

Balancing Traffic Flow Efficiency with IXP Revenue in Internet Peering

Md Ibrahim Ibne Alam, Koushik Kar
Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute
Troy, New York, USA
{alam2, kark}@rpi.edu

Elliot Anshelevich
Computer Science
Rensselaer Polytechnic Institute
Troy, New York, USA
eanshel@cs.rpi.edu

Abstract—We consider a traffic peering game between Internet Service Providers (ISPs) at an Internet Exchange Point (IXP), where each ISP pair has a choice of exchanging traffic through a public IXP switch or sending the traffic through transit providers. We analyze the traffic flow efficiency (measured as social welfare) and the IXP revenue at the equilibrium of this game, as a function of the per-unit price charged by the IXP. We show that there exists a price point at which both social welfare and revenue are high, and the corresponding price-of-anarchy values can be expressed in terms of certain sublinearity measures of the inverse demand curves of the ISPs. Simulations carried out using models based on actual IXP data obtained from PeeringDB demonstrate that the theoretical bounds correctly capture the performance trends against the variation of price, and for a carefully chosen pricing point both social welfare and IXP revenue are within a factor of two of the corresponding optimal values.

Index Terms—Internet Exchange Point, Internet Service Provider, Public Peering, Social Welfare, Revenue

I. INTRODUCTION

Internet eXchange Points (IXPs) carry a large fraction of the traffic between Internet Service Providers (ISPs) in the Internet today. It has been estimated that almost 80% of the IP addresses in the world can be reached via public peering, and 20% of all the traffic traces go through IXPs [1]. Despite a falling trend in transit costs, peering between ISPs (particularly between content and access ISPs) is on the rise [2], [3].

In basic terms, an IXP is a data center with network switches through which ISPs exchange traffic (peer) with each other [4], [5]. A large number of these IXPs, both in Europe and particularly US, operate for profit (e.g., Equinix [6]), charging fees to each ISP for sending/receiving traffic through the IXP. The price charged by the IXP typically depend on the capacity of the port leased by the ISP. The peering decisions are determined bilaterally by the ISPs themselves, taking into account the potential Quality-of-Service (QoS) improvements due to peering at the IXP (as opposed to sending the traffic through transit providers) and the price charged by the IXP.

Arguably, the price charged by the IXP can have a significant impact on the peering relationships formed at the IXP, or equivalently, the traffic flows exchanged through the IXP. As we demonstrate later in this paper, there is an inherent tradeoff between the efficiency of traffic flows exchanged through the IXP and the revenue earned by the IXP, and the pricing policy

at the IXP must be chosen carefully if a balance needs to be attained between these two objectives.

This paper models and analyzes this tradeoff between traffic flow efficiency and IXP revenue at the traffic exchange equilibria between ISPs at an IXP. In the traffic exchange (peering) game that we model, ISPs determine how much traffic to exchange through an IXP, in response to the per-unit traffic pricing policy set by the IXP. This decision in turn determines how much port capacity an ISP must lease at the IXP, and which other ISPs a given ISP will peer with. We analyze the equilibrium of this game as a function of the per-unit price, and show that there is a pricing point that can attain good traffic flow efficiency (measured as social welfare) and IXP revenue at the same time. More specifically, we show that when the per-unit price is chosen appropriately, the Price-of-Anarchy (PoA) for both social welfare and revenue is small, and are characterized by certain worst-case sublinearity measures of the inverse demand curve for the ISP pairs. Simulations carried out using models derived from data on the 28 largest US IXPs obtained from PeeringDB confirm these observations, and show that our theoretical bounds capture the actual performance trends reasonably well.

II. RELATED WORK

Our game-theoretic model is inspired by a prior line of work on network formation games (e.g., [7]–[9]), where two nodes can build links through mutual agreement but can sever links individually. However, unlike prior models, in this work the utility derived from a connection (peering) is not fixed, but depends both on congestion and the price charged by the market maker (IXP). Further, unlike prior work (such as [10]) that considers the transit vs peering question focusing only on ISP costs, we analyze the system with the goal of providing a good balance between social welfare and IXP revenue.

Our work is closely related to the models analyzed in our prior work [11]–[13], but differs from these existing works in several important aspects. While [11] analyzes the effect of proportional pricing on social cost and evaluates the revenue performance through simulation, we propose and theoretically analyze pricing policies that simultaneously attain good social welfare and revenue. The work in [12] considers how the operational cost of a non-profit IXP should be shared among the

member ISPs, whereas we consider a profit-making IXP in this paper. Therefore, while we analyze IXP revenue, and how it can be traded off with social welfare, [12] only considers social cost. The model considered in [13] is very similar to ours, which also shows the existence of pricing points that can attain good social welfare and revenue simultaneously. However, these results rely on certain strong regularity (smoothness) condition of the inverse demand curves, which do not hold for IXPs, as we will see in Section V. Consideration of realistic inverse demand curves for IXPs requires us to consider the sublinearity measures outlined in Section IV, which applies to a very large set of (including non-smooth) inverse demand curves. Further, while [13] only analyzes the problem theoretically and for generic large markets, we consider it in the context of IXPs, and evaluate our results using inverse demand curves estimated from actual IXP data.

III. SYSTEM MODEL

A. Game Theoretic Model

We consider a set N of ISPs (agents in our game-theoretic model) that are involved in traffic exchange through a public switch offered by an IXP. An ISP pair (i, j) has a total traffic demand of B_{ij} between themselves. Part of this traffic, y_{ij} , is routed through the public switch, while the rest is either sent by other means (using the ISPs' transit providers, or through private peering) or not sent at all¹. The strategy of ISPs i and j , each acting in self-interest, involves deciding on the y_{ij} . Clearly, the decisions of ISPs i and j are coupled, as y_{ij} must be jointly decided by the two ISPs.

We denote λ_{ij} to be the per-unit utility each of the two ISPs i and j derive from sending traffic through the IXP.² The traffic that is exchanged through the public switch incurs a congestion cost of $d(y)$ per-unit traffic, which depends on the total traffic y sent through the switch. This congestion cost will typically be reflected in terms of average delay experienced by the traffic (and therefore we will sometimes use the terms 'congestion cost' and 'delay' interchangeably); however, $d(y)$ could also represent other Quality-of-Service (QoS) parameters (or a combination of them) that are affected by the overall load at the public switch. We assume that $d(y)$ is a given function (i.e., not part of the strategy); however, we will explore the efficiency of the equilibrium for different forms of the function $d(y)$. Additionally, each ISP has to pay a price of $p(y)$ to the IXP per-unit traffic, for the use of the public switch. Table I summarizes some of the most commonly used terms and notations in our model.

Before concluding the discussion of our basic model, two points are worth mentioning. Firstly, note that the port capacity needed by ISP i at the IXP would be directly related to $y_i = \sum_j y_{ij}$. Therefore traffic exchange strategy decisions $(y_{ij}, \forall j)$ in our game-theoretic model has a direct bearing on the port capacity provisioning decision faced by the ISP i . Secondly,

¹We assume demand to be undirected, so $y_{ij} = y_{ji}$.

²This utility can be from cost savings derived for not having to pay for transit services or private peering, or from improved traffic delay (QoS) when peering through the IXP.

TABLE I
SUMMARY OF COMMONLY USED NOTATION.

Term	Description
y_{ij}	Traffic of ISP pair (i, j) sent publicly through the IXP.
y_i	$\sum_j y_{ij}$, total traffic of ISP i going through the IXP.
y	$\frac{1}{2} \sum_i \sum_j y_{ij}$, total traffic flowing through the IXP.
$\vec{y}$	Total traffic allocation vector (vector of values y_{ij})
λ_{ij}	Per-unit utility received by (i, j) by sending traffic through IXP
$d(y)$	Congestion cost per-unit traffic incurred at the IXP
$p(y)$	Price per-unit traffic set by the IXP

note that while we allow for y_{ij} to be a fraction of B_{ij} for ease of analysis, the properties of the equilibrium solution imply that in almost all cases, an ISP pair (i, j) will send all of their traffic through the IXP (if they decide to peer with each other), or nothing (not peer at all).

B. Social Welfare and Revenue

Given the above model setup, we define the Social Welfare (SW) and IXP Revenue (Rev) next. Overall, SW can be split into the welfare the ISPs are getting from the IXP, and the welfare the IXPs are receiving (as a form of Revenue). The SW of ISP i , denoted by $SW_i(\vec{y}, p(y), d(y))$, is calculated as

$$\sum_{(ij) \ni i} y_{ij} \lambda_{ij} - \{p(y) \sum_{(ij) \ni i} y_{ij} + d(y) \sum_{(ij) \ni i} y_{ij}\}, \quad (1)$$

where the first term is the utility of the ISP by sending $\sum_{(ij) \ni i} y_{ij}$ peering traffic via the IXP, and the rest denotes the sum of the payments paid to the IXP and (implicit) loss of the ISP's welfare caused by the congestion at the switch. Denoting $c(y) = p(y) + d(y)$ and $W_i(\vec{y}) = \sum_{(ij) \ni i} \lambda_{ij} y_{ij}$, the welfare of ISP i is expressed as

$$SW_i(\vec{y}, c(y)) = W_i(\vec{y}) - c(y)y_i. \quad (2)$$

Note that $c(y)$ can be viewed as the aggregate cost seen by the ISPs per-unit traffic. If we denote $\sum_i W_i(\vec{y}) = 2W(\vec{y})$, then the total Welfare of all ISPs is given by

$$SW_{ISP}(\vec{y}, c(y)) = 2(W(\vec{y})) - c(y)y, \quad (3)$$

where the multiplier of 2 comes from the fact that y_i and y_j both include y_{ij} . The revenue gathered by the IXP is:

$$Rev(\vec{y}, p(y)) = p(y) \sum_i \sum_{(ij) \ni i} y_{ij} = 2p(y)y. \quad (4)$$

Thus, the Social Welfare (SW) for the system is given by

$$\begin{aligned} SW(\vec{y}) &= SW_{ISP}(\vec{y}, c(y)) + Rev(\vec{y}, p(y)), \\ &= 2W(\vec{y}) - 2d(y)y = 2W(\vec{y}) - 2E(y), \end{aligned} \quad (5)$$

where $E(y) = d(y)y$. The first term of this SW is the utility that ISPs are getting by using the IXP, and the second term is the cost of the congestion at the shared switch in the IXP. For the rest of the paper, unless otherwise stated, we will assume $E(y)$ to be a continuous, piece-wise differentiable function with $E(0) = 0$, and $E'(y)$ to be a non-decreasing function with $E'(0) = 0$. Note that SW does not consist of $p(y)$ which is the price of per-unit traffic charged by the IXP to the ISPs.




Fig. 1. Normalized Inverse Demand Curve ($\lambda(y)$).

IV. THEORETICAL ANALYSIS

We state our main theoretical claims in this section; proof outlines of all the results can be found in the Appendix.

A. Properties of the Model

1) *Equilibrium Traffic*: We assume *price-taking* ISPs, i.e., they see the per-unit traffic cost $c(y)$ for sending traffic through the IXP, and will only send traffic which is worth paying that cost. Thus, ISP i will send all the traffic through the IXP as long as $\lambda_{ij} \geq c(y)$. This leads to the following notion of equilibrium traffic flow:

Definition IV.1. A traffic flow $\bar{y}_e$ with $y_e = |\bar{y}_e|$ is said to be an equilibrium flow if and only if all the traffic with $\lambda_{ij} > c(y_e)$ is sent and the traffic with $\lambda_{ij} < c(y_e)$ is not sent.

Based on this definition, we simplify our terminology and use the term ‘equilibrium traffic flow’ to refer to the *total* flow through the IXP at equilibrium (a scalar), and denote it by y_e .

2) *Inverse Demand Curve*: Next, we state two important properties of equilibrium traffic flows that will be useful in our PoA analysis. We first define the notion of the *inverse demand curve*, $\lambda(y)$, constructed as follows. First, the λ_{ij} values are arranged in a decreasing order (ties broken arbitrarily); let λ^k be the k^{th} highest value, and B^k be the corresponding traffic demand. Then, the $\lambda(y)$ curve is a non-increasing step-function, with the step of height λ^k having a width of B^k . Let $\lambda(y^-)$ denote the limit of $\lambda(x)$ as x approaches y from below, and similarly $\lambda(y^+)$ if it approaches y from above. We then have the following property:

Theorem IV.1. y_e is an equilibrium traffic flow if and only if $\lambda(y_e^-) \geq c(y_e) \geq \lambda(y_e^+)$. Moreover, such a flow always exists.

3) *Social Optimum*: We next provide an important property of the optimal traffic flow (OPT), one that maximizes the total Social Welfare (SW), with which the equilibrium solution will be compared. Again, with slight abuse of terminology, by ‘optimal traffic flow’ we refer to the *total* traffic flow at the social optimum, denoted by y_p .

Theorem IV.2. At social optimality, all the traffic with $\lambda_{ij} > E'(y_p)$ flows through the IXP and all traffic with $\lambda_{ij} < E'(y_p)$ does not. Also, $\lambda(y_p^-) \geq E'(y_p) \geq \lambda(y_p^+)$.

B. Definitions

Guided by the inverse demand ($\lambda(y)$) curves derived from data from IXPs (described in Section V-B2), we characterize these curves in terms of two sublinearity properties which will be used in our PoA analysis. These properties are derived from sublinear curves characterized by a single parameter (α) that can provide tight lower support to the ($\lambda(y)$) curves. The two properties, termed *shift factor* (α_1) and *stretch factor* (α_2) of the curves, are formally defined below.

Definition IV.2. An inverse demand curve ($\lambda(y)$) has a shift factor α_1 if $\frac{\lambda(y)}{\lambda_{max}} + \frac{y}{y_{max}} \geq \alpha_1, \forall y$.

Definition IV.3. An inverse demand curve ($\lambda(y)$) has a stretch factor α_2 if $\left(\frac{\lambda(y)}{\lambda_{max}}\right)^{\alpha_2} + \left(\frac{y}{y_{max}}\right)^{\alpha_2} \geq 1, \forall y$.

Figure 1 illustrates the shift factor (α_1) property with the red dash-dot straight line, and the stretch factor (α_2) property with the red solid curved line. In the figure both the axis are normalized with their respective maximum values, i.e., $\bar{\lambda}(y) = \frac{\lambda(y)}{\lambda_{max}}$, and $\bar{y} = \frac{y}{y_{max}}$, hence the maximum possible value on both axis is 1. Note that for a given $\lambda(y)$ curve, for tightness we would like to choose the *maximum* value of α_1 (α_2) that satisfies the condition in Definition IV.2 (Definition IV.3).

Definition IV.4. The price of anarchy (PoA) of Social Welfare (SW) for a pricing policy π is the ratio of SW at optimum to SW at the equilibrium induced by π .

Definition IV.5. The price of anarchy (PoA) of Revenue (Rev) for a pricing policy π is the ratio of maximum achievable Rev to Rev at the equilibrium induced by π .

C. PoA Analysis

To choose a pricing policy for the IXP that achieve good values of $PoA(SW)$ and $PoA(Rev)$ simultaneously, we propose a per-unit pricing $p(y) = \beta_b = \max(\beta_e, \beta_p)$, where $\beta_p = \lambda(y_p) - d(y_p)$ with y_p being the socially optimal traffic flow (OPT). Also, β_e depends on the specific sublinearity measure used, among the two measures defined in Definitions IV.2 and IV.3. If $\lambda(y)$ has a shift factor α_1 , then $\beta_e = K\alpha_1 - d(y_e)$, where y_e is such that $\lambda(y_e) = K\alpha_1$. If $\lambda(y)$ has a stretch factor α_2 , then $\beta_e = K^{(1/\alpha_2)} - d(y_e)$ where y_e is such that $\lambda(y_e) = K^{(1/\alpha_2)}$. Also, the value of K ($0 < K < 1$) is an appropriately chosen constant. For visual aid β_e and β_p are depicted in Figure 1, and for convenience we denote $c_e(y) = \beta_e + d(y)$, and $c_p(y) = \beta_p + d(y)$.

Theorem IV.3. For a $\lambda(y)$ curve with a shift factor α_1 , charging a per-unit price $\beta_b = \max(\beta_e, \beta_p)$ attains at least a $\left(\frac{1}{\alpha_1(1-K)}, \max\left(\frac{1}{\alpha_1(1-K)}, \frac{2}{K\alpha_1}\right)\right)$ approximation of the maximum SW and Revenue respectively.

Corollary IV.3.1. For a $\lambda(y)$ curve with a shift factor α_1 , charging a per-unit price $\beta_b = \max(\beta_e, \beta_p)$ with $K = 2/3$, attains a PoA of $\left(\frac{3}{\alpha_1}\right)$ for both SW and Revenue.

Theorem IV.4. For a $\lambda(y)$ curve with a stretch factor α_2 , charging a base price $\beta_b = \max(\beta_e, \beta_p)$ attains at least

$a \left(\frac{1}{(1-K)^{1/\alpha_2}}, \max \left(\frac{1}{(1-K)^{1/\alpha_2}}, \frac{2}{K^{(1/\alpha_2)}} \right) \right)$ approximation of the maximum SW and Revenue respectively.

Corollary IV.4.1. For a $\lambda(y)$ curve with a stretch factor α_2 , charging a base price $\beta_b = \max(\beta_e, \beta_p)$ with $K = \frac{2\alpha_2}{1+2\alpha_2}$, gives a PoA of $(1 + 2\alpha_2)^{1/\alpha_2}$ for both SW and Revenue.

V. SIMULATION

A. Data Collection

To achieve realistic traffic demand values B_{ij} and the corresponding utility λ_{ij} , data from PeeringDB and CAIDA databases were collected and analyzed. PeeringDB was utilized for obtaining information about the locations of the IXPs, the ISPs peering in that location (also called Point-of-Presence (PoP)), and the port capacity each ISP has purchased. On the other hand, we utilized CAIDA to obtain the number of active routers and their approximate location (at a city level) for each ISP, to approximate the amount of traffic that may be generated for that ISP at that location.

B. Simulation Setup

1) *Generating $\lambda(y)$ curves:* To generate the $\lambda(y)$ curves we need two sets of values: i) the traffic demand between ISPs (B_{ij}), and ii) the per-unit utility λ_{ij} for that traffic. While the exact values for these are very difficult to estimate closely, we make several reasonable approximations based on the PoP locations (obtained from PeeringDB), router densities (obtained from CAIDA) and previously published models on traffic demand and pricing. The traffic demand between two ISPs serving at two different PoP locations is determined using the gravity model [14]. If ISP i has R_A number of routers serving at location A and ISP j has R_B number of routers serving at location B , then the traffic demand between these two ISPs for these two locations is thus approximated as $Y_{AB} = \frac{R_A \times R_B}{d_{AB}^2}$. Then the summation of all these values over all the possible pairs of router locations gives us the total traffic demand between these two ISPs, hence $B_{ij} = \sum_{A,B} Y_{AB}$. The utilities λ_{ij} are calculated based on savings in transit costs, following [15], which models transit costs as being linearly or logarithmically proportional to the distance that traffic has to travel. Since traffic between different locations of the same ISP pair (say Y_{AB}) is going to travel different distances (d_{AB}), we use the weighted average of these distances: for some ISP pair (i, j) , we set $d_{ij} = \frac{\sum_{A,B} Y_{AB} d_{AB}}{\sum_{A,B} Y_{AB}}$. Thus, we have the per-unit utility as, $\lambda_{ij} = a \times d_{ij}$ or $\lambda_{ij} = a \times \log(d_{ij})$, for an appropriately chosen constant a . The total traffic B_{ij} of ISP pair (i, j) is split across the different PoP locations that the two ISPs have in common, in a way that the traffic on any path is inversely proportional to total end-to-end geographical distance of the path. Figure 2 shows some sample $\lambda(y)$ curves from the largest (in terms of number of participating ISPs) 28 IXPs in USA, as generated by this approach.

2) *Simulations:* Simulations were done for the largest 28 IXPs among the 140 IXPs present in USA. Most of the remaining (smaller) IXPs have a very small number of participating ISPs, resulting in a few discrete $\lambda(y)$ values and

making the study of the equilibrium uninteresting. Also, from the PeeringDB port capacity data, it was found that more than 95% of the total port capacities (which can be seen as an indicator of the traffic flowing through these IXPs) are accounted for by considering the largest 28 IXPs. In our simulations, we consider two broad class of delay functions, namely polynomial delay and queuing delay functions. For polynomial delay function with exponent n , $d(y) = ay^n$; the PoA value was calculated considering the value of a that resulted in the worst PoA. Since the PoA value also critically depends on the $\lambda(y)$ curves which differ across IXPs, both worst and average case PoA values were calculated by taking the the worst value and average values over all the $\lambda(y)$ curves, respectively. For the case of queuing delay functions, $d(y) = \frac{a}{\mu - y}$, the results were generated for different utilization factors ($U_f = y/\mu$) by scaling the value of a , and taking the worst case values of PoA. Due to space limitations, we only plot the average PoA values, but include the worst-case PoA values and theoretical bounds in Table II.

C. Results and Discussion

From Theorem IV.3 and IV.4, we expect that both the $PoA(SW)$ and $PoA(Rev)$ values will depend heavily on the value of α (both α_1 and α_2); and the smaller the value of α is the higher the PoA will be. From the simulated $\lambda(y)$ curves, it was found that the α_1 values varied from 0.2650 to 0.6592; whereas the α_2 values had a range of 0.1612 to 0.3872.

The PoA values for $d(y)$ being a polynomial function is presented in Figure 3. In the figure, $PoA(SW)$, $n = 1$ refers to the average of the worst case PoA value of Social Welfare obtained from our simulations when the exponent of the polynomial function is 1 (linear). From the figure we observe that the $PoA(SW)$ has close to optimum value (slightly larger than 1) for $K = 0.3$ and then gradually increase to 2.5 with the increase of K . On the other hand, the $PoA(Rev)$ values at first start decreasing with the increase in K , and start increasing after reaching about $K = 0.45$. Therefore, $K = 0.45$ provides us a good PoA for both SW and Revenue.

For the queuing delay function, similar results can be observed, as shown in Figure 4. In the figure, U_f stands for the utilization factor, which is usually around 50-70% in real world scenario; hence we used those two corner values of that range. For this case we see that a value of $K = 0.4$ to 0.55 results in a small value of PoA for both SW and Revenue.

Table II has the theoretical and simulated (worst-case) PoA values for $K = 0.3, 0.5$, and 0.7. Since, theoretical values depend on both K and α , we used the minimum value of α_1 from the range of values found from the $\lambda(y)$ curves, which is 0.2650. The results with α_2 followed similar trend as α_1 and due to space constraints, is not discussed here. From the table it is apparent that in most cases the worst-case PoA value is well within the theoretical bounds and for a good choice of K (as mentioned in the Corollary IV.3.1 and IV.4.1), we can achieve very good PoA for both SW and Revenue.






Fig. 2. Sample $\lambda(y)$ curves with different α values. Fig. 3. Simulated PoA values (polynomial delay). Fig. 4. Simulated PoA values (queuing delay).

TABLE II
POA VALUES FOR POLYNOMIAL AND QUEUING DELAY FUNCTIONS.

Term	K=0.3	K=0.5	K=0.7
PoA(SW) (polynomial)	3.3851	2.4319	5.8607
PoA(SW) (queuing)	1.9968	2.4902	6.6428
PoA(SW) (Theo)	5.3908	7.5472	12.5786
PoA(Rev) (polynomial)	1.5511	1.8605	5.1399
PoA(Rev) (queuing)	5.5403	1.9556	2.0722
PoA(Rev) (Theo)	25.1572	15.0943	12.5786

VI. CONCLUSION

We analyze the traffic exchange equilibrium between ISPs at a profit-making IXP, and establish the existence of a pricing policy that ensures good social welfare and IXP revenue simultaneously. The policy only requires estimation of a worst-case sublinearity measure of the inverse demand curve, is very straightforward to compute and implement. Theoretical bounds on PoA of social welfare and revenue for the proposed pricing policy is given in terms of the sublinearity measure. PoA values obtained from simulations with inverse demand curves estimated from actual data on IXPs, fell well within these bounds, and the average of the worst case PoA values were less than two when K is chosen appropriately. For two broad classes of delay functions, we found from simulations that a value of K in the range of 0.45 to 0.55 ensures a good balance between social welfare and revenue. These recommended per-unit prices can be readily translated to pricing based on port capacity that IXPs typically implement in practice.

ACKNOWLEDGEMENTS

The authors should like to thank the National Science Foundation for supporting this work through award CNS-1816396.

REFERENCES

- [1] T. Böttger, G. Antichi, E. L. Fernandes, R. di Lallo, M. Bruyere, S. Uhlig, and I. Castro, "The Elusive Internet Flattening: 10 Years of IXP Growth," *arXiv e-prints*, 2018.
- [2] K. Russell, "Here's Why You Should Be Thrilled Netflix Is Paying Comcast For Content Delivery," *Business Insider*, Feb. 25, 2014.
- [3] P. K. Dey and M. Yuksel, "Peering Among Content-Dominated Vertical ISPs," *IEEE Networking Letters*, vol. 1, no. 3, pp. 132–135, 2019.
- [4] M. Chiesa, D. Demmler, M. Canini, M. Schapira, and T. Schneider, "SIXPACK: Securing Internet EXchange Points Against Curious Onlookers," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '17)*. New York, NY, USA: ACM, 2017, p. 120–133.

- [5] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, 2012, pp. 163–174.
- [6] Equinix, "Equinix Internet Exchange," retrieved on May 26, 2021. [Online]. Available: <https://www.equinix.com/interconnection-services/internet-exchange/>
- [7] M. O. Jackson, "A Survey of Network Formation Models: Stability and Efficiency," *Group formation in economics: Networks, clubs, and coalitions*, vol. 664, pp. 11–49, 2005.
- [8] E. Tardos and T. Wexler, "Network formation games and the potential function method," *Algorithmic Game Theory*, pp. 487–516, 2007.
- [9] E. Anshelevich and M. Hoefer, "Contribution games in networks," *Algorithmica*, vol. 63, no. 1-2, pp. 51–90, 2012.
- [10] G. Accongiagico, E. Altman, E. Gregori, and L. Lenzini, "Peering versus transit: A game theoretical model for autonomous systems connectivity," in *Multilevel Strategic Interaction Game Models for Complex Networks*. Springer, 2019, pp. 201–237.
- [11] M. I. I. Alam, E. Anshelevich, K. Kar, and M. Yuksel, "Proportional pricing for efficient traffic equilibrium at internet exchange points," in *Proceedings of the 33rd International Teletraffic Congress (ITC 33)*, 2021 (to appear).
- [12] E. Anshelevich, O. Bhardwaj, and K. Kar, "Strategic Network Formation Through an Intermediary," *Theory of Computing Systems*, vol. 63, no. 6, pp. 1314–1335, 2019.
- [13] E. Anshelevich, K. Kar, and S. Sekar, "Pricing to maximize revenue and welfare simultaneously in large markets," in *International Conference on Web and Internet Economics*. Springer, 2016, pp. 145–159.
- [14] Y. Zhang, M. Roughan, C. Lund, and D. L. Donoho, "Estimating Point-to-Point and Point-to-Multipoint Traffic Matrices: an Information-Theoretic Approach," *IEEE/ACM Transactions on Networking*, vol. 13, no. 5, pp. 947–960, October 2005.
- [15] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, "How Many Tiers? Pricing in the Internet Transit Market," *ACM SIGCOMM Computer communication review*, vol. 41, no. 4, pp. 194–205, 2011.

APPENDIX

A. Proof outline of Theorem IV.1:

From definition of equilibrium and the $\lambda(y)$ curve we know, equilibrium traffic sends demand with highest λ_{ij} first, until it equals the cost $c(y)$. Due to the properties of $c(y)$ and $\lambda(y)$, at least one such intersection point exists. $\square$

B. Proof outline of Theorem IV.2:

Let's consider a traffic vector $\vec{y}_t$ with $y_t = |\vec{y}_t|$, so that the traffic with largest λ_{ij} is sent. Then, from the definition of Social Welfare, we have that:

$$SW(\vec{y}_t) = 2 \int_0^{y_t} \lambda(y) dy - 2 \int_0^{y_t} E'(y) dy.$$

The above applies for both $y_t = y_e$ and $y_t = y_p$. To maximize social welfare, it is clear that y_p should be the intersection of the λ and E' curves (see Figure 1). $\square$

C. Proof outline of SW part of Theorems IV.3 and IV.4

From the discussion of *OPT* we know that if the IXP charges β_p , then we will get the *OPT SW*. Also, if $\beta_e < \beta_p$, then taking the max of these two means taking the price of β_p and $PoA(SW)$ will be 1. Hence, we just need to prove that our claim holds when $\beta_e > \beta_p$. For the case of $\beta_e > \beta_p$, we can easily argue (figure 1) that $y_e < y_p$. Hence, from the definition of $PoA(SW)$, we have,

$$\begin{aligned} PoA(SW) &= \frac{\int_0^{y_p} \lambda(y) dy - \int_0^{y_p} E'(y) dy}{\int_0^{y_e} \lambda(y) dy - \int_0^{y_e} E'(y) dy} \\ &= \frac{\int_0^{y_e} \lambda(y) dy - \int_0^{y_e} E'(y) dy + \int_{y_e}^{y_p} \lambda(y) dy - \int_{y_e}^{y_p} E'(y) dy}{\int_0^{y_e} \lambda(y) dy - \int_0^{y_e} E'(y) dy} \\ &\leq 1 + \frac{(y_p - y_e)(c_e(y_e)) - (y_p - y_e)E'(y_e)}{(c_e(y_e)) \cdot y_e - d(y_e) \cdot y_e} \\ &= 1 + \frac{(y_p - y_e)(\beta_e + d(y_e) - E'(y_e))}{(\beta_e + d(y_e)) \cdot y_e - d(y_e) \cdot y_e} \\ &\leq 1 + \frac{(y_p - y_e)\beta_e}{\beta_e y_e} = 1 + \frac{y_p - y_e}{y_e} \leq \frac{1}{y_e}. \end{aligned} \quad (6)$$

where in the fourth line we have used $\lambda(y_e) = c_e(y_e)$, and in the last line $y_p \leq 1$.

Now, If we use the α_1 sub-linearity property, then we have,

$$\begin{aligned} y_e + c_e(y_e) &\geq \alpha_1, \\ \text{or, } \frac{1}{y_e} &\leq \frac{1}{\alpha_1 - c_e(y_e)}, \end{aligned} \quad (7)$$

where in the second line $\lambda(y_e)$ is replaced by $c_e(y_e) = \beta_e + d(y_e)$. Now, if we choose, $c_e(y_e) = K\alpha_1$, then from equation 6 we get $PoA(SW) \leq \frac{1}{\alpha_1(1-K)}$.

Similarly if we use the α_2 sub-linearity property, we get,

$$\begin{aligned} (y_e)^{\alpha_2} + (c_e(y_e))^{\alpha_2} &\geq 1 \\ \text{or, } \frac{1}{y_e} &\leq \left(\frac{1}{1 - (c_e(y_e))^{\alpha_2}} \right)^{1/\alpha_2}, \end{aligned} \quad (8)$$

Now, if we choose, $(c_e(y_e))^{\alpha_2} = K$, then from equation 6 we get $PoA(SW) \leq \frac{1}{(1-K)^{1/\alpha_2}}$. $\square$

D. Proof outline of Revenue part of Theorems IV.3 and IV.4

We are going to prove the Revenue section for two cases, one where $\beta_e > \beta_p$ and the other where $\beta_e < \beta_p$.

1) *Case I:* $\beta_e > \beta_p$: Since, $\beta_e > \beta_p$, so $\max(\beta_e, \beta_p)$ is β_e , and that is the price which the IXP will charge. Also, we know that $E'(0) = 0$, and $d(y)$ is convex in nature, so $E'(y) \geq 2d(y)$. For, $\beta_e > \beta_p$ we have $E'(y_e) < c(y_e)$, which means $\beta_e > d(y_e)$. Hence, for the case of α_1 property where we chose $\beta_e + d(y_e) = K\alpha_1$, we get $\beta_e > 0.5K\alpha_1$. On

the other hand, with the α_2 property where we chose $(\beta_e + d(y_e))^{\alpha_2} = K$, we get $\beta_e > 0.5 \cdot c_e(y_e) = 0.5K^{1/\alpha_2}$.

Now, with β_e being the price charged by the IXP, we can have the following two sub-cases, $I - A) \beta_e < \beta^*$, and $I - B) \beta_e > \beta^*$, where β^* is the per-unit traffic price charged by the IXP that ensures maximum revenue. Also, let's define $c^*(y) = \beta^* + d(y)$, and y^* be the traffic that attains the maximum revenue when IXP is charging β^* for per-unit traffic.

Case I-A: The ratio of maximum revenue to equilibrium revenue can be bounded by the following way,

$$PoA(Rev) = \frac{\beta^* y^*}{\beta_e y_e} \leq \frac{\beta^* y_e}{\beta_e y_e} \leq \frac{1}{\beta_e} \leq \frac{2}{K\alpha_1}. \quad (9)$$

In the above equation, we have used the fact that $y^* \leq y_e$ because $\lambda(y)$ is a non-increasing curve. Also, the maximum value of β^* can be 1, which is used in the second inequality term. Also, if we had used α_2 property, then we get

$$PoA(Rev) \leq \frac{1}{\beta_e} \leq \frac{2}{K^{(1/\alpha_2)}}. \quad (10)$$

Case I-B: For this scenario, we can bound the $PoA(Rev)$ with similar equation as before,

$$PoA(Rev) = \frac{\beta^* y^*}{\beta_e y_e} \leq \frac{\beta^*}{\beta_e y_e} \leq \frac{1}{y_e}, \quad (11)$$

where we have used the property of $y^* \leq 1$ and $\beta^*/\beta_e \leq 1$. Hence, with our choice of pricing, that is $c_e(y_e) = K\alpha_1$, we get $PoA(Rev) = 1/(\alpha_1(1-K))$. On the other hand, with $c_e(y_e)^{\alpha_2} = K$; we get $PoA(Rev) \leq 1/(1-K)^{(1/\alpha_2)}$.

2) *Case II:* $\beta_e < \beta_p$: According to our pricing policy, β_p will be charged as the per-unit price for this case. Similar to Case I, there are two possible sub-cases, II-A) $\beta_p > \beta^*$, and II-B) $\beta_p < \beta^*$. We will prove that Scenario A is not a possible scenario and then will bound the $PoA(Rev)$ for Scenario B, which is $\beta_p < \beta^*$.

Case II-A: To prove that $\beta_p > \beta^*$ is not possible, let's assume y_t is some traffic such that $y_t > y_p$ (y_p is the *OPT*), and the corresponding Y axis value of y_t is $\lambda(y_t)$. Since, $y_p < y_t$ so, the price β_p for *OPT* should be the larger than or equal to β_t , the price for sending y_t traffic. Also, from the equation of revenue, we know it is $(c(y) \cdot y - E(y))$. Hence,

$$\begin{aligned} &Rev(OPT) - Rev(\lambda(y_t), y_t) \\ &= (\lambda(y_p)y_p - E(y_p)) - (\lambda(y_t)y_t - E(y_t)) \\ &\geq (\lambda(y_p)y_p - \lambda(y_t)y_t) + (E(y_p) - \lambda(y_t))(y_t - y_p) \geq 0. \end{aligned}$$

Hence, for any per-unit price less than the optimum price, the revenue can be at most equal to the *OPT* revenue. So, the maximum revenue per-unit price need to be at least equal to or larger than the optimum per-unit price, which means, we will always have $\beta_p < \beta^*$.

Case II-B: To prove this part we can follow similar procedure as in the proof of $\beta_e \leq \beta^*$. After going through similar arguments we will find $\beta_p \geq 0.5K\alpha_1$ and $\beta_p \geq 0.5K^{1/\alpha_2}$. Then the rest of the proof can be done by substituting β_e with β_p in the $\beta_e \leq \beta^*$ case. $\square$