An Economic Model for a New Broadband Ecosystem Based on Fast and Slow Lanes

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

***Abstract*—Today's residential broadband ecosystem is in stasis – Internet Service Providers (ISPs) suffer from low margins and flat revenues, Content Service Providers (CSPs) have unclear incentives to invest in broadband infrastructure, and users have limited dimensions (speed/quota) in which to compare broadband pricing. This paper explores the use of *service quality* capabilities, in the form of fast- and slow-lanes, for overcoming this stasis. We propose an architecture in which all entities have a say – CSPs request dynamic fast- and slow-lane creation for specific sessions, ISPs operate and charge for these lanes, and users control their broadband bandwidth available to such lanes. We develop an economic model that balances fast- and slow-lane pricing by the ISP with the returns for CSPs and service quality improvement for end-users, and evaluate the parameters of our model with real traffic traces. We believe our proposal based on dynamic fast- and slow-lanes can represent a win-win-win situation for ISPs, CSPs, and end-users alike, and has the potential to overcome the current stagnation in broadband infrastructure investment.**

I. INTRODUCTION

Residential data traffic is growing at 40% per annum, while the average fixed-line broadband bill has been relatively flat for many years [1]. Internet Service Providers (ISPs) have argued that in order to sustain and upgrade their infrastructure to cope with growing traffic volumes, new "two-sided" revenue models are necessary to help narrow the gap between their cost and revenue [2]. Under such a model, a Content Service Provider (CSP), such as Netflix, YouTube, or Hulu, would pay the ISP to create a "fast-lane" to prioritise their traffic over other content, improving quality-of-experience (QoE) for their end-users, leading to increased monetization by virtue of greater user engagement and retention. This new source of revenue for ISPs from CSPs is expected to lead to investment in improving broadband infrastructure.

Understandably, Internet fast-lanes are viewed with suspicion by the public, who view this as giving license to ISPs to block or throttle arbitrary traffic streams of their choice without regard to consumer interest [3], in violation of the so-called "network neutrality" principle. This has led to raging debates amongst policy-makers, activists, economists and researchers [4], [5] on the pros and cons, with the FCC in the US having changed its stance multiple times – currently leaning towards favoring consumers by disallowing fast-lanes.

One of the assumptions that seems to be built into this debate is that the fast-lane negotiation is between the ISP and the CSP, with the consumer having no voice in the matter. Moreover, it is implicitly assumed that the fast-lane prioritization is done statically over a long period of time, and applied in bulk to all traffic from that CSP. These assumptions were indeed exemplified in Netflix's peering payment to Comcast in early 2014 which was reported to be worth \$15-20 million a year [6], to improve Netflix experience for Comcast subscribers. Clearly, users were irate at such back-room deals from which they were shut out, and concerned about the disadvantages for smaller CSPs who do not have the capital to pay the ISP up-front for prioritization of their traffic.

In this paper we consider a new model for fast-lanes that addresses the above two concerns. The *first* aspect of our approach is that fast-lanes are created **dynamically** for specific sessions, triggered by APIs that are open for any CSP to invoke; if accepted, the ISP charges the CSP a micro-payment for the fast-lane, based on duration and bandwidth – pricing model discussed later. The open nature of the API makes the playing-field level for all CSPs, and the micro-payment rather than bulk-payment, ensures that a CSP can invoke fast-lanes in line with their business model, such as only for premium users or during congested periods, and keep costs elastic rather than up-front. The *second* aspect of our approach gives **control to users** to limit the fraction of their broadband capacity that can be used towards fast-lanes; this fraction α, if set to 0, effectively disables fast-lanes for that household and preserves network neutrality, while a setting of 1 gives the ISP full freedom to create fast-lanes at their discretion for that house. An intermediate value of α, say 0.8, gives the ISP access to at most 80% of the broadband bandwidth for fast-lanes, leaving at least 20% for best-traffic that does not request special lanes. This knob is a simple interface for the lay-user to understand, yet lets them customize the extent to which the benefits of fast-lanes can be traded-off against net-neutrality.

We have demonstrated the technical feasibility of API-driven dynamic fast-lane creation, using software defined networking (SDN) technology, in a recent paper [7]. Our goal in the current paper is to explore the economic incentives for this approach, and to show that if tuned appropriately, it can result in a win-win-win situation for end-users, CSPs, and ISPs. Our specific contributions are:

- We show how value flows in this new ecosystem: some CSPs such as Netflix, Youtube pay ISPs for fast-lanes, predominantly for video streaming; ISPs in-turn pay other CSPs such as Dropbox and Zipcloud to offload bulk-transfers to slow-lanes; users can set their α-knob high to get better experience for both video streaming and web-browsing; and CSPs in-turn can increase revenue from improved user-experience. This cycle benefits all entities.

- We show using simulations of traffic traces taken from real networks that user's video-experience improves with fast-lanes, at the cost of increasing web-browsing latencies. We then show that complementing fast-lanes with slow-lanes (that off-load bulk-transfers) improves web-browsing performance, providing incentives to the user to contribute a larger fraction α of their broadband link capacity, which is needed for economic sustainability of this ecosystem.
- We consider realistic pricing models for fast- and slow-lanes, as well as various load conditions under which the fast- and slow-lanes are created, to show via simulation that both ISPs and CSPs can increase their per-user revenue if they appropriately tune their pricing parameters.

The rest of this paper is organised as follows: §II reviews pricing models of user-facing, CSP-facing and two-sided that are relevant to our work. §III outlines our system operation, and choice of ISP-pricing and CSP-revenue models. A simulation study with real trace of over 10 million flows is conducted in §IV, and the economic benefits for the ISP/CSP are studied under various parameter regimes. §V concludes the paper.

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 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. Several ISPs around the world are offering newer pricing schemes such as usage-based pricing, where fee paid is proportional to the volume of data used), tiered pricing that comprising a fixed quota charge and any overage charges for exceeding the quota, and time-of-day pricing where charges are higher during peak-hour usage compared to off-peak hours. Dynamic pricing includes schemes such as day-ahead-pricing whereby charges for the next day are guaranteed the previous day, and congestion-based pricing in which users pay higher prices during higher congestion levels. An excellent survey of the different pricing models aimed at end-users is given in [8].

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 the end-users and CSPs. The work in [9] studies a two-sided non-net-neutral market, and takes into account the 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 CSP would be beneficial to either of them. The work in [10] considers a model comprising a monopoly ISP, a set of CSPs, and end-users. Focusing on the utility of the ISP/CSPs 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.

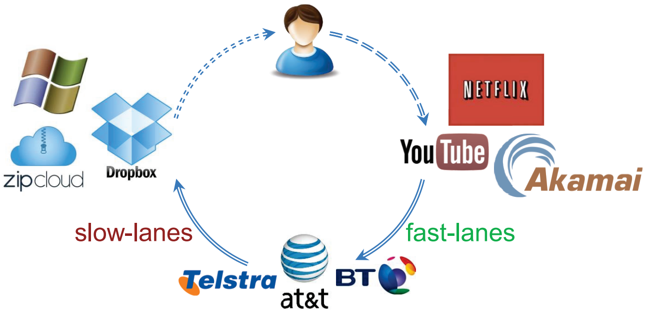

Fig. 1. Broadband economic value chain.

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- and slow-lanes that are created and destroyed on-the-fly, wherein CSPs/ISPs 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 [11], [12] – in this model, the end-user pays a lower fee to the ISP due to CSP induced subsidies, for e.g. Facebook traffic being considered "in-network" and not counting towards the user's quota is an example of this. The CSP can benefit by attracting more traffic from the end-user, while the ISPs can reduce churn 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 CSPs and ISPs, we study the efficacy of a model that permits paid-prioritization at much smaller time-scales, i.e. at per-session granularity.

III. NEW BROADBAND ECOSYSTEM

In Fig. 1 we illustrate the value chain in the new broadband ecosystem. The video CSP, e.g. Netflix, YouTube, generates revenue from users via subscription fees or advertisements, and benefits from fast-lanes by increasing engagement and repeat viewership (technical details in §III-A and economic model in §III-D). The video CSP in turn makes a micro-payment to the ISP for the fast-lane (pricing model in §III-C). We find that creating fast-lanes for video can degrade performance for mice flows, for e.g. web-page loads, and so in §III-B we argue that the ISP pay certain CSPs such as Dropbox and Zipcloud to offload large bulk-transfers on to slow-lanes (pricing model in §III-C). This ecosystem is discussed in detail next, followed by a quantitative evaluation in the following section.

A. Dynamic Fast-Lanes for Video Streams

Our proposal for fast-lanes differs from earlier approaches in being *dynamic* and *open*. The ISP exposes an API, available for any CSP to call, to create a fast-lane for a specific stream. The technical specification of the API, i.e. specifying the end-points of the traffic stream, bandwidth requirement, and duration, and its implementation using software defined networking (SDN) technology, can be found in our prior work [7]. We note that a CSP has full control over API invocation – if network performance is adequate, or if the customer is low-value, the CSP can at their discretion send their traffic as best-effort, much like the way it is today. This gives the CSP 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 CSPs.

If the fast-lane creation is successful, the CSP will make a "micro-payment" to the ISP (pricing model in §III-C). Note that the ISP has every incentive to accept fast-lane calls from CSPs if capacity is available, but can do so only if the user setting permits this. As mentioned earlier, the user has a control knob α that they can set in the range $[0, 1]$, and denotes the fraction of their broadband link capacity that they allow the ISP to carve fast-lanes from. A user wishing to stay network neutral can set their α-knob to 0 to opt out of the scheme, while a user who wants to benefit from fast-lanes can set it to any fractional value up to 1. The ISP can provide an incentive, say in the form of a subsidy, to users for setting their α-knob close to 1, but this is outside the scope of the current work. We will show that fast-lanes can enhance the user's video experience, as also web-browsing performance, provided they are used in conjunction with slow-lanes, as described next.

B. Dynamic Slow-Lanes for Bulk Transfers

Common experience and our evaluations in the next section show that web-browsing experience is degraded when done in parallel with video streams and large downloads. Moving video sessions onto fast-lanes runs the risk that web-browsing performance degrades even further since the "mice" flows share the best-effort queue with bulk transfer "elephant" flows. We therefore propose that bulk transfers be moved to "slow-lanes" that get lower priority than best-effort. To enable this, the ISP opens an API for bulk-transfer CSPs such as Dropbox and Zipcloud to indicate that they are doing a large transfer, and to specify "elasticity" in terms of the delay bound that this transfer can tolerate. As an incentive for calling this API, which will free up network capacity for web-browsing and video traffic, the ISP makes a micro-payment to the bulk-transfer CSP (payment model in §III-C). As we will see in our evaluation section, offloading bulk transfers to slow-lanes can cost the ISP money, but protects video and mice quality, ensuring that the user does not turn their α-knob low which would prevent the ISP from earning revenue from video CSPs.

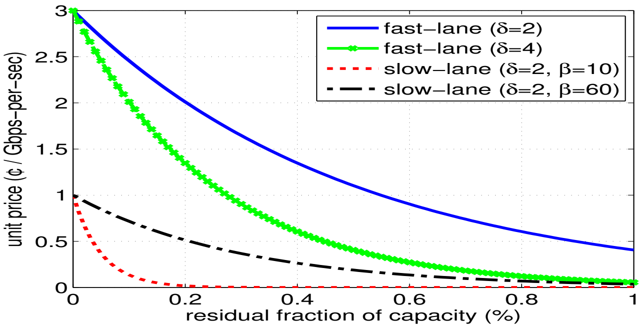

Fig. 2. Price of fast- and slow-lanes.

C. ISP Revenue from Fast- and Slow-Lanes

The price charged/paid by the ISP to special lanes requested by the CSP 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. We 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 [13]. We therefore set the price per Gbps-per-second high when the spare capacity on the broadband link (link rate minus load) is low, and assume it to fall exponentially as the spare capacity increases. Mathematically, the price of fast-/slow-lanes is given by:

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

where P is the spot price of a unit of bandwidth, i.e. for 1 Gbps over a 1-second interval, x is the variable denoting fraction of available broadband link capacity computed by the ISP using an exponential moving average, λ is a constant corresponding to the peak spot-price, and δ is a constant corresponding to the rate at which the spot price of bandwidth falls with available capacity x. It is natural to expect that the ISP prices the fast-lane (the amount they charge the video CSP) higher than the slow-lane (the amount they pay the bulk-transfer CSP). In our study we will be using $\lambda_f = 3\lambda_s$, and $\delta_f = 0.01\beta\delta_s$, where β is the "elasticity" specified by the bulk transfer, corresponding to the factor by which the bulk-transfer is willing to stretch in time compared to a baseline in which the bulk-transfer has access to the entire access link capacity. Note that this payment model incentivizes a bulk-transfer CSP to choose as high an elasticity parameter β as possible to maximize payment from the ISP. Fig. 2 illustrates for these parameters how the price of fast and slow lanes falls steeply with increasing spare capacity.

We emphasize that the pricing model is between the ISP and CSPs, and is neither visible to users nor expected to change their behavior. The ISP's net profit from CSPs is then the revenue from fast-lanes minus the payment for slow-lanes:

$$\text{ISP Profit} = \sum_i (\lambda_f e^{-x\delta_f} . F_i^{size}) - \sum_j (\lambda_s e^{-x\delta_s} . S_j^{size}), \tag{2}$$

where F_i^{size} and S_j^{size} are the size of the flows (in Gb) admitted for fast-lanes and slow-lanes respectively.

D. CSP Revenue Enhancement

The bulk-transfer CSPs have every incentive to call the ISP API for slow-lanes, since they get a payment from the ISP. The video CSP, on the other hand, has to balance the costs of fast-lanes against the returns they obtain in terms of increased revenue from consumers. Several studies [14], [15] 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, for example Figures 2b, 11a, 12, 13 in [15] show that the fraction of content viewed, which is an indicator of user-engagement, falls very 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 CSP's revenue as an exponential function of QoE:

$$R = \mu e^{-\epsilon y}, \tag{3}$$

where R is the overall revenue made by the CSP 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" – we define this as a video flow that does not get its required bandwidth for 10% or more of its duration, μ is a constant representing the potential revenue the CSP can make if video quality were always perfect; for our simulation study we use $\mu = 3$\$ over the 12-hour period, based on Google's average revenue per user (ARPU) of \$45 in Q1 2014 and YouTube's 6% share of Google's revenue in 2014, being scaled to our 10 houses each having an average 3 users, and ϵ is the rate at which the CSP's revenue falls as a function of QoE degradation y. For our simulation study we will use $\epsilon = 2$, i.e. an increase of 1% in unhappy video flows drops revenue by 10%.

IV. EVALUATION USING TRAFFIC TRACE

We now apply our pricing model for fast- and slow-lanes to a traffic trace taken from a campus network, and explore the parameter space to find regions where all three entities benefit.

A. *Simulation Data and Methodology*

Trace data: The trace data was taken from the campus web cache, and contains flow logs on date/time, duration in milliseconds, volume of traffic in bytes, the URL, and the content type, i.e. video, text, image. We used a 12-hour period from the logs, 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 (greater than 10 MB), and the remaining 10.76 million flows were mice (representative of web pages). In terms of traffic volume the three flow types contributed roughly equally (32%, 32% and 36% respectively) to the total traffic.

Simulation Methodology: We developed a native simulation that reads the flow attributes and injects them into the slotted simulation. Flows are serviced slot-by-slot – a slot is of duration 1 second – over an access network emulating a collection of 10 households, each with broadband capacity of 10 Mbps. The video flows that are accommodated by the API are allocated their own "fast-lane" queue, while bulk-transfer flows accommodated by the corresponding API are allocated their own "slow-lane" queue. Mice flows, and all other flows for which either the CSP does not invoke the API or the ISP rejects the API, share a best-effort queue. Within the best-effort queue, the mice flows 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 these 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.

Fast-Lane Strategy: A video CSP can call the fast-lane API at-will based on their business model. To make the study tractable, we assume that they invoke the fast-lane API only when the available bandwidth falls below fraction θ_f of the access link capacity, i.e. when bandwidth is scarce. We believe this to be a reasonable strategy, and also practically feasible since CSPs actively monitor bandwidth anyway; currently using them to adapt their video coding rates. The video CSP can adjust parameter θ_f in the range $[0, 1]$ to make their use of fast-lanes conservative (low θ_f) or aggressive (high θ_f).

Slow-Lane Strategy: The bulk-transfer CSP has a financial incentive to call the slow-lane API for each large download, but the ISP may not always be inclined to accept the request, since they have to pay the CSP for slowing their transfer. For this evaluation we make the assumption that the ISP accepts the slow-lane request only when the available bandwidth falls below fraction θ_s of broadband link capacity. The ISP can adjust this parameter in $[0, 1]$ to either conservatively (low θ_s) or aggressively (high θ_s) off-load bulk transfers to slow-lanes. In this work we will tune this parameter based on the revenue that the ISP obtains from fast-lanes – the ISP therefore tracks parameter ρ_f corresponding to the "fast-lane utilization", measured as the exponentially averaged fraction of broadband link capacity (for a consumer) assigned to fast-lanes, and uses this to adjust θ_s. A natural consequence of this approach, whereby $\theta_s = \rho_f \leq \alpha$, is that a subscriber who contributes a low fraction of their broadband capacity to this ecosystem avails of reduced benefits from slow-lane off-loading.

TABLE I
ROLE OF ESSENTIAL PARAMETERS

Parameter	Description
α	Control knob set by the user to specify the fraction of their broadband link capacity that can be used towards fast-lanes.
β	Elasticity factor by which a CSP making the slow-lane call is willing to stretch the bulk transfer.
x	Instantaneous fraction of available bandwidth on the link.
y	Aggregated fraction of video streams that are deemed "poor quality", i.e. not getting their required bandwidth for at least 10% of its duration.
x and y	x and y are inversely related. Scarce bandwidth (small x) will generally lead to poor quality (large y).
θ_f	A configurable lower threshold on available bandwidth x below which a CSP invokes the fast-lane call.
θ_s	A configurable lower threshold on available bandwidth x below which the ISP admits the slow-lane call.
θ_f and θ_s	θ_s is set dynamically by the ISP depending on the usage of fast-lanes, which in turn depends on θ_f of various CSPs.

Model parameters: Table I summarizes the role of the essential parameters used by our model, which are measured or chosen by the ISP, CSPs and users respectively.

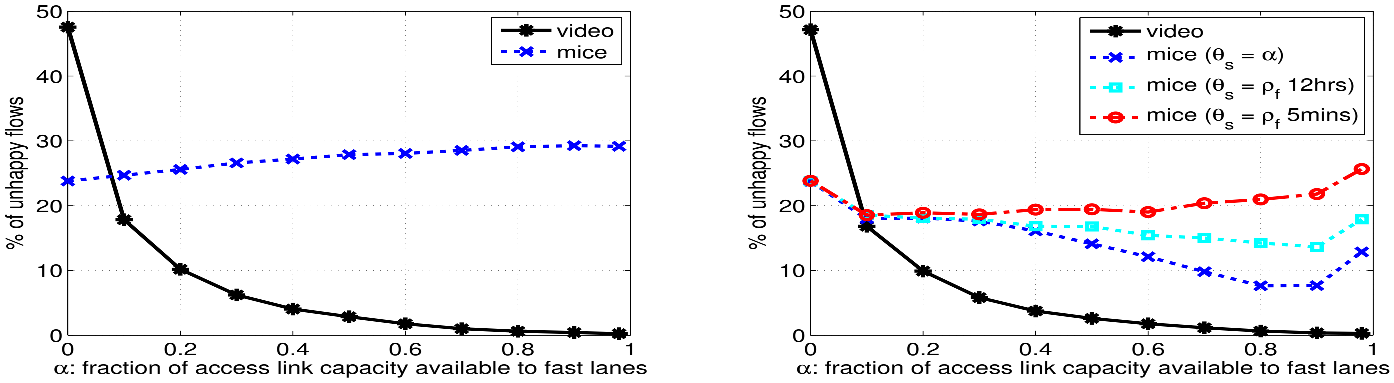

(a)

(b)

Fig. 3. End-user QoE when: (a) only fast-lanes are provisioned, and (b) both fast-lanes and slow-lanes are provisioned. ($\theta_f = 1$)

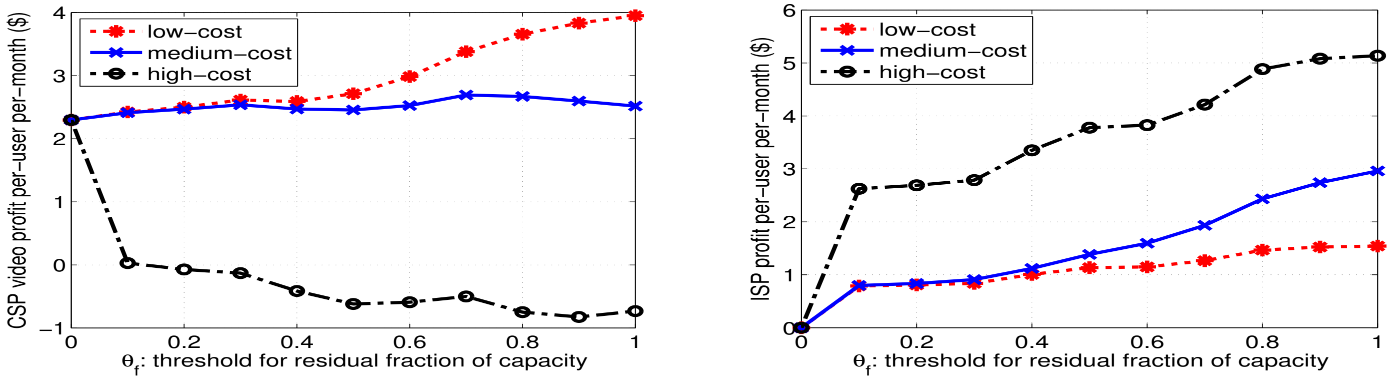

(a)

(b)

Fig. 4. Profit per-user per-month for: (a) video CSP, and (b) ISP. ($\alpha = 0.9$)

B. Performance Results

1) Benefit for the End-User: Our scheme is largely transparent to the 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. To evaluate the impact of α, we plot in Fig. 3(a) the end-user QoE when only fast-lanes are provisioned, and in Fig. 3(b) the end-user QoE when both fast-lanes and slow-lanes are provisioned, for $\theta_f = 1$, i.e. the video CSP invokes the API for every video session. The end-user QoE for video traffic is measured in terms of the fraction of video flows that are "unhappy", i.e. fail to obtain the required bandwidth for at least 10% of the time. As expected, the percentage of video flows that are unhappy falls monotonically with α, falling from 48% at $\alpha = 0$ to 6% at $\alpha = 0.5$, confirming that fast-lanes enhance video QoE at no cost to the end-user.

The QoE for mice traffic is measured in terms of the fraction of flows that do not complete, i.e. web-page does not load within 4 seconds. As shown by the dashed line in Fig. 3(a), the QoE for mice flows worsens with α in the presence of only fast-lanes. This is because a larger α reduces the bandwidth available for the best-effort queue to use, increasing the time needed for the mice flows to complete. However, the QoE is substantially better in the presence of both fast- and slow-lanes, as shown by the dashed lines in Fig. 3(b), attributed to the bandwidth freed up by the elasticity of bulk transfer flows. This improvement in the performance of mice flows depends on the extent to which the ISP admits slow-lanes. As stated earlier, we use an approach whereby the ISP limits slow-lane usage to match fast-lane usage so that costs do not exceed revenues – more specifically, the ISP sets the threshold θ_s on residual bandwidth fraction for slow-lane acceptance to equal ρ_f, the (exponentially averaged) fast-lane utilization. In the most optimistic case when the entire user's available bandwidth fraction α is used for fast-lanes, $\theta_s = \alpha$, and the percentage of mice flows that are unhappy drops to below 10%, constituting a lower bound, as indicated by the dashed blue line, second curve from the bottom in the figure. For our trace, when ρ_f is averaged dynamically at run-time at time-scales of 5-minutes and 12-hours for the curves shown in Fig. 3(b), slow-lane acceptance is reduced, causing mice performance to show more degradation.

2) Benefit for the CSP and ISP: The ISP employs a congestion-based pricing model given by (1). We consider three pricing models for fast-lanes, and accordingly for slow-lanes: (a) *high-cost* lanes corresponding to $\delta_f = 2$ and $\lambda_f = 10$ in which the unit price is set relatively high for the video CSP, (b) *medium-cost* lanes corresponding to $\delta_f = 2$ and $\lambda_f = 3$ that gives freedom to the video CSP to use fast-lanes whenever needed, and (c) *low-cost* lanes corresponding to $\delta_f = 0.5$ and $\lambda_f = 3$ in which the price is less sensitive to the load. In the following results, we use $\theta_s = \rho_f$ that is averaged every 5 minutes.

We plot in Fig. 4 the profit per-user per-month for the video

CSP and ISP as a function of the video CSP's threshold parameter θ_f. The latter's profit is computed as the revenue (Eq. 3) minus cost paid to the ISP (Eq. 1) and falls monotonically when the fast-lane is high-cost, as shown by the bottom curve in Fig. 4(a). This is because the price of the fast-lanes paid to the ISP outweighs the revenue obtained from increased user engagement. This scenario could pose an economic risk to the video CSP, who may choose to not call the API at all given the pricing model. The video CSP's profit when the fast-lanes are medium-cost is shown by the middle curve. In this case, the profit is maximized at \$2.76 per-user per-month when the threshold $\theta_f = 0.7$. Increasing the threshold any further reduces the profit, as the gain from higher user QoE is overridden by the expense incurred for using fast-lanes. Finally, when the fast-lane is low-cost, the top curve shows that the CSP profit increases monotonically because the low price encourages the CSP to call fast-lanes for every video session, and allowing it to capitalise on higher user engagement.

Focusing now on ISP profit, Fig. 4(b) shows that unsurprisingly the profit is seen to increase monotonically with the CSP threshold parameter θ_f for the three pricing models, and it is zero when $\theta_f = 0$. This is because the video CSP does not call the fast-lane API when $\theta_f = 0$, and alongside the slow-lanes are not created as well. The high-cost model provides the largest monetization opportunity for the ISP, but this may not be of interest to the video CSP as mentioned earlier. Both the medium- and low-cost models offer similar returns to the ISP until $\theta_f = 0.4$, following which the former outperforms the latter. The medium-cost model thus seems to be the most reasonable for both the ISP and CSP to use, and under this pricing scheme, the CSP's profit (see Fig. 4(a)) is maximized when $\theta_f = 0.7$, and at the same time earning the ISP a profit of nearly \$2 per-user per-month.

Based on the above results, and from numerous other parameter settings not included here due to space constraints, we believe that for given revenue model parameters (μ, ϵ), which the video CSP can deduce from long-term user-behavior, it is possible to find appropriate pricing model parameters (λ, δ) that will lead to a win-win situation for both the ISP and video CSP in terms of their profits. We believe that market forces will nudge prices towards this region where ISPs have an incentive to offer dynamic fast- and slow-lanes and CSPs the incentive to use them.

V. CONCLUSIONS

In this paper, we have explored the role that service quality can play – in the form of fast- and slow-lanes – to overcome the stasis in today's residential broadband ecosystem. We proposed an architecture wherein all three entities, i.e. CSPs, ISPs, and end-users, have a say. CSPs request the creation of fast- and slow-lanes dynamically for specific traffic streams, ISPs operate and monetize on these lanes, and end-users control the bandwidth made available to these lanes. We developed an economic model that balances fast- and slow-lane pricing by the ISP, with associated revenue generation for CSPs and QoE improvements for the end-users. The parameters of the economic model were evaluated using a real traffic trace. We believe that our approach can lead to a win-win-win situation for all the three parties, and is a solution worth considering seriously given that current proposals are, understandably, stymied.

REFERENCES

[1] Cisco Internet Business Solutions Group. (2012) Moving Toward Usage-Based Pricing. http://goo.gl/pj8Gf.

[2] The European Telecom. Network Operators' Association. (2012) ITRs Proposal to Address New Internet Ecosystem. http://goo.gl/VutcF.

[3] 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/.

[4] H. Habibi Gharakheili, A. Vishwanath, and V. Sivaraman, "Perspectives on Net Neutrality and Internet Fast-Lanes," *ACM CCR*, vol. 46, no. 1, Jan. 2016.

[5] 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.

[6] 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.

[7] V. Sivaraman et al., "Virtualizing the Access Network via Open APIs," in *Proc. ACM CoNEXT*, Dec 2013.

[8] 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.

[9] E. Altman, A. Legout, and Y. Xu, "Network Non-neutrality Debate: An Economic Analysis," in *Proc. IFIP Networking*, Spain, May 2011.

[10] 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.

[11] 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.

[12] 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.

[13] Y. Amir et al., "An Opportunity Cost Approach for Job Assignment and Reassignment," *IEEE Transactions on Parallel and Distributed Systems*, vol. 11, no. 7, p. 760768, 2000.

[14] F. Dobrian et al., "Understanding the Impact of Video Quality on User Engagement," in *Proc. ACM SIGCOMM*, Aug. 2011.

[15] A. Balachandran et al., "Developing a predictive model of quality of experience for internet video," in *Proc. ACM SIGCOMM*, Aug. 2013.