Can ISPs be Profitable Without Violating "Network Neutrality"?

Amogh Dhamdhere
Georgia Institute of Technology
amogh@cc.gatech.edu

Constantine Dovrolis
Georgia Institute of Technology
dovrolis@cc.gatech.edu

ABSTRACT

At the core of the network neutrality debate we find that ISPs, in particular the last-mile Access Providers (APs), are trying to find new ways to be profitable, despite the fact that their transit traffic has been dramatically increasing, while they continue to charge their customers a flat monthly price. In this paper, we consider a simple model of an AP that serves its users traffic from a number of Content Providers (CPs). The AP can communicate with the CPs through a Transit Provider (TP) or through settlement-free peering. We examine the profitability of the AP under a "baseline" model that is based on current practice, considering the heavy tailed variability in per-user traffic and in the popularity of different CPs. Further, we consider other strategies, such as usage-based pricing for heavy hitters, selective peering with popular CPs, and content caching. Our results indicate that an AP can be profitable without the risk of losing users and without violating "network neutrality", through selective peering with CPs and/or content caching.

Categories and Subject Descriptors

C.2.3 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Operations

General Terms

Economics, performance, measurement

Keywords

Network neutrality, ISP pricing, ISP peering

1. INTRODUCTION

The increasing penetration of broadband access, faster last-mile links, the rise of Internet video and peer-to-peer file sharing mean that residential and SOHO (Small Office, Home Office) users download increasingly more content. This content is delivered to users by Internet Service Providers (ISPs) that are known as Access Providers (APs). APs earn their revenues mostly from their users, and they incur costs to operate their network and to purchase upstream connectivity from transit providers. A much discussed trend in recent times is that APs are often not profitable, as the increasing transit traffic leads to escalating costs, while the intense competition in the access market and the commoditization of Internet access leads to falling prices, typically in the form of a flat monthly fee [6, 9, 12].

The APs see their profitability shrink as their role in the Internet ecosystem becomes simply to "move bits around" instead of providing end-to-end services such as IPTV or VoIP. On the other hand, content providers, "over-the-top" services and application providers, collectively referred to as Content Providers (CPs) in this paper, get all the attention recently (and often the profits as well..). This tension has led to the "network neutrality" debate. Despite the many articles in the popular press, articles written by economists and telecommunication policy experts [1, 7, 8, 11, 15, 16, 17], and by computer scientists [5, 14, 18], this debate is still highly misunderstood. We believe that this debate is not really about the "neutrality" of the Internet (a concept that is ill-defined, to say the least), but about the profitability of APs. At the end of the day, it is they that want to change the status quo by charging some CPs, discriminating traffic in priority classes, or entirely blocking certain flows. To understand the network neutrality debate, we need to understand both the economic structure and the traffic characteristics that APs need to work with.

In this paper, we approach this issue quantitatively using a simple model that captures the interactions between an AP, a transit provider, and a number of CPs. The model captures the per-user heavy-tailed traffic distribution, the highly skewed popularity distribution among CPs, and realistic functions for the transit, peering and operating costs incurred by the AP. We first examine a "baseline strategy" that follows current practice, in which the AP charges the same flat rate to all users. Further, the AP does not establish peering sessions with CPs. We then compare this baseline strategy to some strategies that an AP could use to increase its profitability. We focus mainly on strategies that are "network neutral", meaning that the AP does not differentiate between sources of content. These strategies are: usage-based pricing for heavy-hitters, limiting the traffic of heavy-hitters, selectively peering with some CPs, and caching content from selected providers. We also investigate a "non-network neutral" strategy in which an AP charges CPs directly. Our results show that certain strategies are rarely profitable or they are sensitive to factors that are not controlled by the AP (e.g., how would users react to heavy-hitter usage-based pricing?). On the other hand, the strategy of selective peering with CPs is non-disruptive and it can lead to a profit increase, relative to the baseline strategy, for the AP. To increase the effectiveness of such peering, it is important that the AP is co-located with the most popular CPs so that it can reduce peering costs. Caching can also help, even though the profit increase with that strategy depends significantly on the fraction of traffic that can be cached.

NetEcon'08, August 22, 2008, Seattle, Washington, USA.

2. THE NETWORK MODEL

We consider the interactions between three distinct species in the Internet ecosystem.

Access provider (AP): We focus on a single AP that sells Internet access to N paying users.

Content providers (CP): Content providers are the sources of content on the Internet. They do not provide access or transit service to any customers. Instead, CPs earn revenue from sources such as advertisements (out of band revenue). In this work, we do not model the costs and revenues of the CPs, instead focusing on the AP. Further, we take into account only the traffic flow from CPs to the AP, ignoring the requests from customers which are assumed to be small. Also, we assume that the AP does not receive any traffic from other APs, e.g. due to p2p applications. We intend to account for p2p traffic in the extended version of this paper.

Transit providers(TP): TPs provide transit for their customers, which are other ASes. TPs earn revenue by charging their customers for the volume of traffic sent and received. For simplicity, we consider a single TP that can provide transit to any other AS.

All ASes have a certain geographical scope, which is determined by the locations of their points-of-presence (PoPs). Large TPs are typically present globally, while APs and CPs could have only regional presence. As a result, an AP cannot always connect directly to a CP, and the TP is needed to provide reachability. An alternative is for the AP to establish a point of presence in remote locations (using a leased line to that location, for instance) or for the CPs to come closer to the AP by using a content distribution network.

We model two distinct types of inter-AS connections. In a transit (CP) relation, the customer "buys" transit service from the provider. The customer typically pays the provider for traffic sent in both directions on the customer-provider link. In a peering connection, two ASes agree to exchange traffic for free. If the CP and AP have a peering relationship, then the traffic flows directly between the two networks. If both the CP and the AP are customers of the tier-1 provider, then the tier-1 transits traffic that flows from the CP to the AP. Figure 1 illustrates this network model.

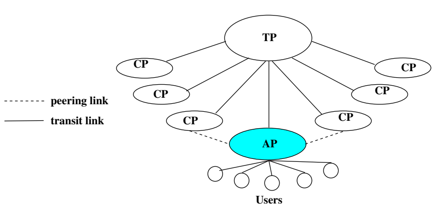

Figure 1: The network model

3. THE BASELINE MODEL

This section describes and evaluates a "baseline" model. We believe this model represents the most common current practices, and it captures pricing, connectivity and traffic distribution among the users of the AP and among CPs.

Connectivity and pricing: We consider a situation where both the AP and CP connect to the TP as transit customers, and there are no peering links between the AP and CPs. The AP has N users and charges each of those users based on a flat monthly rate R, giving it a revenue $N * R$. Based on common prices for Internet access in North America, we set the flat rate charged by the AP to $20/month. The TP charges both the AP and the CPs using the volumes of traffic sent in both directions. The transit pricing function we use for the TP is a concave increasing function of the form $c_t = m_t * V^{0.75}$, where m_t is the transit pricing multiplier used by the TP, V is the charging traffic volume (in Mbps), and c_t gives the monthly price for transit. This pricing function was used in [3] based on pricing data obtained from ISPs, and m_t was around 100 for transit ISPs in North America. Here, we use a transit multiplier $m_t = 100$ for the TP. Using this pricing function, a charging volume of 10Mbps costs $560 ($56/Mbps), while 10Gbps costs $100,000 ($10/Mbps). This illustrates the well known "economies of scale" in transit prices, i.e., the per-Mbps price decreases as the total charged capacity increases.

The TP typically calculates the charging volume V by dividing the month into 5 minute intervals, and V is the 95th percentile of the load on the customer link over all such intervals. Norton [13] notes that the 95th percentile charging model is based on the rule of thumb that the ratio of the 95th percentile to the average load is around 2:1 for web traffic. With the increase of video traffic, however, that ratio could be as high as 4:1. In this paper, we assume that the ratio of 95th percentile to average load is 3:1.

Local costs: The local cost of an AP consists of expenses to lease bandwidth for its network, purchase routers and other equipment, and to hire personnel to operate the network. This local cost is modeled as traffic independent and traffic dependent components of the form $c_l = f_l + m_l * V^{0.5}$. f_l is the traffic independent fixed cost component, and we set f_l =$250000/month for the AP. The local cost multiplier m_l is set to 500. In the absence of data about ISP operational costs, the local cost parameters are chosen to yield a net profit margin of approximately 20% for the AP, which is similar to what was seen in the balance sheet of a large North American access ISP. The local cost exponent is 0.5, which means that the cost incurred to carry traffic scales slower than the transit costs paid by the AP, while also showing economies of scale.

AP profit: The profit of the AP in the baseline model is the total revenues minus the transit and local costs, i.e.,

$$\mathcal{P} = NR - m_t * V^{0.75} - f_l - m_l * V^{0.5} \tag{1}$$

AP users: Each of the N users of the AP downloads a certain amount of traffic every month. To model the user traffic demand, we refer to a study of residential broadband access networks in Japan [4]. That study found that the distribution of the amount downloaded by a user is heavy tailed. In their measurements, approximately 4% of users download more than 75 GB/month (heavy hitters), while the remaining download less than 75 GB/month (normal users). We estimate the average of the normal and heavy hitter users as 300 MB/month and 10GB/month respectively, which gives an overall average of approximately 8GB/month.

Here, we draw the amount downloaded by each user from a truncated Pareto distribution with shape parameter 1.1 and mean 8 GB/month, in which case 2% of the users download more than 75 GB/month. The distribution is truncated from above at a point corresponding to the the access link speed. For example, a user behind a 1.5Mbps connection cannot download more than 486 GB/month. We consider different values of the cutoff point corresponding to the various common access speeds: 300kbps (97 GB/month), 700kbps (226 GB/month), 1.5Mbps (486 GB/month) and 10Mbps (3240 GB/month). Figure 2 shows the complementary CDF (CCDF) of the amount downloaded by each user. Unless noted otherwise, we use a cutoff point corresponding to the 1.5Mbps access speed in the rest of this paper.

Inter-AS traffic matrix: After generating the traffic demands for AP users, we create the distribution of the traffic among CPs as follows. The total incoming traffic for the AP is calculated as the sum of the incoming traffic demands for each of its users. This total is used to obtain the *average* incoming traffic, in Mbps, for

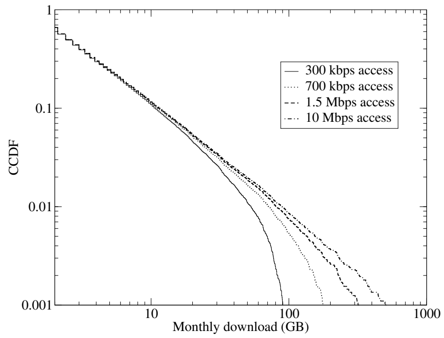

Figure 2: CCDF of the amount downloaded by users (GB/month).

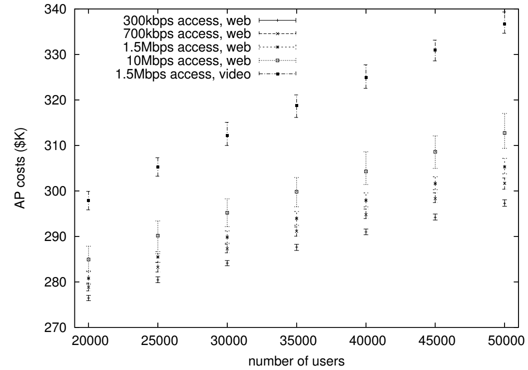

Figure 3: Variability of AP costs with the number of users, access speeds and type of traffic.

the AP. Chang et al. [2] found that the traffic distribution from the top content providers follows a Zipf-like distribution, with shape parameter ranging from 0.9 to 1.1. We assign the actual traffic volumes from each CP i to the AP using a Zipf distribution with shape parameter 1. This produces an effect where certain CPs are "popular" sources of content for the AP.

3.1 Evaluation of the baseline scheme

Here, we evaluate the performance of the baseline model used by the AP. We examine how the profit of the AP varies with the number of users, the different random samples of users, and the increasing amount of video traffic.

Variability in the set of users: We first examine what happens when the number of AP users increases. Recall that the traffic demand of each user is drawn from the heavy-tailed truncated Pareto distribution described earlier. The large variability in the traffic demand of individual users leads to also large variability in the costs incurred by the AP. To demonstrate this effect, we draw 1000 samples (corresponding to different samples of the user population) from a truncated Pareto distribution with shape parameter 1.1 and different cut-off points corresponding to different access speeds. We then create the inter-AS traffic matrices and calculate the costs incurred by the AP. Figure 3 shows the median and the min-max range of the AP costs across 1000 simulation runs. We find that the costs of the AP can vary significantly depending on the amount downloaded by its set of users. Moreover, as the user access speed increases, both the sample mean and the variance increase. This means that the increasing access speeds that users enjoy in the last few years will lead to increasing variability in the AP costs, making it harder for access providers to guarantee their profitability.

The impact of video traffic: A recent trend is that a large fraction of the traffic from content providers is streaming video. Norton [13] notes that video traffic is fundamentally different from web traffic, as the ratio of the 95th percentile to the average load due to video is 4:1, while for web traffic it is roughly 2:1. Consider, for example, that the users download web content at an average rate of V Mbps. Using the 2:1 ratio for web traffic, the AP provisions its network and gets charged by the TP for a charging volume $2V$. If the traffic is video, the AP must provision its network and purchase transit capacity for $4V$. This leads to a significant increase in the costs incurred by the AP, as shown in Figure 3.

4. ISP STRATEGIES

There are various strategies that an AP could deploy to increase its profits. In this section, we evaluate some strategies that are anecdotally mentioned in discussions about network neutrality and ISP economics. We attempt to gain a deeper, quantitative understanding of the pros and cons of these strategies. Further, we compare each strategy with the baseline, and evaluate the conditions under which the AP is able to achieve better profits than the baseline.

4.1 AP charges heavy hitters

In this charging strategy, the AP sets a threshold $\mathcal{T}$ to identify the users that download the largest amounts of traffic. These users are called the "heavy hitters", and the AP uses a volume-based pricing scheme for these users, rather than the flat rate. The price charged to a heavy hitter that downloads an amount of traffic $D > \mathcal{T}$ is given by: $c_v(D) = \frac{D*R}{\mathcal{T}}$, i.e., a heavy hitter is charged proportional to the amount of traffic downloaded.

A volume-based charging strategy is likely to be unpopular with the AP's users. In the presence of sufficient competition in the AP market, customers would switch from an AP that uses volume-based charging to an AP that offers flat-rate, "all you can eat" service. We model the unpopularity of volume-based charging with a probability that a user leaves this AP, referred to as *departure probability*. The departure probability depends on the threshold $\mathcal{T}$ set by the AP and is calculated as follows. For a value of $\mathcal{T}$ set by the AP, it is possible to calculate the number of users N_h that would be classified as heavy hitters. The number of users N_d that are expected to depart at this threshold is assumed to be proportional to N_h, $N_d = d * N_h$, where d is a positive parameter. The departure probability is then set to N_d/N (as long as $N_d \leq N$). The departure probability is applied to all users, not only those that are classified as heavy hitters. This captures the pragmatic fact that users are uncertain about their monthly usage and so they may leave the AP to avoid the possibility of extra fees if they get classified as heavy hitters. The parameter d determines the shape of the departure probability curve, as shown in Figure 4. The parameter d is also related to the degree of competition in the Internet access market. Without competition, users would be bound to a particular AP and d would be quite low as long as users need to have Internet access.

We evaluate the heavy hitter charging strategy by calculating the profit of the AP for different values of the threshold $\mathcal{T}$ and the parameter d. Figure 5 shows the profit of the AP as a function of the threshold $\mathcal{T}$. In the case of $d = 0.1$ and $d = 1$, the user departure probability decreases quickly with $\mathcal{T}$. In this case, even for low values of $\mathcal{T}$, the AP retains a significant fraction of users, and also charges them according to the downloaded traffic. Consequently, it

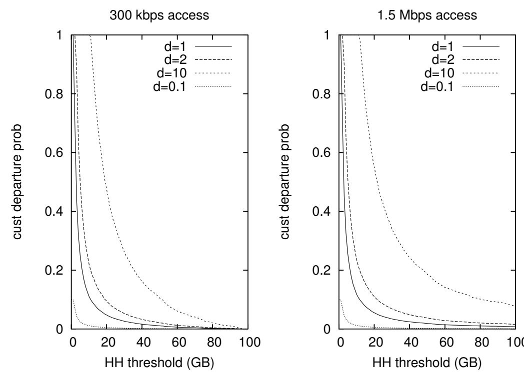

Figure 4: User departure probability as a function of $\mathcal{T}$, N=20000.

can achieve higher profits than the baseline scheme. In the case of $d = 2$ and $d = 10$, the user departure probability decreases slowly, and the optimal value of $\mathcal{T}$ shifts higher. In the most extreme case of $d = 10$, the optimal value of $\mathcal{T}$ occurs when the AP is able to keep all its users. The AP's profit in that case is similar to that of the baseline scheme. The curves for different access speeds are qualitatively similar. As expected, the benefit of heavy hitter charging is smaller if the users are limited by a smaller access speed. This is simply because there are fewer heavy hitters at any given value of $\mathcal{T}$ that the AP would be able to charge.

The previous results illustrate that a volume-based charging strategy is quite sensitive to the user departure probability, which is not controlled by the AP. Even if the departure probability is low, it is difficult to determine the optimal value of $\mathcal{T}$, and hence this strategy is not robust to the selection of this threshold. If the AP sets $\mathcal{T}$ to a sub-optimal point, it could end up with even lower profit than in the baseline scheme.

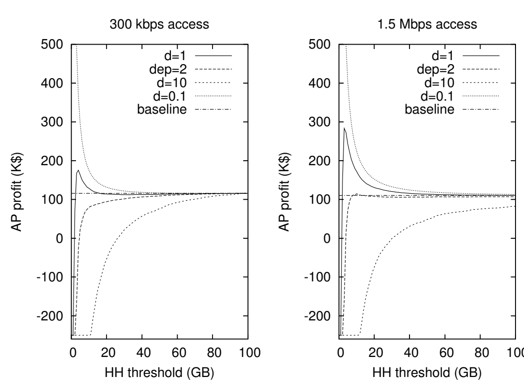

Figure 5: AP profit as a function of $\mathcal{T}$ when the AP charges heavy hitters, N=20000.

4.2 AP caps heavy hitters

In this strategy, the AP imposes "download caps" on its users, i.e., users are not permitted to download more than $\mathcal{T}$ GB/month. If a user reaches that threshold, her account is blocked for the remainder of the month[1]. In this strategy, the AP charges each access customer with the same flat rate R. As with the strategy of charging heavy hitters, capping the amount that a user is allowed to download can be an unpopular strategy. We assume that the departure probability with this strategy is modeled using the same function as in 4. In practice, the departure probability in this model may be higher or lower than in the heavy hitter charging scheme, depending on the user population, the available pricing plans and policies of competing APs, and how APs justify/present these policies to their users.

Figure 6 shows the profit of the AP as a function of the threshold $\mathcal{T}$ used by the AP to cap customers. We find similar trends as in the case of heavy hitter charging. The strategy of capping heavy hitters performs worse than heavy hitter charging, even when the customer departure probability drops quickly (curves marked "d=1" and "d=0.1"). By capping heavy users, the AP is only able to save on its operating costs, and does not gain any additional revenue. With the same user departure profile as in the case of heavy-hitter charging, this strategy would be less profitable for the AP.

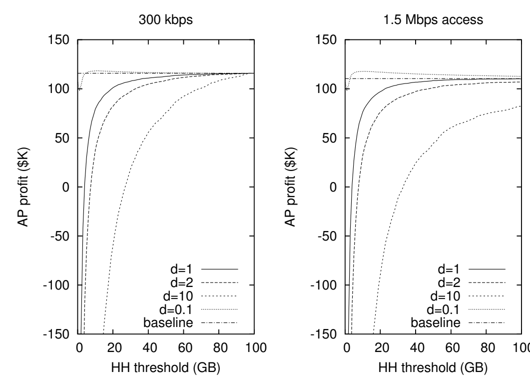

Figure 6: AP profit as a function of $\mathcal{T}$ when the AP caps heavy hitters, N=20000.

4.3 AP charges CPs

There has been much debate on whether an AP should be able to discriminate between CPs. To recover the costs due to increasing traffic volumes, APs would like to charge the CPs that produce most traffic. The AP could rank CPs in decreasing order of traffic volume, and charge a certain fraction of the top providers. We assume that the AP would use its transit pricing function to charge those CPs. This strategy is again likely to be unpopular, and a fraction of the AP's customers may choose to switch to another AP. We model this by making the customer departure probability dependent on the fraction of CPs charged by the AP using a function of the form $y = ax^b + c$. The parameter b determines the shape for the departure probability curve as shown in Figure 7. The values of a and c are adjusted to give a departure probability of 0 when no CPs are charged and 1 when all CPs are charged. We find that the profit of the AP depends stongly on the customer departure probability.

The trends in all the three previous strategies highlight an important tradeoff involved with strategies that can compromise the customer base of the AP. If the AP charges or throttles heavy hitters, or tries to charge CPs instead, it may lose some of its customers. Whether such a charging strategy increases the profitability of the ISP depends heavily on the customer departure probability. As such, the fate of an AP that deploys such a strategy would be highly dependent on user behavior. In the following, we investigate

[1]In practice, the AP may choose to seriously rate-limit a user that exceeds her threshold. For simplicity, we consider the more extreme measure where the user is blocked.

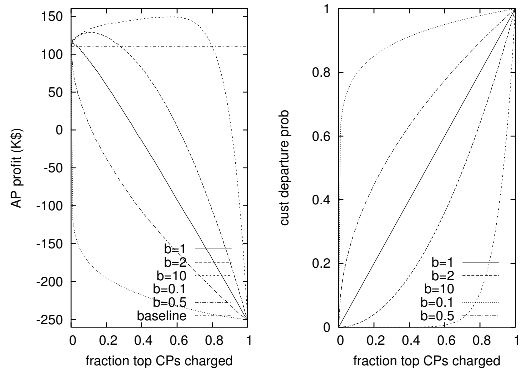

Figure 7: AP profits by charging CPs, as a function of the fraction of CPs charged, N=20000, 1.5Mbps access.

alternate, non-disruptive strategies that the AP could use to increase its profits.

4.4 Selective peering with CPs

So far, we have considered the baseline model in which the AP and CPs are customers of the TP, and there is no direct peering between the AP and CPs. Here, we study a strategy where the AP follows a selective peering policy, peering with a CP depending on the potential benefits and costs associated with peering.

Chang et al. [3] studied the fixed and traffic dependent costs associated with peering. To model the traffic dependent peering costs, we use the function $c_p = f_p + m_p * V^{0.25}$, which is the function used in [3]. The parameters f_p and m_p are different for each CP, and they indicate the difficulty of peering with that CP. For example, some CPs may be colocated in the same city as the AP, in which case the peering costs are low. On the other hand, some CPs may be in entirely different continents, in which case it costs much more (or it may even be impossible) to peer with that CP. We assume that content providers fall into different classes depending on the ease of peering with that content provider. The peering cost multiplier is different for each class of CPs and the values are 10 (easiest), 100 (medium) and 1000 (hardest) peering. The fixed peering costs for these classes are \$500/month (easy), \$5000/month (medium) and \$50000/month (hard). These classes of peering costs are meant to capture the fact that it may be practically impossible for the AP to peer with certain CPs (the "hard" class). For CPs in the "medium" and "easy" classes, it makes sense for the AP to peer, if the traffic volume is sufficiently large. The figures we use for the fixed costs of the easy and medium classes are in the same range as those quoted by Norton [13]. The fixed cost of the "hard" class is very large, to model the fact that it does not make sense for the AP to peer with a CP in that class.

We investigate two distinct divisions of the CPs into the three peering cost classes. In the first, a CP is equally likely to be in any of the three classes. In the second, a CP is in the "easy" and "medium" peering class with probability 0.1 each, and with probability 0.8 is in the "hard" class. We also vary the assignment of CPs to these classes. In one case, the set of CPs in each class is determined randomly. In the second case, the most popular CPs are also the easiest to peer with. This scenario is likely in the case that the popular CPs expand their networks and are thus present in multiple peering points. A recent study gives evidence that some content providers are indeed expanding their networks in recent times [10].

To determine the set of CPs with which to peer, the AP uses the following procedure. The AP considers separately each CP i, and decides whether to peer with CP i based on a simple rule-of-thumb. Let $V(i)$ be the traffic from CP i. The AP calculates the estimated benefit of peering (saving in transit costs) as the amount that would be paid to the TP, assuming a charging volume $V(i)$. This is an approximation, as it does not account for the economies of scale when multiple CPs send traffic to the AP through the same TP.

The AP decides to peer with the CP if the following condition is satisfied:

$$\frac{m_t * V(i)^{0.75}}{f_{p_i} + m_{p_i} * V(i)^{0.25}} > \mathcal{R}$$

The estimated cost of sending the traffic $V(i)$ through the TP is given by $m_t * V(i)^{0.75}$. The cost of peering with CP i is given by $f_{p_i} + m_{p_i} * V(i)^{0.25}$.

Figure 8 shows the profit of the AP as the ratio $\mathcal{R}$ is changed. The left plot is for the case where the CPs are distributed randomly in the three cost classes ("rand"). The number of CPs in each class is either the same (marked "eq"), or is skewed towards "hard" peering (marked "sk"). We repeat the simulations for different number of content providers ("N 50" and "N 200"). All curves show qualitatively similar behavior. If the ratio $\mathcal{R}$ is set too low, the AP forms peering relationships that incur more cost than the transit savings. On the other hand, if the ratio is too large, the AP does not peer with certain content providers that would have reduced the transit costs for the AP. We see that the optimal point for the ratio $\mathcal{R}$ occurs *after* $\mathcal{R} = 1$. This is because of the fact that the AP uses an estimate of the transit savings. Due to the economies of scale in the TP's transit pricing function, the AP *over-estimates* the potential savings in transit. An interesting trend is that above a certain value of $\mathcal{R}$, the profit is fairly robust to changes in $\mathcal{R}$. Also, the profit from peering is larger when the incoming traffic is split into a smaller number of CPs (N=50 vs N=200).

In Figure 8, the profit increase from peering is only 5% over the baseline scheme. Note that the absolute value of the profit (and the improvement over the baseline) depends on certain parameter values, such as the number of AP customers and the fixed local costs. We stress that with an appropriate choice of $\mathcal{R}$, the AP's profit with peering is guaranteed to be equal to or greater than that with the baseline scheme. The peering strategy does not increase the revenues of the AP or affect the fixed local costs. Instead, it reduces the traffic-dependent costs incurred by the AP. For the case of ("N 50 sk sort") in Figure 8, the traffic-dependent cost is \$33,000 with peering and \$40,000 for the baseline scheme, i.e., peering reduces those costs by 17%.

The strategy of selective peering appears to be quite attractive because the parameter $\mathcal{R}$ is controlled solely by the AP. The right graph shows that the benefit from selective peering is larger for the peering cost structure where the CPs with the largest traffic volume fall into the "easy" or "medium" classes. This could happen if the largest CPs expand their networks, and are thus easy to peer with. Given that such expansion by CPs is already happening [10], selective peering could be a profitable strategy for many APs.

4.5 AP caches CP content

We also consider the case where the AP chooses to cache the content that it receives from the major content providers. By caching content, most requests for content by the access customers of the AP are handled locally, and hence can save transit costs for the AP. Here, we assume that there exists a certain fraction h of content from each CP that is "cacheable". When the AP caches content from CP i, the traffic $h * V(i)$ is served locally, while $(1 - h) * V(i)$ has to be downloaded through the transit provider. The fraction h captures the fact that all content from the CP may not

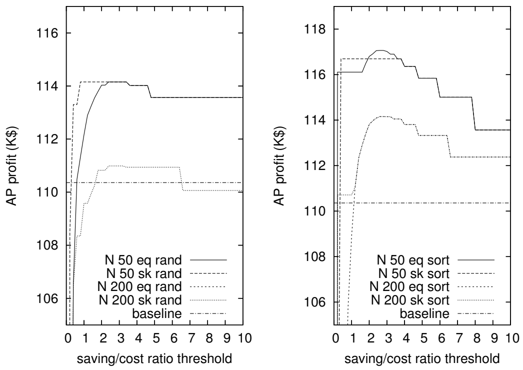

Figure 8: AP profits with selective peering as a function of $\mathcal{R}$. N=20000, 1.5Mbps access.

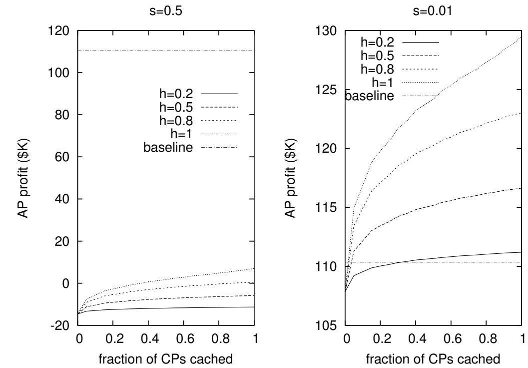

Figure 9: AP profits from caching CP content. n_A=20000, 1.5Mbps access

be cacheable, e.g. dynamically generated content or live video streams. By caching content locally the AP saves transit costs. We model the costs of caching CP content, which involve purchasing servers and bandwidth to serve the content locally. We assume that caching adds to the fixed local cost of the AP according to the relation $f_c = s * f_l$, where s is a parameter that determines how the caching cost relates to the local cost. The AP must decide how many of the largest CPs to cache. The CPs are considered in decreasing order of traffic volume, because caching the largest CPs can lead to the largest potential savings in the transit costs.

Figure 9 shows the profitability of the AP as a function of the fraction of CPs cached. We simulate two cases corresponding to $s = 0.01$ and $s = 0.5$. First, we observe that the profit of the AP increases with the fraction of CPs that it caches, following a concave function. The figure on the left shows the case where the caching cost is large (equal to half of the fixed local cost). In this case, the AP is not able to do better than the baseline scheme, even if it caches all CPs. The right graph shows the case where caching costs are very low in comparison with the fixed local costs. In this case, the AP is able to achieve higher profits than the baseline scheme, depending on how much traffic is cacheable.

This analysis indicates that the attractiveness of content caching depends on the additional local cost incurred by the AP. The AP may be able to optimize its network in such a way that caching costs are small in relation to fixed local costs. In that case, the amount of CP traffic that is cacheable determines whether the AP can obtain higher profits than the baseline scheme. Note that this is again a parameter that is out of the AP's control. The previous scenario represents the case where a CP allows the AP to freely cache its content. It is possible that the CP does not allow the AP to do so due to copyright or privacy concerns. As such, our analysis of this strategy evaluates the *best case* scenario for the AP.

5. CONCLUSIONS

We took a quantitative approach towards understanding the network neutrality issue from the point of view of an access provider. We examined a baseline scheme that follows current practices, and some variants of charging and connection schemes. Our results show that AP strategies based on charging are rarely profitable or are highly sensitive to factors out of the control of the AP. On the other hand, the AP can obtain substantial additional profit by engaging in selective peering with CPs or caching CP content locally.

6. REFERENCES

[1] J. M. Bauer. Dynamic Effects of Network Neutrality. *International Journal of Communication Vol.1, 2007*.

[2] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger. An empirical approach to modeling inter-as traffic matrices. In *Proceedings of the Internet Measurement Conference (IMC)*, 2005.

[3] H. Chang, S. Jamin, and W. Willinger. To Peer or Not to Peer: Modeling the Evolution of the Internet's AS-Level Topology. In *Proceedings of IEEE Infocom*, 2006.

[4] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proceedings of ACM SIGCOMM*, 2006.

[5] J. Crowcroft. Net Neutrality: The Technical Side of the Debate: A White Paper. *ACM SIGCOMM CCR*, 2007.

[6] N. Economides. The Economics of the Internet Backbone. *Handbook of Telecommunications Economics, Elsevier Publishers*, 2006.

[7] N. Economides and J. Tag. Net Neutrality on the Internet: A Two-Sided Market Analysis. *SSRN eLibrary*, 2007.

[8] E. W. Felten. Nuts and Bolts of Network Neutrality. Technical report, 2006.

[9] R. Frieden. Network Neutrality or Bias? Handicapping the Odds for a Tiered and Branded Internet. *Bepress Legal Series. Working Paper 1755.*, 2006.

[10] P. Gill, M. Arlit, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of Passive and Active Measurement Conference (PAM)*, 2008.

[11] R. W. Hahn and R. E. Litan. The Myth of Network Neutrality and What We Should Do About It. Technical report, 2006.

[12] W. B. Norton. The Evolution of the U.S. Internet Peering Ecosystem. *Equinix white papers*, 2004.

[13] W. B. Norton. Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem. *Equinix white papers*, 2007.

[14] Robert Beverly and Steven Bauer and Arthur Berger. The Internet Is Not a Big Truck: Toward Quantifying Network Neutrality. In *Proceedings of Passive and Active Measurement Conference (PAM)*, 2007.

[15] J. G. Sidak. A Consumer-Welfare Approach To Network Neutrality Regulation of the Internet. *Journal of Competition Law and Economics, Vol. 2, pp. 349-474, 2006*.

[16] B. Van Schewick. Towards an Economic Framework for Network Neutrality Regulation. *Journal on Telecommunications and High Technology Law, Vol. 5, 2007*.

[17] T. Wu. Network Neutrality, Broadband Discrimination. *Journal of Telecommunications and High Technology Law, Vol. 2, p. 141, 2003*.

[18] X. Yang, G. Tsudik, and X. Liu. A Technical Approach to Net Neutrality. In *Proceedings of Fifth Workshop on Hot Topics in Networks (Hotnets-V)*, 2006.