Pricing User-Sanctioned Dynamic Fast-Lanes Driven by Content Providers

Hassan Habibi Gharakheili†, Arun Vishwanath*, Vijay Sivaraman†
†University of New South Wales, *IBM Research-Australia
Emails: {h.habibi@unsw.edu.au, arvishwa@au1.ibm.com, vijay@unsw.edu.au}

***Abstract*—The notion of "fast-lanes" for prioritising certain Internet content on residential broadband access links is being vigorously debated today. While Internet Service Providers (ISPs) see fast-lane revenue as an imperative for their economic sustainability, end-users and several Content Providers (CPs) feel threatened by the violation of network neutrality. In this paper we argue that fast-lanes can present a win-win-win situation for ISPs, CPs, and end-users alike, if implemented in a way that gives each party appropriate control knobs. To this end, we present an architecture in which fast-lanes are dynamically created and destroyed on-the-fly on a per-session basis using input from all three parties – this departs from current proposals that have largely focused on static fast-lanes. We then present a simplified economic model that demonstrates the benefits for each entity: users can get enhanced Quality of Experience (QoE) at no extra cost, CPs can monetise user-satisfaction using business-models of their discretion, and the ISP can experiment with two-sided pricing models of their choice. We evaluate our proposal using real traffic traces and multiple pricing models, and set the ground for a deeper study into the economics of dynamic fast-lanes.**

I. Introduction

Internet "fast-lanes" – driven by broadband Internet Service Providers (ISPs) – are currently the subject of a raging debate amongst policy-makers, activists, economists and researchers [1], [2]. Fast-lanes allow ISPs to charge Content Providers (CPs) such as Netflix, YouTube and Hulu to prioritise their traffic. Such an approach could open doors for improved quality-of-experience (QoE) for end-users, while giving ISPs a new degree of freedom (i.e. service quality) to exploit for increasing their revenue.

Residential data traffic is growing at 40% per annum [3], while average revenue per-user (ARPU) for fixed-line ISPs is growing at only 4% per annum [3]. ISPs argue that in order to sustain and upgrade their infrastructure to cope with growing traffic volumes, new business models such as fast-lanes are necessary to help narrow the gap between their cost and revenue [4], [5]. The Federal Communications Commission (FCC) has also expressed interest in this proposal [6]. Understandably, there is a huge backlash from the public because fast-lanes are perceived to give license to ISPs to violate net-neutrality by throttling or blocking arbitrary traffic streams of their choice without regard to consumer interest [7], [8].

A. Static Fast-Lanes

We believe that current proposals for relatively *"static"* fast-lanes are unlikely to lead to sustainable models that will be palatable to all the three parties involved. In one version of this model, the CP pays the ISP a lump-sum (or annual) amount for creation and maintenance of long-term fast-lanes – NetFlix's peering payment to Comcast in early 2014, believed to be in the order of $15-20 million a year [9], can be seen as an example of this model. However, this "bulk-payment" model is unfair towards smaller CPs, who are disadvantaged by not having the up-front capital to pay the ISP to boost their traffic (in other words, this model lacks elasticity). Moreover, end-users are understandably irate at such back-room deals between ISPs and CPs from which they are completely shut-out, and in which they have no voice.

To counter the consumer backlash, AT&T proposed an alternative in October 2014, which empowers the FCC to prohibit the creation of fast-lanes by ISPs, and puts the onus on the end-users to decide which sites and services (video, VoIP, gaming, and such) should receive priority treatment [10]. While the proposal has received measured support [11] from a few quarters, others remain sceptical. We believe that while engaging end-users in fast-lane creation is a worthwhile idea, the static nature of the envisaged fast-lanes does not overcome several critical obstacles:

- *Complexity:* A vast majority of end-users lack the sophistication to configure fast-lane parameters corresponding to each Content Provider.
- *Unfairness:* Smaller CPs with lower consumption volume are less likely to be configured by users compared to large CPs (Netflix, YouTube).
- *Granularity:* Performance cannot be controlled on a per-session basis (e.g. for a specific movie, rather than all content from Netflix).
- *Monetisation:* If the user is charged for creation of the fast-lane, uptake is likely to be low, limiting the ISPs revenue growth. If the cost for user-configured fast-lanes is expected to be passed on to the CP, the negotiation mechanism remains unclear.

B. Dynamic Fast-Lanes

To overcome the various issues listed above, we propose a radically new model in which fast-lanes are *"dynamic"*, i.e. created on-the-fly on a per-session basis. Further, every party has a say in this. The end-user specifies a static number to the ISP, which is the fraction of their broadband link capacity that they allow the ISP to carve fast-lanes out of. This single knob keeps complexity low for the end-user; more importantly, the user has the freedom to completely disable fast-lanes (by setting the knob to 0), or give the ISP full control of their access bandwidth for fast-lane creation (by setting the knob to 1), or choosing any intermediate fraction in-between (recommendations on choosing this parameter will be discussed further on).

In the meantime, the ISP opens a dynamic API that *any* CP can call to carve out bandwidth for their ongoing traffic streams. The API call, if successful, is associated with a "micro-payment" from the CP to the ISP. The CP is by no means forced to call this API for fast-lane creation – if network performance is adequate, or if the customer is low-value, the CP can at their discretion send their traffic in a best-effort manner (the way it is today). This gives the CP granular flexibility in choosing if (and how much) they want to pay on a per-session basis, and the increased elasticity eliminates "bulk-payments" that traditionally disadvantage smaller CPs.

We demonstrated the technical feasibility of API-driven dynamic fast-lane creation, using software defined networking (SDN) platforms, in a recent paper [12]. In this paper our goal is to explore the economic incentives for this approach, and to show that it results in a win-win-win situation for end-users, CPs, and ISPs. We consider flat-rate and congestion-based pricing models for the micro-transactions associated with dynamic fast-lane creation, and conduct simulations using a real trace of over 10 million flows taken from a campus network. Our study reveals the following:

- *End-users* can control, via a single knob, the extent to which they want to avail of fast lanes. We show that permitting fast-lanes in general gives the user better video streaming quality (at no extra cost), but users who value other traffic types (such as mice-flows or elephant transfers) have the ability to restrict fast-lane creation.
- *Content providers* can exercise fine-grained control over their usage of (and payment for) fast-lanes, e.g. by using fast-lanes only when network load is high, or only for high-value customer traffic; we show, using an exemplar revenue-model, that their fast-lane costs can be more than offset by their gains from higher customer satisfaction (e.g. improved QoE for video streaming sessions).
- *ISPs* can monetise dynamic fast-lanes, via micro-transactions, from any and all CPs. Importantly, this imposes no extra cost on the end-user, and allows end-users to adjust (or opt-out of) fast-lanes if they so choose. We investigate the revenue potential for ISPs under both flat-rate and congestion-based pricing models.

II. Related Work

We now briefly review the different smart data pricing (SDP) models and the economics around fast-lanes (touching upon aspects including net-neutrality and sponsored content).

A. Pricing Models for End-Users

Pricing of broadband Internet, i.e. what an ISP charges the end-user, has been extensively investigated. Broadly, these pricing schemes can be classified as being static or dynamic. Static pricing includes flat-rate pricing, where a user only pays a fixed charge in a billing period regardless of the volume of data used in that period. To bridge the growing gap between ISP costs and revenue, several ISPs around the world are offering newer pricing schemes such as usage-based pricing (fee paid is proportional to the volume of data used), tiered pricing (a fixed quota charge and any overage charges for exceeding the quota), and time-of-day pricing (higher charges during peak-hour usage compared to off-peak hours).

Dynamic pricing includes schemes such as day-ahead-pricing (charges for the next day are guaranteed the previous day), and congestion-based pricing (users pay higher prices during higher congestion levels). An excellent survey of the different pricing models aimed at end-users is given in [13].

Our work is orthogonal to the above studies on user-pricing, since we do not aim to affect user-prices or user-behavior, and indeed want to keep fast-lane economics largely transparent to users. Consequently, our scheme is oblivious to the data plans that the end-users have contracted with their ISPs, and we do not make any attempt to affect user behavior by time-shifting their traffic demands.

B. Two-Sided Pricing Models

Several recent works have considered two-sided pricing models, wherein the ISP charges both end-users and CPs. In [14], it is shown that under certain circumstances, net-neutrality regulations can have a positive effect in terms of total surplus under monopoly/duopoly ISP regimes. The work in [15] also studies a two-sided non-net-neutral market, but additionally takes into account QoS provided by the ISP to the end-user. By defining a model for total end-user demand, and using the mean delay of an M/M/1 queue as the QoS metric, the authors theoretically evaluate the conditions under which a charge made by the ISP to the CP would be beneficial (to either of them).

The work in [16] considers a model comprising a monopoly ISP, a set of CPs, and end-users. Focusing on the utility of the ISP/CPs and the resulting social welfare, the authors argue in favour of establishing priority-based pricing and service differentiation rather than on effecting net-neutrality regulations. Using game-theoretic analysis and incorporating models for congestion control algorithms such as TCP, [17] arrives at a number of interesting conclusions: most notably, when regulations are beneficial and when they are not. The authors also introduce the notion of Public Option ISPs, which could be an alternative to enforcing tight regulations.

These works largely consider (semi-)static payment arrangements and evaluate the resulting utility gains using game-theory; by contrast, our model differs by considering dynamic fast-lanes that are created and destroyed on-the-fly, wherein CPs make per-session decisions based on run-time factors such as network load.

C. Economics of Sponsored Content

The concept of "sponsored content" has been studied before [18], [19] – in this model, the end-user pays a lower fee to the ISP due to CP induced subsidies (Facebook traffic being considered "in-network" and not counting towards the user's quota is an example of this). The CP can benefit by attracting more traffic from the end-user, while the ISPs can reduce churn

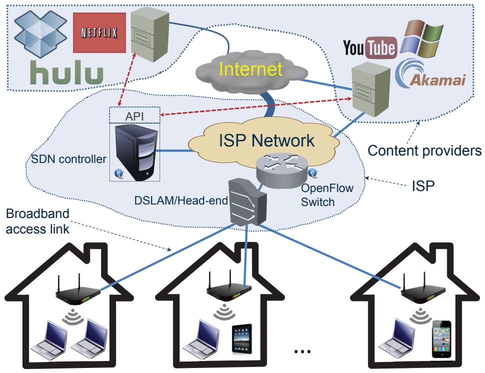

Fig. 1. Illustrative access network topology.

and retain customers. Although our work is loosely linked to this concept, it differs in not ascribing any subsidies to the end-users; moreover, unlike sponsorship models that are long-term contracts between CPs and ISPs, we study the efficacy of a model that permits paid-prioritisation at much smaller time-scales (i.e. at per-session granularity).

III. System Model

Fig. 1 shows a typical access network topology. Each household has a wireless home gateway to which devices within a household (e.g. laptops, smartphones, smart TVs, and others) connect. The home gateways offer broadband Internet connectivity via a line termination equipment (DSLAMs) at the ISP local exchange. The DSLAM is connected to an SDN-enabled Ethernet switch (e.g. an Openflow switch) which in turn connects to the ISP's backhaul network, providing access to the global Internet. The ISP network could peer directly with a CP or indirectly via a CDN or other ISPs. Finally, the ISP network houses an SDN controller that exposes the APIs for a CP to call to establish the fast-lanes.

A. API for Video Streaming Content

The ISP exposes an API for creation of a dynamic fast-lane, available to any CP. The API requires the CP to specify the end-points of the traffic stream (server and client IP addresses and TCP/UDP port numbers), along with the minimum bandwidth (on the access link between the ISP and the end-user) to be reserved for a video streaming application, together with the duration for which the reservation is sought. We have intentionally kept the API specification minimal at this stage, future work can extend it to incorporate peak rate, burstiness, etc. for the video flow. A more detailed description of the API can be found in our prior paper [12], and a prototype implementation was demonstrated in [20].

B. Dynamic Negotiation Framework

The ISP charges the CP each time the latter invokes the reservation API (the cost model is discussed in the next subsection). The decision on whether or not to invoke the API is entirely up to the CP, and can be based on user-class, network conditions, etc. For the purposes of this study, we will assume that the CP does not discriminate amongst users (i.e. does not preferentially call the API for premium users over regular users). However, we do assume that the CP calls the API only when it estimates that the available bandwidth on its path to the end-user (as a fraction of access link capacity) is below a certain threshold θ. A CP that never wants to use the API would therefore set its $\theta = 0$, whereas a CP that wants to invoke the API for every video session (irrespective of network load) would set $\theta = 1$. In general, a CP could choose an intermediate value, say $\theta = 0.2$, implying that it will take its chances with best-effort video-streaming when 20% or more of access link bandwidth is available, and only invoke the API to reserve bandwidth when the residual capacity falls below 20%. We note that techniques for bandwidth estimation such as packet-pairing [21] and packet-dispersion [22] are well-known, and CPs like YouTube and Netflix do active bandwidth monitoring to adapt their streaming video coding rates. For the purposes of our simulation study, we make the further assumption that the CP makes the decision (on whether or not to invoke the API) only once at the beginning of each video session, and does not modify that decision mid-stream.

C. Price and Revenue Models

The price charged by the ISP to reserve dynamic bandwidth requested by the CP via the API is assumed to be a function of the access link load, in-line with "congestion-based pricing" schemes that have been used in the literature. Several researchers have used a two-tier pricing structure based on "peak" and "off-peak" hours; in this work we instead choose a pricing structure in which the price of the resource changes as a continuous function of its availability. A convenient pricing function is the exponential, which has been used by other researchers [23]. We therefore set the price (per Kbps-per-second) high when the spare capacity (link rate minus load) is low, and assume it to fall exponentially as the spare capacity increases. Mathematically, the price is given by:

$$P = \lambda e^{-\delta x}, \tag{1}$$

where P is the spot price of a unit of bandwidth (i.e. for 1 Kbps over a 1-second interval), x is the variable denoting fraction of available link capacity (computed by the ISP using an exponential moving average), λ is a constant corresponding to the peak spot-price (we use $\lambda = 0.5$ cents-per-Kbps-per-sec in our simulations), and δ is a constant corresponding to the rate at which the spot price of bandwidth falls with available capacity x. Our simulations will use $\delta = 0$ corresponding to constant bandwidth-price invariant to load, and $\delta = 2, 5$ corresponding to different degrees of congestion-based pricing. We emphasize that the pricing model is applied by the ISP to charge the CP; users neither see this price, nor change their behavior as a consequence of this price.

The use of the API to increase video QoE for users is expected to lead to some (long-term) returns for the CP. Several studies, e.g. [24], [25], have shown that improved QoE increases user-engagement and user-retention. Putting a price on this is however tricky. The model we use is based on the observation that user-engagement seems to fall rapidly with QoE-decay - this is borne out in several large-scale studies (e.g. Figures 2b, 11a, 12, 13 in [25]) that show that the fraction of content viewed (an indicator of user-engagement) falls very

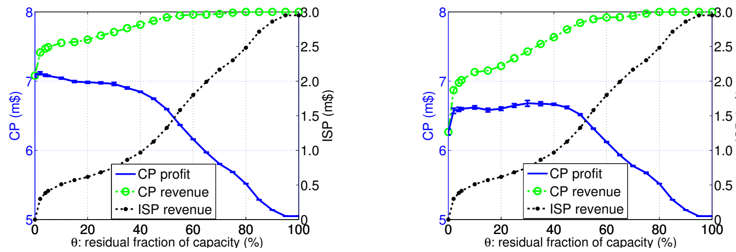

(a) (b)

Fig. 2. Flat-rate pricing ($\delta = 0$) with: (a) $\epsilon = 10$, and (b) $\epsilon = 20$.

steeply as rebuffering rates increase from 0 to 0.2 events-per-minute, by which time most of the harm is done; subsequent increase in rebuffering rates only marginally reduces content-viewing time. This leads us to approximate the CP's revenue as an exponential function of QoE:

$$R = \mu e^{-\epsilon y}, \quad (2)$$

where R is the overall revenue made by the CP (over a stipulated time-period, chosen to be 12 hours in our simulation study corresponding to the length of our traffic trace), y is the fraction of the user's streaming video flows that are deemed "poor quality" (a video flow that does not get its required bandwidth for 10% or more of its duration is deemed to be of poor quality), μ is a constant representing the potential revenue the CP can make if video quality were always perfect (for our simulation study we use $\mu = 8$ million dollars over the 12-hour period, based on YouTube's stated revenue of $5.8b in 2012), and ϵ is the rate at which the CP's revenue falls as a function of QoE degradation y. For our simulation study we will use $\epsilon = 10, 20, 30$ corresponding to different degrees to which degraded QoE affects the CP's revenue.

IV. Simulation Evaluation and Results

Our objective is to evaluate if our proposal for dynamic fast-lanes can lead to a win-win-win situation for the CPs, ISPs, and end-users. We therefore take a 12-hour traffic trace from our University, and feed it into our simulation that applies dynamic fast-lanes with the price and revenue models described above. In addition to varying the parameters of the two models, we also investigate the impact of the two control-knobs; (i) θ, the threshold on residual link capacity at which a CP invokes the API, and (ii) α, the fraction of link capacity a user contributes towards fast-lanes, on overall benefits.

A. Simulation Trace Data

The trace data for the simulation was obtained from our university's web cache. Each row in the trace contains flow-level information such as date and time of arrival, duration (in milliseconds), volume of traffic (in bytes) in each direction, the URL, and the content type (video, text, image). The logs span a 12 hour period (12pm-12am) on 16th March 2010, comprising 10.78 million flows and 3300 unique end-user clients. Of these flows, $11,674$ were video flows (predominantly from YouTube, identified by the content type field), $9,799$ were elephant flows (defined as transfers of size greater than 1 MB), and the remaining 10.76 million flows were mice (defined as transfers of size 1 MB or less, representative of web pages). Though the number of flows is largely skewed towards mice, in terms of traffic volume the three flow types constituted roughly equally (32%, 32% and 36% respectively) to the total traffic.

In terms of video rates, 98% of video flows required less than 5 Mbps, and only 0.2% of the flows required more than 10 Mbps; in terms of duration 90% of the video flows last under 3 minutes, and only 1% of the flows last for longer than 10 minutes. For completeness, we note that the file transfer size of elephant flows exhibits a heavy tail, with 99% of the flows transferring under 100 MB, and the maximum file size was about 1 GB; further, 93% of the mice flows complete their transfers within 1 second, and about 0.3% of the flows transferred more than 300 KB. The characteristics of our traffic trace are consistent with prior findings [26].

B. Simulation Methodology

We developed a native simulation that reads the flows attributes (start time, duration, type, rate/volume) and injects them into the slotted simulation. Flows are serviced slot-by-slot (a slot is of duration 1 second) over a broadband access link of capacity 100 Mbps. This access link emulates a collection of households, each having an average DSL connection of 2-10 Mbps. Note that the traffic trace is static, and user-demand pattern is neither assumed nor required to change since our pricing mechanism is largely transparent to the end-user. The video flows that are accommodated by the API – assumed to be constant bit rate – are allocated their own queue, while the other flows (mice and elephants) and those video flows that are denied by the API share a best-effort queue. Within the best-effort queue, the mice flows (that transfer less than 1 MB) are given their bandwidth first, since they are typically in the TCP slow-start phase. The remaining bandwidth is divided fairly amongst the video and elephant flows, because such flows are usually in the TCP congestion avoidance phase. The scheduling is work-conserving, so any bandwidth unused by the reserved bandwidth queues are given to the best-effort queue.

C. Performance Results

1) Flat-Rate Pricing Model: We first present results for the flat-rate pricing model, i.e. $\delta = 0$ in Eq. (1), and depict the results in Fig. 2. We start with user-parameter α set to 1, meaning that the fast-lane can utilise the full access link capacity. The ISP revenue (equal to CP cost) is shown by the dotted black line, and is in units of millions of dollars (over the 12-hour period of our trace), shown on the y-axis on the right-side. It is seen that when the threshold θ is zero

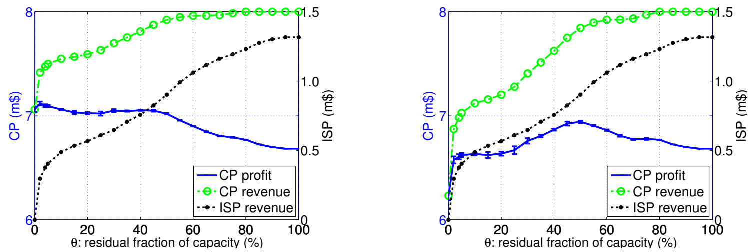

(a) (b)

Fig. 3. Congestion-based pricing ($\delta = 2$) with: (a) $\epsilon = 10$, and (b) $\epsilon = 20$.

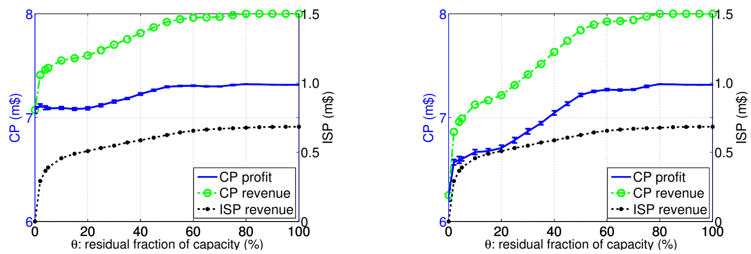

(a) (b)

Fig. 4. Congestion-based pricing ($\delta = 5$) with: (a) $\epsilon = 10$, and (b) $\epsilon = 20$.

(x-axis), the ISP makes no revenue, since the CP never calls the API (recall that the CP only calls the API if the fractional residual capacity on the access link falls below threshold θ). As the CP's threshold θ increases, the ISP's revenue increases roughly linearly as well, peaking at about \$3m. Meanwhile, the CP's revenue (dash-dotted green curve, left-side y-axis) also increases with threshold θ. When $\theta = 100\%$, the CP makes the API call to the ISP to reserve bandwidth for every video flow; this ensures that no video flow experiences QoE degradation (in other words y=0 in Eq. 2), and the CP's revenue therefore takes on the maximum possible value of $\mu = \$8\text{m}$ dollars. As the CP reduces their threshold θ, the API gets called only for higher network congestion states, and a larger fraction of video flows experience QoE degradation, thereby reducing user-engagement – indeed, as θ approaches zero, the number of "unhappy" flows (that experience degradation for at least 10% of their duration) increases to $y = 1.2\%$. The negative consequence of this reduction in video QoE on the CP's revenue depends on the parameter ϵ: revenue drops to \$7m when ϵ=10 (Fig. 2(a)), and to \$6.2m when ϵ=20 (Fig. 2(b)).

The CP's profit (revenue minus cost) is depicted by the solid blue lines in the plots (confidence interval in the form of standard deviation is also shown as error bars, and is found to be very small). The CP's gain depends highly on the sensitivity of user-engagement on video quality (i.e. parameter ϵ). When ϵ=10, we find that CP profit declines with threshold θ (Fig. 2(a)), implying that there is little incentive for the CP to ever call the API; on the other hand, when $\epsilon = 20$, CP profit increases with threshold θ and then decreases, peaking at around θ=40%. This suggests that the CP gains from calling the API only when the residual capacity falls below 40%, and is better off transmitting video as best-effort traffic otherwise.

2) Congestion-Based Pricing Model: We saw above that flat-rate pricing gives limited benefit to the CP, and the ISP may therefore employ a congestion-based pricing model, i.e. one that reduces price when residual capacity is more abundant, as an enticement to get CPs to call the fast-lane API more often. We therefore consider the exponentially-falling price model of Eq. (1) with the exponent δ taking values 2 and 5, shown respectively in Figures 3 and 4.

Once again, ISP revenue starts at zero when $\theta = 0$ (since the CP does not call the API in this case), and rises steadily as the CP increases their threshold θ. Not surprisingly, the maximum revenue for the ISP, when the CP calls the API for every video flow (i.e. $\theta = 100\%$), is lower in this case than before, being under \$1.5m for $\delta = 2$ (Fig. 3) and under \$0.7m for $\delta = 5$ (Fig. 3), due to the exponentially falling price of bandwidth as availability increases. Also not surprisingly, the CP's maximum revenue is again \$8m at $\theta = 100\%$ since all video flows experience perfect QoE. As θ reduces, the CP's profit (revenue minus cost) depends on the relative values of the exponents in the price and revenue models: the CP's profit is found to be decreasing with θ when $\delta = 2$, $\epsilon = 10$ (Fig. 3(a)), indicating that in this regime the CP does not increase profits by using fast-lanes. When $\delta = 5$ (corresponding to a rapid decay in fast-lane price with the available bandwidth) and $\epsilon = 10$ or 20 (Figures 4(a) and 4(b) respectively), the CP gains by always using fast-lanes. In the case when $\delta = 2$, $\epsilon = 10$ (Fig. 3(a)), the CP maximizes profit by setting threshold $\theta = 50\%$ – for lower θ, the ISP's price rise is too steep to be offset by revenue growth, whereas larger θ gives marginal revenue growth for the CP.

Based on the above results (and on those from numerous other parameter setting not included here due to space constraints), we believe that for given revenue model parameters (μ, ϵ), which the CP can deduce from long-term user-behavior, it is possible to find appropriate price model parameters (λ, δ) that lead to a win-win situation where both the ISP and CP increase their profits. We believe that market forces will nudge prices towards this region where ISPs have an incentive to offer

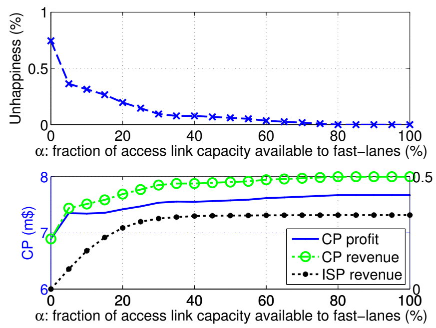

Fig. 5. Impact of α on end-user video flow quality and ISP/CP revenues.

dynamic fast-lanes and CPs the incentive to use them.

3) End-User Benefit: So far the results have focused on the economic incentives of fast-lanes for the ISP and the CP. The scheme is largely transparent to end-users, who are not expected to change their behavior. They do however have one control-knob: the fraction $\alpha \in [0, 1]$ of their access link capacity that they allow the ISP to carve fast-lanes from. The larger they set their α to, the more likely the ISP will be able to successfully create the fast-lanes when demanded by the CPs. To evaluate the impact on video QoE, we plot in Fig. 5 the impact of α on end-user QoE (top plot), and the corresponding ISP/CP revenues (bottom plot), for the case when the CP always calls the API for fast-lane creation (i.e. $\theta = 100\%$). The QoE is expressed in terms of the fraction of video flows that are "unhappy", i.e. fail to get their required bandwidth for at least 10% of their duration. Unsurprisingly, unhappiness falls monotonically in α (e.g. unhappiness falls by 91.5% when $\alpha = 0.5$, and becomes zero at $\alpha = 1$), suggesting that if video QoE is paramount to the user, they should set their α as high as possible (recall that end-users do not pay for fast-lanes; indeed future study will consider users getting a subsidy from the ISP for allowing fast-lanes). Simultaneously, it is seen that the ISP and CP both increase profits with increasing α, demonstrating the win-win-win for all.

In spite of the increased video QoE for end-users for higher α settings, we believe it is important to provide this control-knob to users, particularly those who want the benefits of network neutrality for non-video traffic. We have shown in our prior work [12] that setting α too high (close to unity) runs the risk of allocating all the bandwidth to video traffic, potentially reducing bandwidth for elephant transfers (e.g. peer-to-peer) and increasing delays for small transfers (e.g. web-page loads). Our proposal gives users the flexibility to determine the setting of the single knob most suited to their preferences.

V. Conclusions

The current debate surrounding Internet fast-lanes focuses on relatively static settings, which we believe are unlikely to be palatable to all three entities involved. In this paper, we have advocated an approach that creates dynamic fast-lanes, which are initiated by the CP, operated by the ISP, and sanctioned by the end-user. Using simple but representative models for fast-lane pricing by ISPs and associated revenue-generation for CPs, we have shown that it is possible for our approach to lead to a win-win-win situation for all the three entities, specifically for end-users who have to neither pay nor change their behavior, and can opt out at any point. While we are not necessarily saying that our solution addresses all aspects of the complex problem of Internet fast-lanes, we believe it is at least worth considering seriously given that current proposals are, understandably, stymied.

References

[1] Wall Street Journal. (2014, Apr) FCC to Propose New 'Net Neutrality' Rules. http://online.wsj.com/news/articles/SB10001424052702304518704579519963416350296.
[2] The New Yorker. (2014, Apr) Goodbye, Net Neutrality; Hello, Net Discrimination. http://www.newyorker.com/tech/elements/goodbye-net-neutrality-hello-net-discrimination?currentPage=all.
[3] Cisco Internet Business Solutions Group. (2012) Moving Toward Usage-Based Pricing. http://goo.gl/pj8Gf.
[4] The European Telecom. Network Operators' Association. (2012) ITRs Proposal to Address New Internet Ecosystem. http://goo.gl/VutcF.
[5] M. Nicosia et al., "Rethinking Flat Rate Pricing for Broadband Services," Cisco Internet Business Solutions Group, White Paper, Jul. 2012.
[6] The New York Times. (2014, Apr) F.C.C., in a Shift, Backs Fast Lanes for Web Traffic. http://www.nytimes.com/2014/04/24/technology/fcc-new-net-neutrality-rules.html?_r=1.
[7] GIGAOM. (2014, May) Opposition to FCCs controversial "fast lane" plan is gaining steam. https://gigaom.com/2014/05/08/opposition-to-fccs-controversial-fast-lane-plan-is-gaining-steam/.
[8] ——. (2014, May) Amazon, Netflix and tech giants defend net neutrality in letter to FCC. https://gigaom.com/2014/05/07/amazon-netflix-and-tech-giants-defend-net-neutrality-in-letter-to-fcc/.
[9] FinancialTimes. (2014, Nov) Netflix wants to put Comcast genie back in fast lane bottle. http://www.ft.com/cms/s/0/0bc54d54-639e-11e4-8216-00144feabdc0.html#axzz3KreEKl59.
[10] The Washington Post. (2014, Sep) AT&Ts fascinating third-way proposal on net neutrality. www.washingtonpost.com/blogs/the-switch/wp/2014/09/15/atts-fascinating-third-way-proposal-on-net-neutrality/.
[11] Fox2Now. (2014, Oct) AT&T wants you to design your own Internet fast lane. http://fox2now.com/2014/10/13/att-wants-you-to-design-your-own-internet-fast-lane/.
[12] V. Sivaraman et al., "Virtualizing the Access Network via Open APIs," in *Proc. ACM CoNEXT*, Dec 2013.
[13] S. Sen et al., "A Survey of Smart Data Pricing: Past Proposals, Current Plans, and Future Trends," *ACM Comp. Surveys*, vol. 46, no. 2, 2013.
[14] N. Economides and J. Tag, "Network neutrality on the Internet: A two-sided market analysis," *Inf. Econ. and Policy*, vol. 24, pp. 91–104, 2012.
[15] E. Altman, A. Legout, and Y. Xu, "Network Non-neutrality Debate: An Economic Analysis," in *Proc. IFIP Networking*, Spain, May 2011.
[16] J. Wang, R. T. B. Ma, and D. M. Chiu, "Paid Prioritization and Its Impact on Net Neutrality," in *Proc. IFIP Networking*, Norway, Jun 2014.
[17] R. T. B. Ma and V. Misra, "The Public Option: A Nonregulatory Alternative to Network Neutrality," *IEEE/ACM Transactions on Networking*, vol. 21, no. 7, pp. 1866–1879, 2013.
[18] L. Zhang and D. Wang, "Sponsoring Content: Motivation and Pitfalls for Content Service Providers," in *Proc. IEEE INFOCOM workshop on Smart Data Pricing*, Canada, Apr/May 2014.
[19] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring Mobile Data: An Economic Analysis of the Impact on Users and Content Providers," in *Proc. IEEE INFOCOM*, Hong Kong, Apr/May 2015.
[20] J. Matthews, C. Russell, and V. Sivaraman, "SDN API for Access Network Virtualization," in *Demo at IEEE LCN*, Oct 2013.
[21] K. Lai and M. Baker, "Measuring Link Bandwidths Using a Deterministic Model of Packet Delay," in *Proc. ACM SIGCOMM*, Aug/Sep 2000.
[22] C. Dovrolis and P. Ramanathan, "Packet-Dispersion Techniques and a Capacity-Estimation Methodology," *IEEE/ACM Transactions on Networking*, vol. 12, no. 6, pp. 963–977, 2004.
[23] Y. A. et al., "An Opportunity Cost Approach for Job Assignment and Reassignment," *IEEE Transactions on Parallel and Distributed Systems*, vol. 11, no. 7, pp. 760–768, Jul 2000.
[24] F. Dobrian et al., "Understanding the Impact of Video Quality on User Engagement," in *Proc. ACM SIGCOMM*, Aug. 2011.
[25] A. Balachandran et al., "Developing a predictive model of quality of experience for internet video," in *Proc. ACM SIGCOMM*, Aug. 2013.
[26] S. Ramachandran. (2010) Web metrics: Size and number of resources. https://developers.google.com/speed/articles/web-metrics.