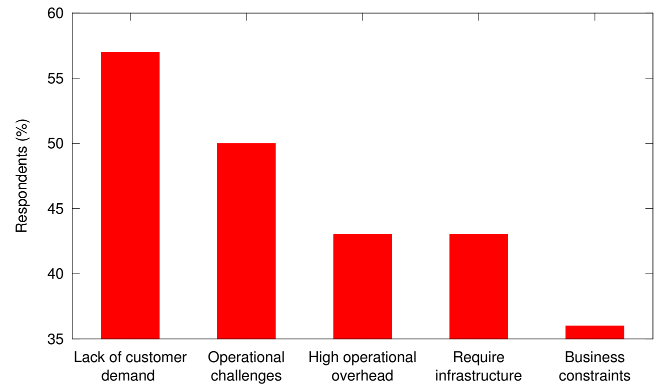

Fig. 4. Main barriers for selling low priority unused traffic services.

demand among customers, it is unlikely commercial operators will engage. Operators also perceived high organisational difficulties, such as effectively forming arrangements with third parties willing to contribute to costs (*e.g.,* charities). Sufficient levels of revenue were also queried, with many worrying that costs and overheads would be too high. This exacerbates other factors, most notably the required expenditure for any necessary underpinning technologies.

B. *Involving New Stakeholders*

Stakeholders and value chains are relatively well defined in traditional Internet economics (*e.g.,* money flows from customer→ISP). Whereas this has generally worked well in affluent areas, this seems ill suited to developing regions. Limiting stakeholder interactions to a small set of players centres an undesirable levels of power within a small group of organisations, and prevents innovation in terms of incentivisation. For example, local government, charities or businesses all have reasons to decrease the cost of human-centered services, and replace them with low cost online services [9]. As such, these organisations become part of the Internet deployment ecosystem. Potentially such an approach could be beneficial across society; for instance, it is reported that the UK could achieve an annual savings of upto £3bn by enabling more online services [10].

We therefore argue that new models that enable such stakeholders to become involved are vital. Novel value chains could be then created for mutual benefit. For instance, advertisement brokers might help subsidise network operating costs, with the longterm intent on gaining more eyeballs on their adverts. Virtual network operators (VNOs) could be a key component in this, enabling smaller scale virtual networks to build up these more innovative and diverse pricing models that involve multiple stakeholders [4]. VNOs could, for example, be run by local government agencies, charities, third party companies or grassroot user communities.

Turning to our survey, 43% said that they were happy to accept money from third parties (such as local government or charities), rather than consumers themselves. This is encouraging to see, as this is the first step towards building models

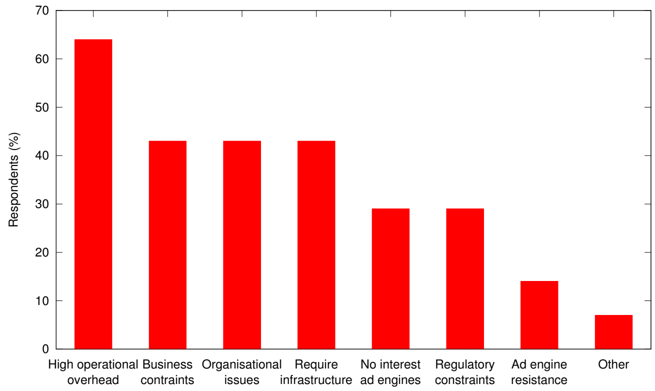

Fig. 5. Main barriers for charging third parties (for e.g ad providers, local government).

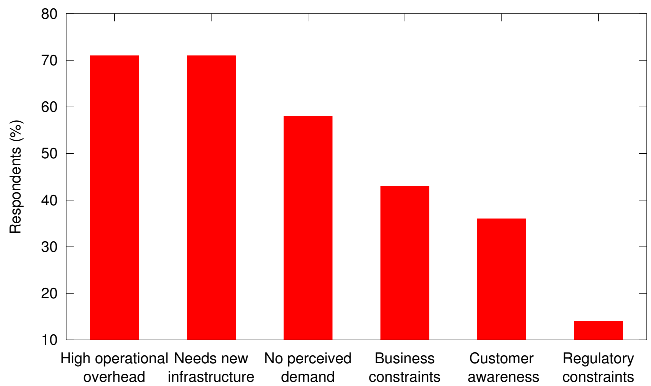

Fig. 6. Main barriers for supporting micropayments.

based on multiple stakeholders. That said, the majority of the operators (57%) disagreed that charging third parties (such as advertisement brokers) was a good idea. Figure 5 lists the variety of reasons that were given. Most prominently, operators were concerned that the organisational costs and complexities would be too high. This indicates that introducing standards for enabling inter-stakeholder cooperation (and billing) would be a positive first step towards motivating operators. Many also noted that there would be a perceived resistance from the third parties: operators are not keen on pursuing these mechanisms without bilateral engagement.

C. Micropayments

Micropayments are small one-off payments for particular services. They turn always-on access into an on-demand model. Most prominently, they are used in mobile networks and have been a major driver in the mobile ecosystem especially in developing regions [2]. However, micropayments are less prevalent in fixed home broadband networks. This is largely because there can be a significant per-customer cost in wired environments (*e.g.,* sending access equipment, laying cable). Micropayments, however, could take place in fixed networks if low income users could acquire connectivity via an existing customer (*i.e.,* by connecting to their WiFi). Innovative incentives could be devised here, where a portion of the micropayment is given to the customer providing the edge connectivity.

A variety of short fixed usage services could be provided via micropayments. Usage could be based on time, volume, application or QoS. For example, "turbo boosters" have been pioneered by some network operators, allowing users to press a button on their device, which gives them a short period of higher quality of service. This can allow users to only pay for high bandwidth when wishing to watch videos. This would be highly beneficial for occasional users, with diverse needs. Micropayments could also be combined with other pricing mechanisms such as reverse payments. For example, a remote doctor could pay for a short period of additional capacity when video conferencing with economically disadvantage patients [4].

To explore operator opinions of micropayments, we inspect the survey. The majority of the respondents (71%) do not currently support any form of micropayments whatsoever. However, 70% of those who did not support them, stated that they would like to support it in the future. This encouraging response led to us to investigate what operators perceive to be the key barriers for supporting micropayments in the future. Figure 6 presents an overview of the responses. Once again, many operators complained about the high operational costs associated with running such a complex pricing scheme. Interestingly, for micropayments, an equal number of respondents (71%) also said that the need for new infrastructure was a key challenge. A relatively consistent set of respondents stated that a lack of customer demand was also a vital barrier ($\approx$60%). Once again, engaging customers to demand such models appears to be a surprisingly dominant problem in the area of SDP.

D. Timeshifting of Content

Timeshifting is the process by which traffic delivery is moved in time to a point other than its (ideal) point of generation. For example, a user might be asked to delay sending a large file until the off peak period. Nowadays, the primary source of traffic generation is static content delivery, particularly multimedia content (*e.g.,* video). This is predicted to constitute over 85% of Internet traffic by 2018 [11]. In our survey, 43% of the network operators stated that video traffic consumes more than 50% of their peak network capacity. Content delivery scenarios would not be well handled by utilising unused capacity in the traditional sense of the word, because they are often very traffic hungry. Further, often these applications are highly redundant in their transmission of the same bytes multiple times across the same network [12].

An interesting observation about content consumption is that often it is highly predictable. For example, a user who watches an episode of a serial is likely to watch the next one [13]. We therefore argue that this semantic insight could be effectively used to timeshift content *distribution*. Note that this is distinct from timeshifting content *consumption*: we wish to place in-network storage such that content can be pre-loaded during the periods of under utilisation. As such, in the ideal case, users will be able to immediately access content when they request it because it will be stored nearby. For example, if a user is predicted to consume x tomorrow, a local copy of x can be pushed to their home access point during the night. A particularly strong use case for this is highly popular objects such as adverts and homepages that get repeatedly requested. Remarkably, from our survey 36% of network operators state that adverts consume more than 25% of their peak network capacity.

According to our survey results, 79% of the network operators were willing to offer cheaper Internet access if the users timeshifted their Internet usage to off peak hours, with 43% of the operators stating they are ready to give a 50% or more discount to users for incentivising them to time shift their video consumption to off peak hours. It therefore seems that this scheme is one of the most attractive. Discussions with the operators reveal that this is largely driven by the nature of their transit pricing schemes: ISPs are keen to reduce their peak, in favour of increasing their average traffic rates. One operator in particular was willing to give upto 100% discount to incentivise users to timeshift their video consumption. 57% agreed that timeshifting lower priority traffic to off peak hours would also be beneficial.

V. Discussions and Conclusion

A number of findings from our survey are worth discussing. A notable observation across most smart data pricing schemes is the lack of perceived customer demand. In many cases, operators reported that their customers were not aware/interested in these advanced pricing schemes. When combined with the complaints of high operational overheads, many respondents were reluctant to engage. Currently, this is best exemplified by the near universal lack of uptake from our respondents. That said, there were various nuggets of positivity. For example, 71% of respondents said they were willing to sell unused capacity at a discount rate (assuming a third party pays). Operators are not willing to take on the burden of creating additional low cost data plans (because the cost increase will be greater than the revenue increase), however, they are willing to add some incentive mechanisms in their existing plans to regulate demand. For instance, 79% of respondents showed interest in performing time-based pricing to incentivise users to move activity away from their traffic peaks. More complex schemes were less popular, though, with only 43% saying they would be prepared to include new stakeholders (*e.g.,* advertisement brokers) in charging schemes. Finally, 43% of the network operators stated that their network infrastructures do not have the ability to support the above mentioned access and pricing schemes. Operators also find operational overheads to be too high to justify either billing advertisers or using third parties to subsidise customers. However, this may be due to the relatively small size of these ISPs and the regulatory requirements of the regions they operate in. Sponsored data is already being offered by AT&T and Verizon in the US, so potentially the overhead costs of creating such plans are justified for large operators but too high for smaller ones.

There are number of important future research strands in this realm. Our study offers just a first step in this direction, with a need to expand drastically both the breadth and depth of our findings. Content providers have already made efforts to bridge the digital divide with zero-rating, app-based pricing, and sponsored content, all of which open up further interesting questions for SDP research. For example, what implications do such measures have on net-neutrality? In terms of real-world deployments and evaluations, perhaps most critical is the education of citizens in the potential benefits of SDP. Without pressure from consumers, it is unlikely that incumbent network operators would offer such schemes on a large-scale. In terms of research, it is therefore important to perform future user-facing studies that understand how people with little Internet experience could exploit advanced pricing mechanisms. Hence, future SDP research requires multi-disciplinary as well multi-stakeholder collaboration, *e.g.,* via the recently charted Internet Research Task Force Global Access to the Internet for All (GAIA) research group.

VI. Acknowledgements

The work was supported by the EC H2020 RIFE project Grant No. 644663. We would like to thank Ermanno Pietrosemoli, Marco Zennaro, Steve Song, John Klensin and Jon Crowcroft for their useful feedback and help with our survey.

References

[1] ICT Facts and Figures, ITU Report, 2014

[2] Measuring the Information Society, ITU Report, 2014.

[3] A. Sathiaseelan et al, A Feasibility Study of an In-the-Wild Experimental Public Access WiFi Network, ACM DEV 5, San Jose, December 2014.

[4] A. Sathiaseelan and J. Crowcroft, LCD-Net: lowest cost denominator networking, ACM SIGCOMM CCR Vol. 43, No. 2, pp. 52-57, April 2013.

[5] S. Sen, C. Wong, S. Ha, M. Chiang, A Survey of Broadband Data Pricing: Past Proposals, Current Plans, and Future Trends, ACM Computing Surveys, 2013.

[6] M. Andrews, U. Ozen, M. Reiman, Q. Wang, Economic models of sponsored content in wireless networks with uncertain demand, INFOCOM Smart Data Pricing Workshop, 2013.

[7] J. Huang, R. A. Berry, M. L. Honig, Auction-based Spectrum Sharing, ACM/Springer Mobile Networks and Applications, 11(3), 405-418, 2006.

[8] R. Stanojevic, N. Laoutaris, P. Rodriguez. "On economic heavy hitters: shapley value analysis of 95th-percentile pricing." Proceedings of the 10th ACM ICM, 2010.

[9] A. Sathiaseelan et al, Virtual Public Networks, 2nd IEEE European Workshop on Software Defined Networking (EWSDN), Berlin, October 2013.

[10] Cabinet Office, Digital by Default proposed for government services, www.cabinetoffice.gov.uk/news/digital-default-proposed-government-services, 2010.

[11] Cisco Visual Networking Index (VNI).

[12] G. Tyson et al. "A trace-driven analysis of caching in content-centric networks", IEEE ICCCN, 2012.

[13] N. Gianfranco, et al. "Understanding and decreasing the network footprint of catch-up tv", 22nd international conference on World Wide Web, 2013.