

T4P: Hybrid Interconnection for Cost Reduction

Ignacio Castro and Sergey Gorinsky

Institute IMDEA Networks, Madrid, Spain

Abstract—Economic forces behind the Internet evolution have diversified the types of ISP (Internet Service Provider) interconnections. In particular, settlement-free peering and paid peering proved themselves as effective means for reducing ISP costs. In this paper, we propose T4P (Transit for Peering), a new type of hybrid bilateral ISP relationships that continues the Internet trend towards more flexible interconnections at lower costs. With a T4P interconnection, one ISP compensates the other ISP for their peering by providing this other ISP with a partial-transit service. In comparison to paid peering, T4P is able to reduce the combined transit/peering costs of an ISP due to the subadditive nature of transit billing. As a cost-effective alternative to existing interconnection types, T4P expands and strengthens the connectivity of the Internet, e.g., between content and eyeball networks. After analyzing incentives of ISPs to adopt T4P, we use real traffic data from several IXPs (Internet eXchange Points) to quantify the T4P economic benefits. Our evaluation confirms the promising potential of T4P.

I. INTRODUCTION

The Internet is an evolving ecosystem where a multitude of interconnected ISPs (Internet Service Providers) supports global connectivity of end users. In its early years, the ecosystem was essentially a hierarchy where smaller ISPs paid bigger ISPs for the universal Internet reachability via transit links. Subsequent massive emergence of peering enabled many ISPs to exchange their customer traffic over cost-effective settlement-free peering links [1], [2]. The evolution kept increasing the diversity of inter-ISP connection types and introduced partial-transit and paid-peering links [3], [4]. In contrast to the full transit, a partial-transit link offers access to only a fraction of the global Internet address space. With paid peering, one of the peering ISPs pays the other peer for exchanging their customer traffic.

To a large extent, the driving forces behind the interconnection evolution are economic. For example, if two ISPs exchange their traffic via a transit provider, their payments to the provider significantly exceed the cost of communicating the same traffic over a settlement-free link. The costs of the peering are mostly related to the infrastructure and labor of maintaining the physical interconnection, either as a direct link or through an IXP (Internet eXchange Point) [5]–[7]. The potential of settlement-free peering to reduce the costs for both peers does not mean that the ISPs will indeed establish and sustain such a relationship. For instance, the ISPs might view each other as competitors and be unwilling to reduce the costs of the counterpart. Furthermore, the costs of each party depend on the peering-link traffic and ISP sizes. Agreements for settlement-free peering commonly stipulate that the traffic flows in the two directions of the peering link should be balanced within a certain ratio (e.g., ratio 2:1) and

that the geographic scopes of the peering networks should be similar [2]. Loosening the above requirements, a paid-peering interconnection enables peering of diverse ISPs through monetary payments, e.g., by allowing one ISP to send more traffic and pay the other ISP a monetary compensation for the traffic imbalance.

In parallel to the interconnection evolution [8], the types of ISPs have evolved as well. Some ISPs run eyeball networks that primarily serve residential users. Other ISPs concentrate on providing Internet access for content providers such as Yahoo or YouTube [3], [9]. While popular content providers are the major sources of Internet traffic, an eyeball network acts mostly as a traffic sink. Peering between content and eyeball networks is potentially problematic not only because their traffic flows are unbalanced but also due to the heterogeneity of network types. The differences between content and eyeball networks complicate the issue of whether and how much one network should pay the other for their peering. Because the costs associated with last-mile infrastructures are typically high for the eyeball ISP (and significantly higher than the infrastructure costs of the content ISP), the eyeball ISP can view the high costs as a just cause for demanding a compensation from the content ISP. Moreover, since these high costs represent substantial barriers to entry in the eyeball-network market, the eyeball ISP can try leveraging its significant market power when negotiating a peering agreement with the content ISP [3], [9], [10]. On the other hand, the content ISP can be reluctant to compensate the eyeball ISP and even perceive such compensation demands as a violation of network neutrality [11]. The lack of clarity about proper conditions for eyeball-content peering has led to so-called peering wars [12]–[15] which disrupted the Internet connectivity.

In this paper, we propose *T4P (Transit for Peering)*, a new type of hybrid bilateral ISP interconnection that can reduce interdomain traffic costs and strengthen the Internet connectivity. In T4P, the link between two ISPs Y and Z carries not only peering but also some transit traffic. In particular, ISP Z communicates a portion of its transit traffic not through its transit providers but over the T4P link and through the transit providers of ISP Y. While ISP Y effectively becomes a partial-transit provider for this traffic of ISP Z, the partial transit serves as an in-kind traffic-delivery compensation paid by ISP Y to ISP Z for the peering. Unlike paid peering, T4P does not employ monetary payments. In comparison to paid peering, T4P is able to reduce the combined transit/peering costs of an ISP due to the subadditivity of transit billing. Our paper uses real traffic data to demonstrate the cost-reduction potential of T4P.




Fig. 1. Different types of interconnection between ISPs Y and Z: while double-arrow lines depict traffic flows, single-arrow lines show monetary flows.

The main contributions of the paper are the following ones:

- We propose T4P, a novel type of hybrid ISP interconnections where partial transit serves as a compensation for peering.
- Using real traffic data available at several IXPs, we quantify the financial incentives of ISPs to adopt T4P.

The rest of the paper is structured as follows. Section II provides additional background and motivation for the studied problem. Section III presents the concept of T4P in more detail. Section IV reports the economic model and analyzes the ISP incentives for T4P adoption. Section V evaluates T4P using the real IXP data. Section VI discusses related work. Finally, section VII concludes the paper with a summary of its contributions.

II. INCREASING DIVERSITY OF INTERCONNECTIONS

This section presents the background information on the Internet interconnection evolution that sets the stage for the T4P proposal.

Transit is the interconnection type that was heavily predominant in the early years of the commercial Internet. A transit link connects two ISPs called a provider and customer. The customer pays the provider for the traffic communicated in both directions of the interconnection. In exchange for the monetary payments, the provider offers the customer access to the global Internet. In a typical transit relationship, the provider is a larger network with a broader geographic scope. Figure 1a shows two examples of transit interconnections, with ISPs Y and Z acting as customers.

Subadditive billing is the standard settlement arrangement for transit [16]. Based on the traffic-rate samples collected for all short-term intervals (e.g., of 5 minutes long each) during the billing period (e.g., 1 month), the 95th-percentile traffic rate is calculated for each of the two link directions. The largest of the two 95th-percentile traffic rates serves as a billed traffic rate (although some billing versions use the sum rather than the maximum of the unidirectional rates). Then, a subadditive function is applied to the billed traffic rate to compute the monetary settlement. With the subadditive billing, larger traffic rates are usually billed at lower transit prices per Mbps.

Settlement-free peering has emerged as a cost-effective interconnection where two ISPs Y and Z exchange their customer traffic directly without any monetary compensation. Whereas the peering link carries only traffic of own customers, a vast majority of ISPs still needs transit links to reach the global Internet. Nevertheless, by reducing the traffic on the transit links, settlement-free peering reduces the transit costs for both ISPs Y and Z. Figure 1b depicts settlement-free peering.

Despite its definite potential for cost reduction, settlement-free peering has struggled to fully accommodate the increasingly diverse population of ISPs. In particular, the Internet evolution produced the two ISP types of content and eyeball networks with very different profiles in regard to the number and sizes of customers, dominant direction of traffic flows, cost structure, and market power. Eyeball networks receive more traffic than they send, serve more users, incur higher traffic-delivery costs, and enjoy a stronger bargaining position

because vendor lock-in is arguably easier with residential users than with content-providing customers [3], [9]. These differences make eyeball ISPs hesitant to peer with content ISPs on the settlement-free basis. For example, after content provider Netflix became a customer of Level 3, the traffic imbalance on the peering link between Level 3 and eyeball-network Comcast increased, and Comcast threatened to de-peer (i.e., terminate the peering agreement) with Level 3. Although Level 3 offered to resolve the conflict by upgrading its communication infrastructure and making its routing more beneficial for Comcast, the latter rejected the offer and de-peered.

Paid peering is a more flexible interconnection that allows ISP Y to monetarily compensate ISP Z for their peering. With respect to the traffic flows, paid and settlement-free peering are identical. Figure 1c illustrates paid peering.

III. T4P CONCEPT

While section II exposed the increasing interconnection diversity as well as economic factors behind this evolution, our T4P proposal continues the diversification trend to find an economically viable niche in the evolving Internet ecosystem. T4P is a novel type of hybrid bilateral ISP interconnection between diverse ISPs, such as content and eyeball networks. Unlike with paid peering, the T4P interconnection does not involve any monetary settlement. Instead, T4P employs in-kind compensations in the form of partial traffic transit. Specifically, ISP Y compensates ISP Z for their peering by providing a transit service for some traffic of ISP Z. Figure 1d depicts the T4P interconnection between ISPs Y and Z.

T4P is fundamentally different from paid peering not only because of eliminating any monetary compensation between ISPs Y and Z but also due to changes in the traffic flows. While paid peering is identical to settlement-free peering in restricting the link between ISPs Y and Z to own customer traffic, the hybrid T4P link combines the peering traffic with transit traffic. Hence, the T4P interconnection affects both transit routes and traffic rates along these routes.

The subadditive nature of transit billing is the reason why T4P is able to reduce the traffic costs of both ISPs Y and Z in comparison to paid peering. Although the overall transit traffic of ISPs Y and Z does not change with T4P, serving some transit traffic of ISP Z through ISP Y can decrease the overall transit costs of the two ISPs due to the billing subadditivity.

Whereas the subadditive billing can reduce the combined transit/peering costs of ISPs Y and Z, the attractiveness of T4P vs. paid peering for an individual ISP depends on the monetary settlement in the paid-peering interconnection. In particular, ISP Z finds T4P more attractive than paid peering only if the transit bill reduction for ISP Z with T4P is at least as large as the monetary settlement of paid peering.

IV. INCENTIVE ANALYSIS

In this section, we expand and formalize the incentives of ISPs to adopt T4P. While we envision T4P as a cost-effective

interconnection between diverse ISPs, paid peering serves as a natural baseline in our analysis.

The lack of real data on paid-peering settlements makes it problematic for our model to explicitly represent the monetary compensation paid by ISP Y to ISP Z in the paid-peering interconnection. Instead of treating the monetary compensation as an explicit parameter, our analysis relies on the key insight that the combined traffic costs of ISPs Y and Z are independent from this monetary compensation: the compensation paid by ISP Y is the same as the compensation received by ISP Z. Hence, the first step of our analysis focuses on the overall economic efficiency of T4P vs. paid peering for the two ISPs together.

Our model accounts for the subadditive billing of transit services. The subadditive billing method is a positive factor for T4P because adding the partial-transit traffic of ISP Z to the own transit traffic of ISP Y reduces the price paid by ISP Y per Mbps of the aggregated transit. We consider the method variant that bills the two directions of a link together by applying function f to the sum of the 95th-percentile traffic rates in the individual directions. In our model, billing function f always selects the CDR (Committed Data Rate) value that yields the smallest possible per-Mbps price for any traffic pattern. Using i to denote either ISP Y or ISP Z, we express transit cost C_i of ISP i as

$$C_i = f(A_i) \quad (1)$$

where A_i represents the billed bidirectional traffic pattern.

To represent the transit costs of ISPs Y and Z interconnected with T4P, let T denote the partial-transit traffic on the T4P link, and W_i be the own transit traffic of ISP i on the link with its normal transit provider. Then, transit costs F_Y and F_Z of ISPs Y and Z with the T4P interconnection are

$$F_Y = f(W_Y + T) \quad \text{and} \quad F_Z = f(W_Z - T). \quad (2)$$

With paid peering, transit costs N_i of ISP i equal

$$N_i = f(W_i). \quad (3)$$

Then, *transit cost reduction* R_i provided by T4P to ISP i is

$$R_i = N_i - F_i. \quad (4)$$

With the subadditive billing, f is a non-decreasing function. Thus, equations 2, 3, and 4 imply that T4P does not decrease the transit costs for ISP Y (i.e., $R_Y \leq 0$) and does not increase the transit costs for ISP Z (i.e., $R_Z \geq 0$).

We define *aggregate T4P gain* G as

$$G = R_Y + R_Z. \quad (5)$$

Combining equations 2, 3, 4, and 5, we express the aggregate T4P gain as

$$G = f(W_Y) + f(W_Z) - f(W_Y + T) - f(W_Z - T). \quad (6)$$

Due to the subadditive billing, the aggregate T4P gain is non-zero in general and depends on traffic T that ISP Z transits through ISP Y.

TABLE I
IXPS USED IN OUR DATA-DRIVEN EVALUATION OF T4P INSTANCES

Acronym	Name	Number of ISP members	Peak traffic rate, Gbps	Average traffic rate, Gbps
BIX	Budapest Internet eXchange	53	152	92
FICIX	Finnish Communication and Internet eXchange	25	32	19
InterLAN	InterLAN	63	22	11
NIX	Neutral Internet eXchange	54	116	76
SIX	Slovak Internet eXchange	52	42	23
IIX	Israeli Internet eXchange	17	2	1




Fig. 2. Network-management image for the peering traffic at an IXP: the x-axes are in hours; the y-axes are in Mbps; the blue lines and filled green areas represent the two directions of the traffic between two peering ISPs.

In comparison to paid peering, the overall economic efficiency of T4P for the two ISPs together is better when $G > 0$. Hence, we have proved the following theorem:

Theorem 1: For the two ISPs together, T4P is economically more attractive than paid peering when

$$f(W_Y) + f(W_Z) > f(W_Y + T) + f(W_Z - T). \quad (7)$$

Switching from the aggregate gain to the individual perspectives of ISPs Y and Z, one can think of R_Z as a monetary equivalent of the in-kind traffic-delivery compensation provided to ISP Z with T4P. Hence, ISP Z favors T4P when R_Z is larger than monetary compensation x paid by ISP Y to ISP Z in the paid-peering relationship. Even when R_Z is strictly greater than x , ISP Y also finds T4P more attractive as long as the transit cost increase ($-R_Y$) of ISP Y with T4P is smaller than x . There is a continuum of such mutually beneficial settings when aggregate gain G is positive.

V. DATA-DRIVEN EVALUATION

In section IV, we analyzed ISP incentives for adopting T4P. To quantify the economic attractiveness of the T4P, we evaluate it using real traffic data and real transit pricing. Section V-A presents our evaluation methodology. Section V-B

TABLE II
VOXEL TRANSIT PRICING

Committed Data Rate, Mbps	Price per Mbps
10	\$25
50	\$15
100	\$10
1000	\$5
10000	\$4

illustrates the potential benefits of T4P with an example. Section V-C evaluates T4P in more detail.

A. Evaluation methodology

Our evaluation relies on real traffic at six European IXPs presented in table I. Each of the IXPs reports peering traffic for its ISP members in the form of network-management images, such as the one in figure 2. Obtained by applying OCR (Optical Character Recognition) to the images [17], our numeric data for the peering traffic serve as a basis for approximating the transit traffic of the member ISPs [18]. Specifically, we scale up the peering traffic of a member ISP with the factor of 1.5 to represent the transit traffic of the ISP (real traffic data available at some academic ISPs validate such correspondence between the peering and transit traffic patterns). To evaluate T4P instances, we consider all possible pairs of ISPs at each IXP.

The 95th-percentile billing for transit services utilizes the Voxel prices in table II. Transit providers tend to treat their prices as confidential. Voxel, a North American ISP, is a rare exception and publishes its transit pricing [19]. Table II sums up the transit prices of Voxel with respect to the CDR (Committed Data Rate) chosen by the customer. The transit payment is calculated as the product of the price per Mbps for the chosen CDR (our paper uses \$ or USD to refer to United States dollars) and either 95th-percentile traffic rate or CDR when the latter is larger.

B. Illustrative example

In this section, we present an example that illustrates the potential benefits of T4P for a pair of ISPs Y and Z at NIX (Neutral Internet eXchange). The peering traffic of the ISPs is imbalanced with ISP Y sending more traffic at the ratio of 4:1.

TABLE III
ILLUSTRATIVE EXAMPLE OF T4P RELATIONSHIPS FOR TWO ISPs AT NIX

Interconnection type	Parameter	ISP Y	ISP Z
Paid peering	Billed transit traffic rate, Gbps	7.7	5.6
	Transit costs	\$38K	\$28K
T4P	Transit traffic rate, Gbps	12.2	1.1
	Transit costs	\$49K	\$6K
	Transit cost reduction	−\$11K	\$22K
	Maximum compensation to ISP Z	N/A	\$11K
	Maximum compensation reduction for ISP Y	\$11K	N/A

We examine the T4P instance that maximizes the aggregate T4P gain by shifting 80% of the transit traffic of ISP Z to the T4P link. As shown in table III, T4P decreases the transit costs of ISP Z by \$22K but raises the transit costs of ISP Y by \$11K. Thus, the maximum aggregate T4P gain attained by T4P is \$11K. ISP Y can provide up to a \$11K compensation to ISP Z without increasing its own overall traffic costs. This amount of $G = \$11K$ represents the budget of the T4P transit-cost benefits that can be distributed between ISPs Y and Z in various ways. The other extreme of this continuum is financially equivalent to settlement-free peering for ISP Z: the latter does not see any changes in its traffic costs but ISP Y decreases its overall traffic costs by \$11K.

C. Evaluation results

Now, we take a detailed look at the T4P relationships between all ISPs at the six IXPs. Figure 3 shows the distributions of aggregate T4P gain G . Such gains are in addition to those resulting from peering and arise due to the subadditive billing and the aggregation of transit traffic. At FICIX, 60% of its ISP pairs gain at least \$1K, with almost 5% of its ISP pairs gaining beyond \$10K. The percentage of ISP pairs with G above \$1K is greater than 40% for NIX and InterLAN, and larger than 20% for BIX and SIX. IIX provides the lowest gains among the 6 IXPs, with less than 5% of its ISP pairs gaining more than \$1K. Figure 4 focuses on the top 20% of the ISP pairs with the biggest G across all 6 IXPs. Around 150 ISP pairs have gains a larger G than \$5K, and about 45 ISP pairs gain more than \$10K.

VI. RELATED WORK

Throughout the paper, we already referred to essential pertinent research. This brief section takes a broader look at the related work. The T4P niche in the current Internet ecosystem should be viewed in the context of the evolving Internet ecosystem. The Internet has developed from a fundamentally hierarchical net and flattened due to the pervasive peering that characterized the Internet evolution [1], [2], [20]. The growing heterogeneity of ISPs is fundamental to the plausibility of T4P [3]. This growing diversity was accompanied by new forms of interconnections such as partial transit and paid




Fig. 3. Distributions of aggregate T4P gain G for T4P at the six IXPs.




Fig. 4. Top 20% of the ISP pairs with the T4P relationships that have the biggest aggregate gains across all six IXPs.

peering [2], [4]. While resource allocation and corresponding cost recovery in a distributed environment is a complex matter in general [21], the Internet evolution complicated the situation even further. Although tussles affected the Internet from its early days [22], this growing diversity together with the ubiquitous peering resulted in frequent disagreements over peering cost allocations. Traffic imbalances [23] in peering relationships led to demands of monetary compensations [2] and to a large set of techniques for minimizing costs and maximizing revenues [24]–[26]. The tensions were specially acute between networks that primarily connect residential users and those networks that connect content providers. Such rifts caused so-called peering wars [12]–[15]. These tussles resulted in heated debates about network neutrality and the role of paid peering in it [6], [9], [27]–[30]. Instead of imposing a fee based on the different revenue or cost structure as in [9], [10], T4P exploits the different cost structure of heterogeneous ISPs and succeeds in reducing overall costs without monetary payments.

With peering being a prominent cost-reduction technique, ISPs' profit seeking behaviors have led to other cost-reduction solutions such as IP multicast [31], CDNs (Content Distribution Networks) [32], P2P (Peer-to-Peer) localization [33], [34] and traffic smoothing [35], [36]. Unlike the above solutions that change the rate or timing of the overall transit traffic, T4P reduces the transit costs by redistributing the same traffic over the two transit links.

VII. CONCLUSION

This paper proposed and evaluated T4P (Transit for Peering), a new type of hybrid bilateral ISP interconnection. In T4P, one ISP compensates the other ISP for their peering by providing this other ISP with a partial-transit service. Leveraging the subadditive billing of transit services, T4P offers economic benefits over paid peering and thereby opens new opportunities for ISP interconnections. We modeled and analyzed T4P relationships with paid peering as the natural comparison baseline. Then, we evaluated T4P based on the real Voxel transit pricing as well as real traffic data available at the six European IXPs. Our results confirmed the potential of T4P to expand and strengthen the Internet connectivity through the more flexible cost-effective interconnection.

ACKNOWLEDGMENTS

We are truly grateful to Rade Stanojevic for providing data, software, and feedback. Syed Hasan helped us in preparing the original submission. This work was supported in part by Department-of-Education grant S2009/TIC-1468 from the Regional Government of Madrid, Ramon-y-Cajal grant RYC-2009-04660 from the Spanish Ministry of Science and Innovation, FP7-PEOPLE grant 229599 and FP7-ICT grant 288021 from the European Commission.

REFERENCES

- [1] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The Flattening Internet Topology: Natural Evolution, Unightly Barnacles or Contrived Collapse?" in *Proceedings of the International Conference on Passive and Active Network Measurement (PAM)*, 2008.
- [2] A. Dhamdhere and C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh," in *Proceedings of ACM CoNEXT*, 2010.
- [3] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy," in *Proceedings of the 35th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2007.
- [4] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. Vazirani, "How Many Tiers? Pricing in the Internet Transit Market," in *Proceedings of ACM SIGCOMM*, 2011.
- [5] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *Proceedings of ACM SIGCOMM*, 2009.
- [6] A. Dhamdhere and C. Dovrolis, "Can ISPs Be Profitable Without Violating Network Neutrality?" in *Proceedings of ACM NetEcon*, 2008.
- [7] A. Dhamdhere, C. Dovrolis, and P. Francois, "A Value-based Framework for Internet Peering Agreements," in *Proceedings of International Teletraffic Congress (ITC)*, 2010.
- [8] A. Dhamdhere and C. Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem," *ACM/IEEE Transactions on Networking*, vol. 19, no. 5, pp. 1420–1433, 2011.
- [9] R. Ma, D. Chiu, J. Lui, V. Misra, and D. Rubenstein, "Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement," in *Proceedings of ACM NetEcon*, 2008.
- [10] J. Musacchio, G. Schwartz, and J. Walrand, "A Two-Sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue," *Review of Network Economics*, vol. 8, no. 1, 2009.
- [11] R. Ma, D. Chiu, J. Lui, V. Misra, and D. Rubenstein, "On Cooperative Settlement Between Content, Transit, and Eyeball Internet Service Providers," *IEEE/ACM Transactions on Networking*, vol. 19, no. 3, pp. 802–815, 2011.
- [12] M. A. Brown, A. Popescu, and E. Zmijewski, "Peering Wars: Lessons Learned from the Cogent-Telia De-peering," in *NANOG 43*, 2008.
- [13] M. A. Brown, C. Hepner, and A. Popescu, "Internet Captivity and the De-peering Menace," in *NANOG 45*, 2009.
- [14] A. Popescu and T. Underwood, "D(3)peered: Just the Facts Ma'am," in *NANOG 35*, 2005.
- [15] N. Anderson, "Peering Problems: Digging into the Comcast/Level 3 Grudgematch," <http://arstechnica.com/tech-policy/news/2010/12/comcastlevel3.ars>, Ars Technica, 2011, [Online].
- [16] X. Dimitropoulos, P. Hurley, A. Kind, and M. Stoecklin, "On the 95-percentile Billing Method," *Proceedings of the International Conference on Passive and Active Network Measurement (PAM)*, 2009.
- [17] "T4P: Hybrid Interconnection for Cost Reduction," http://fourier.networks.imdea.org/~ignacio_decastro/t4p.html, [Online].
- [18] R. Stanojevic, I. Castro, and S. Gorinsky, "CIPT: Using Tuangou to Reduce IP Transit Costs," in *Proceedings of ACM CoNext*, 2011.
- [19] Voxel dot net, pricelist. Accessed in January 2011. <http://www.voxel.net/ip-services>. [Online].
- [20] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An Agent-based Model of Interdomain Network Formation, Traffic Flow and Economics," in *Proceedings of IEEE INFOCOM*, 2012.
- [21] S.-W. Cho and A. Goel, "Pricing for Fairness: Distributed Resource Allocation for Multiple Objectives," *Algorithmica*, vol. 57, no. 4, pp. 873–892, 2010.
- [22] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 462–475, 2005.
- [23] W. B. Norton, "Internet Service Providers and Peering," in *NANOG 45*, 2001.
- [24] —, "The Internet Peering Playbook: Connecting to the Core of the Internet," 2002.
- [25] R. Stanojevic, N. Laoutaris, and P. Rodriguez, "On Economic Heavy Hitters: Shapley Value Analysis of 95th-percentile Pricing," in *Proceedings of ACM IMC*, 2010.
- [26] P. Banger and S. Gorinsky, "Impact of Prefix Hijacking on Payments of Providers," in *Proceedings of COMSNETS*, 2011.
- [27] J. Crowcroft, "Net Neutrality: the Technical Side of the Debate: a White Paper," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 1, pp. 49–56, 2007.
- [28] T. Wu, "Network Neutrality, Broadband Discrimination," *Journal on Telecommunications and High Technology*, vol. 2, pp. 141–141, 2003.
- [29] R. M. Frieden, "Network Neutrality or Bias? Handicapping the Odds for a Tiered and Branded Internet," *Bepress Legal Series*, 2006, working paper 1755.
- [30] M. Yuksel, K. Ramakrishnan, S. Kalyanaraman, J. Houle, and R. Sadhvani, "Class-of-service in IP Backbones: Informing the Network Neutrality Debate," in *Proceedings of ACM SIGMETRICS*, 2008.
- [31] S. Deering, "Multicast Routing in a Datagram Internetwork," Ph.D. dissertation, Stanford University, 1991.
- [32] L. Qiu, V. Padmanabhan, and G. Voelker, "On the Placement of Web Server Replicas," in *Proceedings of IEEE INFOCOM*, 2001.
- [33] D. R. Choffnes and F. E. Bustamante, "Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-peer Systems," in *Proceedings of ACM SIGCOMM*, 2008.
- [34] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider Portal for Applications," in *Proceedings of ACM SIGCOMM*, 2008.
- [35] N. Laoutaris, G. Smaragdakis, P. Rodriguez, and R. Sundaram, "Delay Tolerant Bulk Data Transfers on the Internet," in *Proceedings of ACM SIGMETRICS*, 2009.
- [36] M. Marcon, M. Dischinger, K. Gummadi, and A. Vahdat, "The Local and Global Effects of Traffic Shaping in the Internet," in *Proceedings of COMSNETS*, 2011.