Project Summary

The Internet is an ecosystem of about 50,000 *Autonomous Systems* (or ASes) that operate independently, having different objectives and operational constraints. What glues the Internet together is the bilateral techno-economic agreements that form the interconnections between these ASes. These interconnections have evolved over the last 20 years or so, since the commercialization of the Internet in the mid-nineties, in a rather ad-hoc manner, often resulting in bilateral or multilateral disputes about who should peer with whom, whether one of the two parties should pay the other, and about the conditions that an interconnection should satisfy (e.g., balanced traffic ratios). These problems result in congested interconnections and, in some cases, unreachability problems that can affect millions of Internet users.

The main premise of this research is that the currently deployed framework of Internet interconnection has fundamental weaknesses and systemic problems that inevitably will continue causing peering disputes between ASes. Instead of looking at each peering conflict as an isolated incident, we need to investigate thoroughly the limitations of the interconnection framework that is currently in place, and probably re-design it through an inter-disciplinary, techno-economic perspective.

Intellectual Merit: The research objectives of this project are three-fold. First, to analyze four common distinct interconnection scenarios, and to investigate their limitations under a general but realistic modeling framework. These scenarios cover the cases of monopoly and oligopoly in the Internet access market, as well as the possibility of one or more transit providers in the path between access and content providers. The modeling framework is based on recent developments in economics such as the framework of two-sided markets and the theory of incomplete contracts. The analytical modeling results will be complemented by computational results using the agent-based simulator GENESIS that was developed by the PIs in earlier work.

The second objective is to analyze data provided by Comcast and other sources so that we can estimate the key parameters of the previous economic models and validate their predictions. These parameters include the elasticity of demand for Internet traffic at the access and transit markets, the sensitivity of the demand to congestion, and the traffic growth rate for different Internet service plans.

The third objective is to develop a new techno-economic interconnection framework that can provide a broader and more economically efficient set of interdomain relations than just transit and (settlement-free) peering. The proposed interconnection framework should also provide the right incentives so that all relevant Internet firms continue to invest sufficient resources to Internet infrastructure. The design of the new interconnection framework will be guided by the modeling and empirical results of the first two research objectives.

Broader Impact: The educational activities for this project will enhance undergraduate, graduate and professional education with concepts that lie at the intersection of networking, game theory, economics, and techno-economic agent-based modeling. Outreach activities will disseminate the results of this project to operator communities (e.g., NANOG), regulator bodies (e.g., FCC) and related techno-economic forums (e.g., the World Economic Forum's "Global Agenda Council on the Future of the Internet"). The project will have a significant inter-disciplinary component in both its inputs and outputs. Specifically, much of the prior work behind this research (inputs) has originated in economics and game theory. Conversely, the project's results (outputs) are likely to be relevant to any other discipline that is dealing with large-scale dynamic networks of autonomous and strategic agents.

Key Words: Internet economics; Transit and peering interconnections; Peering disputes; Two-sided markets; Contract theory; Economics of complex networks.

1 Motivation

The Internet, a network of about 50,000 Autonomous Systems (ASes), has undergone major transformations in the last decade. Most of today's Internet traffic originates from a few large content providers and their (own or partner) CDNs In 2013, Sandvine [103] reported that about one half of all downstream consumer traffic came from Netflix or Youtube. The transition from a *hierarchical* to a *flat* Internet [39, 50, 5], along with the rise of "supergiant" content providers constitutes the rise of a new market structure in the Internet ecosystem. Content providers and CDNs can control the source (and thus the path) of data coming into an ISP, and so they can increase (or alleviate) loading and congestion on different points of interconnection. At the same time, ISPs along with their peering partners jointly control the capacity of incoming links, which can limit the actual options for the delivery of high-volume traffic. This tussle between content and access providers has led to interest from techologists and policy makers in the bargaining, money flow, and market power issues behind Internet interconnections. Peering disputes over traffic imbalances are not new – several peering disputes between large transit ISPs have happened over the years, e.g., [3, 36, 99]. More recent peering disputes, however, have been fueled by exploding demand for streaming video, and growing concentration of content among a few content providers and distribution networks [18, 45, 46, 115, 23, 105, 4, 24, 64], raising questions about appropriate network management, interconnection business strategies, and the impact of these peering disputes on end-user performance.

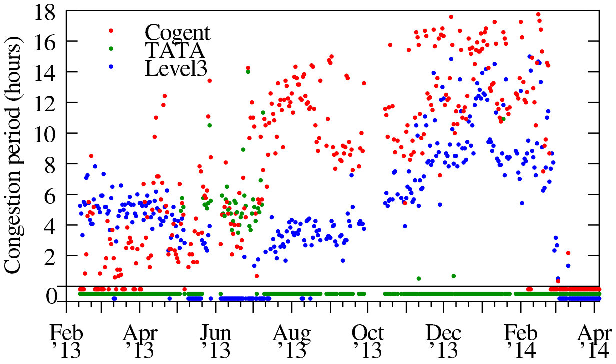

Figure 1: Evidence of congestion on three links between Comcast and neighbor networks that carried Netflix traffic in 2013 and 2014.

Peering disputes can persist for weeks or months; a case in point is the ongoing dispute between Netflix and some large broadband access providers in the US. Netflix and these providers have not been able to agree on the appropriate settlement structure for direct interconnection. This led to a situation where Netflix routed traffic towards these access networks via intermediate transit providers. Then, the transit and access ISPs could not agree on who should pay for the required capacity upgrades, and so several peering links between them showed evidence of increasing congestion over time. Figure 1 shows the duration of persistent congestion on links between Comcast and three transit providers (Cogent, Tata and Level3) that Netflix used to route traffic towards Comcast. Congestion on these links increased steadily from early 2013 to early 2014. At its worst, some of these links were congested for up to 18 hours a day. The Measurement Lab (M-Lab) project has measured similarly poor performance at several US broadband access networks during 2013 and 2014 [85]. Modern-day peering disputes are much more than minor squabbles between Internet companies; these disputes directly affect millions of people who rely on the Internet as a critical infrastructure for work and entertainment.

An important point to note about the aforementioned peering disputes is that while they result in performance degradation, they are mostly *economic issues*. The technical solutions are straightforward – to increase the capacity of the network (or shift traffic to other routes) as demand increases. The key questions, however, relate to which party should pay for infrastructure upgrades and how the costs and benefits of interdomain relationships should be split among the different parties.

The main premise of this research is that the currently deployed framework of Internet interconnection has fundamental weaknesses and systemic problems that inevitably will continue causing peering disputes between ASes. Instead of looking at each peering conflict as an isolated incident, we need to investigate thoroughly the limitations of the interconnection framework that is currently in place, and probably re-design it through an inter-disciplinary, techno-economic perspective. The research objectives of this project are three-fold. First, to analyze four common distinct interconnection scenarios, and to investigate their limitations under a general but realistic modeling framework (Section 4). These scenarios cover the cases of monopoly and oligopoly in the Internet access market, as well as the possibility of one or more transit providers in the path between access and content providers. The modeling framework is based on recent developments in economics such as the framework of two-sided markets and the theory of incomplete contracts. The analytical modeling results will be complemented by computational results using the agent-based simulator GENESIS that was developed by the PIs in earlier work. The second objective is to analyze data provided by Comcast (see attached letter of commitment) and other sources so that we can estimate the key parameters of the previous economic models and validate their predictions (Section-5). These parameters include the elasticity of demand for Internet traffic at the access and transit markets, the sensitivity of the demand to congestion, and the traffic growth rate for different Internet service plans. The third objective is to develop a new techno-economic interconnection framework that can provide a broader and more economically efficient set of interdomain relations than just transit and (settlement-free) peering (Section-6). The proposed interconnection framework should also provide the right incentives so that all relevant Internet firms continue to invest sufficient resources to Internet infrastructure. The design of the new interconnection framework will be guided by the modeling and empirical results of the first two research objectives.

2 Background

The Internet is a dynamic and self-organized *network of networks*. The larger individual networks are often registered as *Autonomous Systems* or ASes, so that they can have their own provider-independent addresses, and to have more than one direct interconnection with other ASes. Individual users or smaller businesses, on the other hand, are typically connected to the Internet through another AS (e.g., their residential ISP or the AS they work for). The interconnections between these ASes are also highly dynamic because they are determined by the economic, performance or strategic objectives of ASes.

Autonomous Systems are often classified based on their main functional role or business objective. We summarize this classification next, emphasizing at the outset that many ASes which at one time fit neatly into a single one of the categories below today play multiple roles in the ecosystem.

Enterprise Networks: Most ASes fall in this category. They are typically corporations or organizations that connect to the Internet with their own, provider-independent addresses.

Access Providers: These are firms that sell Internet access to residential and business customers (mostly through broadband technologies such as DSL or DOCSIS, but possibly also through fiber optics or wireless connections).

Content Providers: These are firms that generate Internet content, such as online video, news, e-commerce, online social networking, or Web search. The revenues of these firms are generated mostly through user subscriptions, advertisements, and online sales.

Transit Providers: These are firms that operate geographically large and high-capacity backbone networks. Historically, transit providers were paid by all other types of ASes to transfer data over large distances.

Content Distribution Networks (CDNs): These are firms that replicate Internet content in their distributed storage infrastructure ("caches"), serving download requests from locations (typically third party IXPs, explained below, where they have deployed caches) that are close to end-users. The customers of CDNs are typically content providers.

Internet Exchange Points (IXPs): These companies operate facilities in which different ASes can be co-located and interconnect with each other (if they choose to do so). IXPs are are paid by the ASes that use these facilities.

The previous classification system can be misleading today, because the largest ASes try to be more independent and versatile, playing multiple roles at the same time. For instance, major content providers (e.g., Netflix, Amazon) have developed their own CDNs. In some cases (e.g., Google), content providers operate their own international backbone networks. Some transit providers (e.g., Level3) have also diversified their role by offering CDN services, while certain access providers (e.g., Comcast) have deployed large, high-capacity backbone networks so that they now offer transit services. This diversification of the business roles and functions of ASes has major implications for the economics of Internet interconnections. We can no longer determine who should pay whom based on the single business function that has been historically associated with each AS. By the same token, old classifications no longer have much relevance, including the notion of “Tier-1” ASes. Today, every interconnection must be evaluated independently by the interconnecting parties to determine what its terms should be, examining the costs and benefits that that interconnection brings to each party.

An interconnection between two ASes can be of different types. In the following, we review the major types of AS interconnections today:

(Global) Transit: This is an asymmetric relation in which one AS is the “customer” and the other is the “provider”. The provider offers the customer routes that can reach any network in the Internet, and it *advertises* the addresses of the customer to the rest of the Internet. The customer pays the provider for the traffic it sends to and receives from the Internet.

Settlement-Free Peering (SF-peering): This is a symmetric relation wherein the two peering ASes agree to exchange traffic that is destined to them or their customers “for free”, based on the mutual exchange of transport value each AS obtains from the other. Traditionally, a rough balance of traffic has been deemed a necessary element for many SF-peering relationships, as well as roughly equivalent network facilities and sufficient traffic to merit dedication of one or more high capacity links.

Paid Peering (Paid-peering): This type of interconnection (sometimes also referred to as “on-net transit”) can be thought of as an intermediate solution between transit and SF-peering relations. A paid-peering interconnection is asymmetric: one party is the customer and the other is the provider; Similar to SF-peering relations, the only traffic that can be exchanged is traffic flowing between the two ASes or their customers.

Other kinds of interconnections: As the Internet is evolving, the level of sophistication in the available interconnection types has been increasing to meet more specialized needs and cost structures. For instance, a Partial-Transit relation provides transit service but for only a subset of the global routes or ASes.

3 Related work

Two-sided markets and net neutrality: The analysis of Internet interconnections can benefit from the recent modeling breakthroughs in analyzing two-sided platforms [26, 15, 100, 101, 57, 117]. Some of the more recent papers in this vein have analyzed the incentives of two-sided platforms to invest in capacity and quality of service under demand uncertainty and network congestion [47, 21, 27, 66].

Some of the issues that will be considered in this project are related to the ongoing debate about the value of net-neutrality regulation for the Internet. Most of the economic models addressing the costs and benefits of net neutrality regulation employ the two-sided markets framework. Good surveys of the theoretical and empirical literature on net neutrality can be found in [104] and [67]. One of the concerns about net neutrality regulation is that it may lower the incentive of ISPs to invest. Some recent works study models with congestion-sensitive content providers and a monopolistic ISP who undertakes investment in capacity and can provide access to a prioritized service at a fee [96, 33, 42, 66, 98]. Bourreau et al. [21] and Njoroge et al. [89] analyze models with investments undertaken by competing ISPs and study whether the policy concerns surrounding the net neutrality debate are alleviated when there is competition between Internet platforms. Economides and Tåg consider a two-sided market model in which one price can be positive while the other is constrained to zero under network neutrality regulation [43]. Within that framework they analyze the implications of two-sided pricing made possible by the violation of the network neutrality principle and the interaction between the two sides generated by cross-group externalities.

Economics of interconnections: The proposed research also relates to the broader economics literature on interconnection in the telecommunication market, which was started by Armstrong [14] and Laffont, Rey, and Tirole [70, 71]. These papers show that if network operators are constrained to offer only linear retail tariffs to their subscribers, they will negotiate collusive interconnection charges exceeding the associated interconnection costs and the efficient level of the interconnection charge. Laffont et al. analyzed access charges chosen by rival interconnected networks and their effect on communication costs between content providers and Internet subscribers [69].

Incomplete contracting and investments under uncertainty: The choice of peering arrangement and of the parameters of the interconnection contract between two ASes is closely related to their incentives to undertake investments. The characteristics of these investments (large scale, spillovers of investment benefits between the interconnecting ASes, uncertainty of demand) suggest that the peering parties may have incentives for cooperative or joint investments [22]. The incomplete contracting literature has demonstrated that the investment "hold-up problem" can be remedied by various organizational interventions [65, 54, 6, 8]. Moreover, under certain conditions, an appropriately chosen simple (incomplete) contract can provide the right incentives for the efficient level of selfish investments [34, 7, 90, 44]. However, the investments by interconnected ASes are complicated by the fact that such interconnections create not only direct (internal) benefits for the investor but also external benefits for the interconnected partner [31]. The more recent theoretical literature has shown that the first-best investment can be attained only if the external investment benefits are small relative to the internal benefits, and sufficiently symmetric across the interconnected parties [55].

Internet-related network formation and pricing models: In earlier work, we have proposed a fair paid-peering framework [41]. Game theoretic justifications of SF-peering have been establilshed by Baake et al. [17] and Badasyan et al. [19]. Shrimali and Kumar have evaluated conditions and pricing schemes for paid-peering [107, 108, 109]. Lippert et al. analyzed asymmetric peering interconnections [73]. Ma et al. proposed the use of the "Shapley value" framework to avoid peering disputes [82, 84], while the evolution of the Internet from a game theoretic perspective has been studied by Meirom et al. [86] and Ma et al. [83]. Shakkottai et al. asked whether simple pricing schemes can approximate complex revenue-maximizing pricing [106]. Transit pricing and the provisioning of different service tiers in that market has attracted significant interest recently [59, 37, 111, 113]. Anshelevich et al. proposed a network formation model considering the business effects of transit versus peering relationships [10, 11]. Other Internet-related network formation models have been proposed by Johari et al. [63, 62], and Arcaute et al. [12, 13].

GENESIS and computational agent-based network formation models: In earlier work, have developed and applied the GENESIS simulation framework - an agent-based computational model that captures detailed aspects of strategic formation and removal of both transit and SF-peering links [74, 75, 77]. Lodhi et al. use a variant of GENESIS to investigate the reasons and consequences of "Open Peering" adoption by many transit providers [76]. Other papers that have relied in a similar computational approach are [29, 60, 39]. This project will leverage GENESIS in the modeling of both existing (Section-4) and proposed (Section-6) interconnection scenarios.

Measurement work that is relevant to Internet interconnection: A major effort in the last 15 years has attempted to map the structure of the Internet topology, particularly at the AS granularity [110, 95, 40, 51, 81]. Although these public data sources provide a reasonable view of customer-provider relationships between networks, they miss a great deal of pervasive regional and non-revenue-based peering activity [118, 30, 28, 93, 102, 94]. In particular, this kind of instrumentation does not typically capture peering at Internet Exchanges (IXes), much of which is not governed by a formal contract [116]. Most larger IXes are public Internet Exchange Points (IXPs) where participants can interconnect across a shared and sometimes distributed fabric. Recent studies by us and others have discovered evidence of more peering connections at IXPs [5, 53] than were previously documented to exist in the entire Internet [16, 40, 32].

Characterization of the interdomain traffic flow is another important problem, as the flow of traffic influences the flow of money. Labovitz et al. report a consolidation of Internet traffic to few ASes [68]. Gehlen et al. report similar results about web traffic flow [49]. Mikians et al. focus on statistical properties of the interdomain traffic matrix [87]. Gursun et al. investigate the relationship between an AS's position in the Internet hierarchy and its ability to infer elements of that matrix [56].

4 Proposed work: Economic analysis of interconnection scenarios

In this section, we describe four increasingly more realistic and complex interconnection scenarios. We are planning to model and analyze these four scenarios based on an economic framework, as described next. The economic models will be parameterized and validated based on the empirical work described in Section 5. Additionally, the analytical results will be complemented with computational results using the agent-based AS-level simulator GENESIS, which was developed by the PIs in earlier work [74, 75, 77].

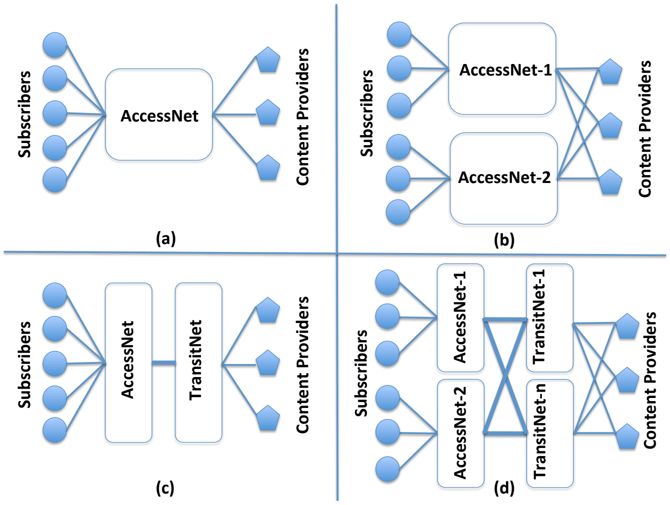

Figure 2: The four interconnection scenarios we plan to model and analyze.

4.1 Access monopoly with direct Access-Content routes (Fig-2a)

Suppose that in a certain market there is only one Internet Access provider (AccessNet). This is true for instance in some rural areas in the US. We can assume that even though AccessNet is a monopoly in this market, it operates under the oversight of a regulator (e.g., FCC). Internet users have the option to subscribe or not subscribe for service. The Internet content that subscribers want to download is provided by one or more Content providers (lets call them ContentNets). A ContentNet may be an original content provider or a Content-Distribution-Network (CDN).

In this first model, ContentNets are directly connected with AccessNet. Their interconnection can be SF-peering or paid-peering. The motivation for the former is to reduce the upstream transit fees that both parties would need to pay if they were not directly connected. If AccessNet has an expanded backbone network however, it may be able to connect directly with ContentNets without requiring an intermediate transit provider. In that case, AccessNet can request a paid-peering interconnection with each ContentNet in which the latter will pay AccessNet for the traffic they send to AccessNet's subscribers.

An important sequence of questions that we will investigate with this model in a two-sided market framework: Should both sides of the market (subscribers and ContentNets) pay for the services provided by AccessNet? And if both sides pay, how should AccessNet allocate the prices between subscribers and ContentNets? How does this price allocation depend on the subscribers' or the ContentNets' elasticity of demand for access? How does it depend on the costs incurred by the ContentNets and AccessNet for transferring this content to subscribers? Is the answer to the previous questions dependent on the traffic growth rate, and consequently, on how the costs of each party increase over time?

In the following, we present a first attempt to model this interconnection scenario; we emphasize that more sophisticated and realistic models will be investigated during the course of this research. Let us consider a two-sided market with a monopolistic access platform (AccessNet) serving exogenously fixed populations of subscribers and ContentNets, with the size of each population normalized to 1 (as done in [15, 100, 101]). In the absence of fixed fees, the subscribers' (or ContentNets') demand for access depends only on the price p^S (or respectively, p^C) charged by AccessNet: $D^S(p^S)$ (or respectively, $D^C(p^C)$). The access demand functions for the two sides are assumed to be concave.

AccessNet aims at maximising its own profit. It will choose a price structure for subscribers p^S and ContentNet p^C so as to maximise the profit function:

$$\pi = \left(p^S + p^C - c\right) D^S(p^S) D^C(p^C)$$

where c is the marginal cost of carrying a unit of traffic from a ContentNet to a subscriber. The profit maximizing prices for a two-sided monopoly access provider are given by a variation of the Lerner equation:

$$p^S = \frac{\eta^S}{\eta^S + \eta^C - 1} c, \qquad p^C = \frac{\eta^C}{\eta^S + \eta^C - 1} c$$

where $\eta^S = \frac{p^S (D^S)'}{D^S}$ and $\eta^C = \frac{p^C (D^C)'}{D^C}$ are the elasticities of demand for access for subscribers and ContentNets, respectively.

The previous model can be considered over-simplified because it assumes that a ContentNet will have no other option but to agree to this payment to AccessNet. An advanced version of this model will also consider *the effects of congestion* when a ContentNet is not willing to pay AccessNet for a paid-peering interconnection. The result would probably be an under-provisioned interconnection between AccessNet and that ContentNet in which the capacity is not upgraded regularly and so peering links become congested. Such congestion can hurt both AccessNet and ContentNet for the following reasons: subsribers may choose to terminate their service with AccessNet if their Internet access performance is too bad. Or, subscribers may download less content from that ContentNet because of congestion. Here, we need to consider the sensitivity of subsribers to congestion and how elastic is the demand in the presence of congestion. Additionally, the magnitude of congestion depends on the number of remaining subscribers (a negative externality). Who "hurts more" in this situation will largely affect whether AccessNet or ContentNet will back-off from their position and agree to the other party's terms.

4.2 Access oligopoly with direct Access-Content routes (Fig-2b)

In this interconnection scenario, subscribers can choose between a small number of access providers. In most US markets, there are 2-3 choices for Internet access. Let us consider two access providers, AccessNet-1 and AccessNet-2. These two providers may differ in terms of their service plans (how many Mbps of access capacity they provide) and in terms of price. Further, they may differ in terms of how congested their (SF-peering) interconnections with different content providers are. Congestion is a dynamic effect subject to negative externalities and so the relative performance of the two access providers can vary over time.

Subscribers, in this model, can switch between the two access providers. They can also choose to not subscribe to any service. To be more realistic, we will capture that subscribers encounter a penalty every time they switch – this may be due to a contractual penalty or simply due to the overhead of switching to a different provider or access technology.

Both AccessNet-1 and AccessNet-2 are directly connected to ContentNets. However, these interconnections may be different. For instance, the connection between a ContentNet and AccessNet-1 may be paid-peering while the connection between that ContentNet and AccessNet-2 may be SF-peering. Or, the paid-peering price of the two interconnections may be different. This implies that an AccessNet-1 subscriber may encounter congestion when downloading some content from that ContentNet, while an AccessNet-2 subscriber does not see any congestion for the same content.

Another important consideration in this scenario is that interconnection agreements are dynamic in nature. So, at the same time that subscribers can switch between access providers, the latter can also change dynamically their peering agreements with ContentNets. For instance, an interconnection may switch between paid-peering and SF-peering over time. The questions that we asked in the previous section apply also to this interconnection scenario. The results will probably be qualitatively different however, as a result of competition between the two access providers and of the resulting dynamics.

Next, we describe a first attempt to model this market structure with a highly stylized two-sided market framework that captures the competition between the two access providers in a static context (i.e., we do not capture any of the dynamics mentioned earlier); more realistic models will be developed during the course of this research. We assume that there is a unit mass of subsribers and a continuum of ContentNets. When a subsriber receives content from a ContentNet, that connection generates a surplus of $a > 0$ to the subsriber and a surplus of $b > 0$ to ContentNet. Each subsriber also obtains an intrinsic value u from connecting to the Internet irrespective of the AccessNet she uses. The AccessNets' costs of providing service to a ContentNet (Subscriber) is $c^{CN}(c^S)$.

ContentNets are assumed to be able to multi-home to both access providers while Subsribers are connected to one access provider, as commonly done in practice. ContentNets multi-home as long as this gives them a higher profit than single-homing. We assume that each ContentNet should incur a fixed cost per AccessNet and this cost is uniformly distributed over $[0, \infty)$ with density $f > 0$. The total mass of ContentNets is large enough so that there are always ContentNets who decide to not connect to any platform.

Subsribers are uniformly distributed on a line between zero and one. The horizontal product differentiation between AccessNets is captured by the assumption that AccessNet-1 (AccessNet-2) is located at the left (right) extreme point of the line. For simplicty we assume that AccessNet $i = 1, 2$ charges an access fee s_i to subscribers and p_i to ContentNets (a more elaborate model will include a usage fee). Let n_i denote the number of ContentNets with a paid-peering connection to AccessNet i.

Given (s_1, s_2), the Subscriber who is indifferent between the two AccessNets is given by

$$u + an_1 - tx - s_1 = u + an_2 - t(1 - x) - s_2, \tag{1}$$

which implies that the number of subscribers of AccessNet-1 is given by

$$x_1 = \frac{1}{2} + \frac{a(n_1 - n_2) - (s_1 - s_2)}{2t} \tag{2}$$

where t is a parameter that determines how closely substitutable are the services of the two AccessNets from the point of view of subscribers. The number of subscribers of AccessNet 2 is $x_2 = 1 - x_1$.

The fixed cost of the ContentNet who is indifferent between subscribing for paid-peering with AccessNet i or not is equal to n_i/f and is determined by:

$$x_i b - p_i - n_i/f = 0. \tag{3}$$

Given (s_1, s_2, p_1, p_2), we determine the allocation (x_1, x_2, n_1, n_2) from equations (2) and (3):

$$x_i = \frac{1}{2} - \frac{af(p_i - p_j) + (s_i - s_j)}{2(t - abf)}.$$

AccessNet i maximizes its profit given by

$$\pi_i = (p_i - c^{CN})n_i + (s_i - c^S)x_i.$$

Assuming that the cost values and utility parameters are such that the AccessNets' profit functions are concave, the equilibrium prices are given by

$$p_1 = p_2 = \frac{(b - a + 2c^{CN})}{4}; \quad s_1 = s_2 = t + c^S - \frac{bf(3a + b - 2c^{CN})}{4}.$$

Therefore, at equilibrium, AccessNets 1 and 2 obtain equal shares of the subcriber market ($x_1 = x_2 = \frac{1}{2}$) and profits:

$$\pi_1 = \pi_2 = \frac{t}{2} + \frac{f[c^{CN}]^2}{4} - \frac{f[4ab + (a + b)^2]}{16}.$$

The number of multi-homed ContentNets is given by

$$n = \frac{f(b + a - 2c^{CN})}{4}.$$

The previous theoretical model can be used as a basis for formulating testable hypotheses regarding the price structures that arise under various paid-peering arrangments between the access provider and content providers.

4.3 Access monopoly with a single Access-Transit-Content route (Fig-2c):

In this interconnection scenario, we return to the case of a single access provider (AccessNet). Content is provided by a number of content providers (ContentNet-x). The difference is that AccessNet is connected with ContentNets *indirectly*, through a *transit provider* (TransitNet). The interconnection between TransitNet and ContentNet-x is a transit link; the latter pays TransitNet for the generated traffic. The interconnection between TransitNet and AccessNet can be a SF-peering link, or a paid-peering arrangement in which one party pays the other. Today, such interconnections are typically SF-peering when AccessNet is a large network that does not require the transit services of TransitNet. When AccessNet is smaller, the interconnection is typically paid by AccessNet.

In this case, AccessNet and TransitNet can be thought of as "platforms" that offer some value to the end-to-end path. AccessNet is the only way to reach the subscribers, while TransitNet is the only way to reach the content providers. But this does not mean that AccessNet and TransitNet are equally important in this end-to-end path or that they encounter the same costs. Consequently, one of them can request a payment from the other, transforming their interconnection from SF-peering to paid-peering. What is an efficient way to split any economic surplus between the two platforms? How should AccessNet and TransitNet charge their customers (subsribers and content providers, respectively) and under which conditions should one of them charge the other? And how can we ensure that the two platforms have the incentive to keep investing in their infrastructures so that they avoid congestion?

The following highly stylized model does not capture congestion effects or investment incentives but it shows some major issues that emerge in this market structure (and that are not present in the two previous interconnection scenarios). The model captures the two-way interconnection between two platforms with captive customer bases (a similar framework to analyze the case of international call termination by two national monopolistic providers is presented in [72]). There are two platforms: A (AccessNet) with a continuum of identical customers of mass 1, and B (TransitNet) with a continuum of customers of mass 1 as well.

Suppose that the cost of originating a unit of traffic in platform i (to be terminated in the other platform) is c_i^O and the cost of terminating a unit of traffic in platform i (originated in the other platform) is c_i^T. Therefore, the total cost of delivering a unit of traffic from i to j is $c_i = c_i^O + c_j^T$. The price charged by platform i to its customers for delivering a unit of traffic to platform j is p_i. The customers' demand for traffic from i to j is $D_i(p_i)$. If traffic termination is priced at marginal cost, the retail profit function for platform i is $\pi_i(p_i) = (p_i - c_i)D_i(p_i)$.

Suppose the originating platform j pays the terminating platform i a paid-peering charge a_i per unit of traffic ($a_i = a_j = 0$ under SF-peering). Then the profits of platform i are

$$\Pi_i = (p_i - c_i^O - a_j)D_i(p_i) + (a_i - c_i^T)D_j(p_j)$$

Let us assume that the timing of moves is as follows: (1) the platforms choose their settlement charges a_i (through a bilateral agreement or by regulation), (2) taking the settlement charges as given, each platform non-cooperatively sets its retail price p_i to maximize its profit, taking the other's retail price as given.

We can start with the competitive retail benchmark assumption, i.e., given the termination charge a_j, the price of a unit of traffic originating in platform i is equal to its marginal cost: $p_i = c_i^O + a_j$. We can consider later a more realistic assumption that the AccessNet and TransitNet platforms have market power vis-a-vis their customers. Under that scenario, any ineffeciency in the determination of the settlement charges will be exacerbated due to the *double-marginalization problem.*

If both platforms choose their settlement rates according to $a_i = c_i^T$, then at the retail competition stage both of them would set the ideal prices $p_i = c_i$. Therefore, cost-based settlement charges induce the best outcome from the point of view of overall welfare.

However, the competitive pricing in each of the two platforms' retail markets does not mean that AccessNet and TransitNet would select their settlement rates in an efficient manner. Profits of platform i under the competitive retail market assumption are based solely on traffic termination: $\Pi_i = (a_i - c_i^T)D_j(p_j) = (a_i - c_i^T)\widehat{D}_j(a_i)$, where $\widehat{D}_j(a_i) = D_j(c_j^O + a_i)$. If the two platforms choose their settlement charge non-cooperatively, they will simply maximize their traffic termination profits with respect to their settlement charge $\frac{d(a_i - c_i^T)\widehat{D}_j(a_i)}{da_i} = 0$. This implies that $a_i > c_i^T$. Therefore, non-cooperative setting of paid-peering charges will cause them to be set at excessively high level, due to the double-marginalization problem.

Can the two networks negotiate a mutually beneficial contractual arrangement which sets their mutual settlement charges at the efficient level? To address this question, we consider a *bilateral bargaining framework* with the possibility of *incomplete contracting* between the networks.

If platforms have complete information about each others' costs and demands and side-payments can be made costlessly between them, the platforms can negotiate efficient termination charges equal to costs: $p_i = c_i$ and $a_i = c_i^T$. How the *first-best* surplus is divided between the platforms and whether such a division is sustainable depends on the details of the negotiation procedure and the contractual environment.

If no such side-payments are possible, the two platforms will simply bargain over the pair of settlement charges (assuming that prices in the retail market are then set competitively). Suppose for the sake of simplicity that platforms must choose symmetric paid-peering fees, so that $a_A = a_B = a$. Assuming pefectly competitive retail markets with a reciprocal peering fee a :

$$\Pi_i = \left[a_i - c_i^T\right] D_j(c_j^O + a_i)$$

and so platform i's ideal reciprocal settlement charge, denoted a_i^*, is given by the expression

$$a_i^* = c_i^T + \frac{D_j(c_j^O + a_i^*) - D_i(c_i^O + a_i^*)}{-D_j'(c_j^O + a_i^*)}.$$

In a rather rare situation when $c_A^T = c_B^T$, $c_A^O = c_B^O$ and $D_A \equiv D_B$ the interests of AccessNet and TransitNet coincide, and each is willing to agree to impose the settlement charge equal to marginal cost. However, in alternative and more likely scenarios, the platforms will have asymmetric interests. Therefore, given the asymmetry in preferences, we can assume that the equilibrium reciprocal settlement charge will lie somewhere in between the platforms' ideal levels, a_A^* and a_B^*, and its exact level will depend on the relative strength of platforms' "bargaining power" .

The expression for a_i^* shows that platform i would set a reciprocal settlement charge above its termination cost either when the platforms' traffic origination costs are similar and platform j has higher demand for traffic than platform i, or when demands are similar and platform j has lower traffic origination costs than platform i. Therefore, platforms with either higher costs or a net traffic inflow will prefer a higher reciprocal settlement charge. The practical implication of this is that network providers with a smaller customer base will prefer higher reciprocal settlement charges than providers with a larger customer base.

As in the earlier scenarios, we can also consider the case of congestion in the interconnection between AccessNet and TransitNet. We expect that congestion will be a systemic problem if the two parties have a broken contractual agreement and they cannot agree on who should pay whom or how much.

4.4 Access oligopoly with mesh-like Access-Transit-Content routes (Fig-2d):

This model is where "it all comes together." Here, we capture competition in both the Access market and the Transit market. We will consider the case of two access providers, AccessNet-1 and AccessNet-2. The number of transit providers is larger in practice. It is common today that content providers are multi-homed to several transit providers. Let us refer to transit providers as TransitNet-x. As it usually happens in practice, each transit provider is connected to all access providers, i.e., there is a direct interconnection between AccessNet-x and TransitNet-y for any x and y. The type of this direct interconnection can be transit, settlement-free peering or paid-peering, depending on (x,y). An important question in this interconnection scenario is how to to distribute any economic surplus between the population of access and transit providers.

Note that there is an important difference between the access market and the transit market: a subscriber is only connected to one access provider at any time. On the other hand, content providers are typically multi-homed to many or all transit providers. We will assume that ContentNet-x is connected simultaneously to *all* transit providers. This flexibility allows content providers to dynamically route their traffic towards each access provider through one or more transit providers. Such dynamic load balancing of the traffic can happen in short timescales (seconds). For instance, if the interconnection between AccessNet-x and TransitNet-y is congested, a ContentNet can re-route some (or all) of the traffic towards AccessNet-x through TransitNet-z, if that path is less congested. The fact that there are several routes from content providers to access providers, and the ability to re-route traffic dynamically, makes the analysis of this model much more complex.

The presence of competition in both the access and transit market means that the two end-points (subscribers and content providers) have more power in getting the connectivity and performance they want, compared to the previous interconnection scenarios. Given that competition is more intense in the transit market, however, and given that content providers can multi-home to several transit providers, makes the former even more powerful.

The impact of congestion is also more interesting in this model. For instance, consider the interconnection between AccessNet-1 and TransitNet-x. Suppose that AccessNet-1 requests a paid-peering interconnection with TransitNet-x, and the latter refuses to pay. If their interconnection becomes congested, Accessnet-1 risks to lose subscribers to AccessNet-2 but *at the same time* TransitNet-x risks to lose the traffic that goes from content providers to AccessNet-1.

5 Proposed work: Empirical work to parameterize models

In this section we describe the empirical research that we will perform in parallel with the modeling work described in the previous section. The main objective of this research will be to realistically parameterize the previous economic models and to validate their predictions. The data for this research will be provided by Comcast (see attached letter of commitment) as well as by various other sources described in the Data Management Plan. The Comcast dataset will contain anonymized information about the aggregate traffic that the corresponding customer uploads/downloads over successive 5-min periods, and it will cover several months in 2014 and 2015. This dataset will include some broadband subscribers as well as some transit enterprise customers. The broadband subscribers will be classified based on location and service plan.

Demand for different service plans: Most ISPs today offer several service plans in terms of upstream and downstream capacity. An important parameter is how the demand for different Internet access service plans varies with the price of those service plans. This is something we can estimate based on the Comcast dataset. Additionally, we will leverage a technique that has been recently developed by PI Dovrolis' group through which the service plan of a user can be passively inferred from (publicly available) MLab data, as long as that user has run MLab's NDT tool a few times. These additional data will allow us to examine the demand for different service plans in ISPs other than Comcast.

Traffic usage for different service plans: Another important parameter is how the demand for Internet content relates to the service plan of a user. Is it true that the more costly a service plan is, the more traffic the corresponding subscriber downloads? If so, the currently deployed multi-tier Internet service plans approximate, at least in some degree, usage-based service plans in which a user pays based on the amount of traffic she generates and consumes.

Sensitivity to congestion: The Comcast dataset will include a time period in the early part of 2014 in which the peering links between Comcast and Cogent were congested for several hours every day, according to data from CAIDA [79]. These congestion episodes will allow us to estimate how the demand for Internet traffic varies with the level of congestion that a subscriber experiences.

Traffic growth: Persistent traffic growth requires investments in Internet infrastructure. The economic modeling of such investments requires reliable estimates for the traffic growth rate. These estimates may be different in access versus core networks. We can estimate the traffic growth rate in access links based on the Comcast dataset. For the network core, CAIDA has access to longitudinal traffic data from SWITCH (we have used this data in an earlier publication [97]) and from Internet2. Additionally, several IXPs publish their aggregated traffic statistics online, e.g., LINX [2], DE-CIX [1].

Marginal costs: In prior work we have developed a cost model for network traffic [88], which accounts for the various cost components of the ISP. We will parameterize this cost model based on the Comcast dataset so that we can estimate the marginal cost of an Internet subscriber under different conditions related to the usage patterns of that user. This exercise will be repeated for Comcast transit customers (enterprise networks) which are billed differently (based on the well-known 95th-percentile scheme).

Interconnection density: The results of the fourth model of Section 4 depend on the interconnection density between access, transit and content providers. Instead of assuming that each access provider is connected to all transit providers, or that each content provider is connected to all transit providers, we will perform an empirical evaluation of this interconnection density for at least the largest access, content and transit providers. The interconnections between these ASes can be detected through well-established BGP data analysis methods [80].

Co-location density: Another relevant question is whether two ASes have the potential to interconnect: are they co-located in at least one public or private peering facility? We have been collecting and archiving snapshots of the peeringDB database since 2010, which has provided important insights into co-location, peering policies, peering requirements, and the evolution of these properties over time [78].

As the models of Section 4 show, the degree of competition in the access, transit or content markets plays a major role in the economics and pricing of the corresponding interconnections. We will try to estimate the degree of competition in each of these markets, at least in the US, as follows.

Competition in the access market: How many viable alternatives exist for Internet broadband access in the US and in other countries? How does this differ across states and regions in the same country? Answering these questions will involve mining data collected by the FCC [48], OECD [92], and the Berkman center [20].

Competition in the transit market: How many transit alternatives do access and content providers have in the US? We can answer this question based on data that is already being collected at CAIDA: AS-level topology annotated with business relationships and AS types [25], the size (in terms of the size of the customer base) of various transit providers, and geolocation information that maps networks to specific geographical regions in which they operate.

Competition in the content market: Most content is served from Content Distribution Networks (CDNs). The largest content providers (e.g., Google) operate their own CDNs. To estimate the degree of competition in the CDN market, we will use data from Alexa [9], which provides a list of the top web domains. By doing DNS lookups for those domains and mapping their IP addresses to ASes, we can estimate the number of domains from among the top-1M sites that are hosted by a given CDN.

6 Proposed work: Towards a new interconnection framework

A weakness of the current interconnection framework is that SF-peering relationships are based on various "peering conditions" (e.g., a balanced traffic ratio) that are often only indirectly related to the costs and benefits of that interconnection [91]. Such conditions are controversial because they often do not have a clear justification and they are not globally adopted, leading to frequent peering disputes. Another weakness of the current interconnection framework is that transit interconnections typically have the same price for all traffic, independent of the traffic's destination or network path. While more complex interconnections are occasionally established, such as "partial transit" or SF-peering only in certain locations [52], they are the exception rather than the norm and they are not well understood in terms of their economic implications.

The third part of our research agenda is to develop a new techno-economic framework for Internet interconnection that can address the previous weaknesses, and provide a broader and more economically efficient set of interdomain relations than just transit and SF-peering. Additionally, the proposed interconnection framework should provide the right incentives so that all relevant Internet firms (access, network and content providers, CDNs, IXPs, etc) continue to invest sufficient resources to Internet infrastructure. We refer to this interconnection framework as "techno-economic" because it will relate to both technological issues (e.g., which BGP routes are advertised to a neighboring AS) as well as economic issues (e.g., how to calculate the marginal cost of transferring a traffic flow over a given network path). The design of the new interconnection framework will be guided by the results and insights that the modeling and empirical components of this project will produce. We describe next the main ideas behind this framework.

A first feature of the proposed interconnection framework is that it will enable a broader class of relations between neighboring ASes, instead of just transit and peering relations. An AS X will be able to announce to another AS Y each route r together with a price $P_X(r)$; this is the price that X will charge to Y for sending traffic specifically to r. The price $P_X(r)$ may be zero for some local routes within X, a wide range of positive values for other routes, and even infinite for routes that X does not want to announce to Y. This interconnection framework includes the case of transit relations (if $P_X(r)$ is the same for all routes) and SF-peering relations (if $P_X(r)$ is zero for routes that originate at X or customers of X, and infinite otherwise) but it also enables a wider spectrum of AS relations. Additionally, the price of each route may also depend on the specific location l at which X and Y interconnect; for instance, a global provider may charge less for US-destined traffic in Canada than in Australia. In general, the price $P_X(r, l)$ will depend on several factors such as the actual cost encountered by X when transferring traffic from l to r, the network performance of that particular path (e.g., its capacity or propagation delay), whether that path traverses only X or other ASes that X will need to pay, or even the degree of competition for that particular route r at location l.

We believe that the flexibility that is provided by such a (route,location)-specific pricing scheme will provide the right incentives to all relevant Internet firms for the following reasons. First, transit providers will have the incentive to provision their network paths with sufficient capacity and low delay so that they can also charge more for the routes that traverse these paths. Arguably, content or application providers will be willing to pay more for a given route if they knew that that route is going to be of sufficient performance for the content or services they sell. At the same time, an access provider or enterprise network will have the incentive to be densely interconnected with other ASes so that it can receive multiple routes to each destination; it can then choose the route that best satisfies its performance and price objectives and constraints. This framework will also be beneficial for Internet users because it will result in better performance for network paths that carry high-value traffic and to lower prices for network paths that do not carry performance-sensitive traffic.

There is some prior work that is relevant to the proposed framework. MINT [112] proposes a market for interdomain routes, where a central *mediator* matches sellers with buyers based on an auction mechanism. MINT's focus is on the required routing protocols, and not on the underlying economic properties of such a scheme. Valancius [114] proposed that transit providers can structure their transit pricing into just 3-4 tiers based on traffic volume and cost of routing that traffic. The PIARA IETF working group [61] discussed financial incentives for providers to perform route aggregation and efficient assignment of address space but was not concerned with how to charge for traffic along those routes.

In the following, we summarize some questions about the proposed interconnection framework that we plan to address in the course of this research:

1) We need to understand in a rigorous manner the economic implications of this framework through analytical modeling and GENESIS simulations. How will it affect the welfare of both Internet users and the various firms that provide Internet content and infrastructure? And which are the implications of the proposed framework for Internet investment incentives?

2) How will this scheme, if deployed globally, affect the routing density as well as the performance of Internet paths? GENESIS is an ideal tool for answering such questions computationally [74].

3) While the proposed interconnection framework allows a provider to set a different price for each route, in practice it is likely that routes will be classified in a small number of pricing tiers. How many tiers would be sufficient in practice? This question was also investigated in [114] but only considering global-transit relations – not paid-peering relations.

4) In the proposed framework, each interconnection between two ASes is a fine-grained paid-peering relationship, where each route has a different price. We will develop algorithms that a provider can use to calculate the optimal price for each route. What cost, performance, and competition parameters would a provider need to estimate in order to compute such optimal prices?

5) There are also a number of practical issues we will need to investigate. For instance, how will a provider be able to announce pricing information along with each route? One approach will be to use BGP communities, where a route is *tagged* with a community value that also encodes a specific price. Another practical issue concerns traffic accounting and billing. With the proposed framework the traffic flowing through a link can have different prices depending on its origin and destination. This issue may be resolved through existing infrastructure, e.g., the use of SNMP and Netflow.

References

[1] DE-CIX Traffic Stats. `https://www.de-cix.net/about/statistics/`.

[2] LINX Traffic Stats. `https://www.linx.net/pubtools/trafficstats.html`.

[3] France Telecom severs all network links to competitor Cogent, 2005. `http://morse.colorado.edu/~epperson/courses/routing-protocols/handouts/cogent-ft.html`.

[4] YouTube sucks on French ISP Free, and French regulators want to know why, 2013. `http://gigaom.com/2013/01/02/youtube-sucks-on-french-isp-free-french-regulators-want-to-know-why/`.

[5] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *Proceedings of ACM SIGCOMM*, 2012.

[6] P. Aghion and P. Bolton. An incomplete contracts approach to financial contracting. *The review of economic Studies*, 59(3):473–494, 1992.

[7] P. Aghion, M. Dewatripont, and P. Rey. Renegotiation design with unverifiable information. *Econometrica: Journal of the Econometric Society*, pages 257–282, 1994.

[8] P. Aghion and J. Tirole. Formal and real authority in organizations. *Journal of political economy*, pages 1–29, 1997.

[9] Alexa. The top sites on the Web. `http://www.alexa.com/`.

[10] E. Anshelevich, B. Shepherd, and G. Wilfong. Strategic Network Formation Through Peering and Service Agreements. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, FOCS. IEEE Computer Society, 2006.

[11] E. Anshelevich and G. T. Wilfong. Network formation and routing by strategic agents using local contracts. In *WINE*, pages 386–393, 2008.

[12] E. Arcaute, R. Johari, and S. Mannor. Network formation: Bilateral contracting and myopic dynamics. In *Internet and Network Economics*, pages 191–207. Springer, 2007.

[13] E. Arcaute, R. Johari, and S. Mannor. Local two-stage myopic dynamics for network formation games. In *Internet and Network Economics*, pages 263–277. Springer, 2008.

[14] M. Armstrong. Network interconnection in telecommunications. *The Economic Journal*, 108(448):545–564, 1998.

[15] M. Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691, 2006.

[16] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2009.

[17] P. Baake and T. Wichmann. On the Economics of Internet Peering. *NETNOMICS*, 1(1):89–105, 1999.

[18] Backdoor Santa. Some Truth About Comcast - WikiLeaks Style, 2010. `http://www.merit.edu/mail.archives/nanog/msg15911.html`.

[19] N. Badasyan and S. Chakrabarti. A simple game-theoretic analysis of peering and transit contracting among internet service providers. *Telecommunications Policy*, 32(1):4–18, 2008.

[20] Y. Benkler. Next Generation Connectivity: A Review of Broadband Internet Transitions and Policy from Around the World. *The Berkman Center for Internet and Society Technical Report*, 2010.

[21] M. Bourreau, F. Kourandi, and T. M. Valletti. Net neutrality with competing internet platforms, 2014.

[22] M. Bourreau and M. Verdier. Cooperative and noncooperative r&d in two-sided markets. *Available at SSRN 2442166*, 2014.

[23] J. Brodkin. Why YouTube buffers: The secret deals that make-and break-online video, July 2013. `http://arstechnica.com/information-technology/2013/07/why-youtube-buffers-the-secret-deals-that-make-and-break-online-video/`.

[24] S. Buckley. France Telecom and Google entangled in peering fight, 2013. `http://www.fiercetelecom.com/story/france-telecom-and-google-entangled-peering-fight/2013-01-07`.

[25] CAIDA. AS-rank. `http://as-rank.caida.org`.

[26] B. Caillaud and B. Jullien. Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328, 2003.

[27] R. Casadesus-Masanell and G. Llanes. *Investment Incentives in Proprietary and Open-Source Two-Sided Platforms*. 2012.

[28] H. Chang, R. Govindan, S. Jamin, S. Shekker, and W. Willinger. Towards capturing representative AS-level Internet topologies. *Computer Networks*, 44(6):737–755, 2004.

[29] H. Chang and S. Jamin. To Peer or not to Peer: Modeling the Evolution of the Internet AS-level Topology. In *In IEEE INFOCOM*, 2006.

[30] H. Chang and W. Willinger. Difficulties Measuring the Internet's AS-Level Ecosystem. In *Proceedings of the 40th Annual Conference on Information Sciences and Systems*, 2006.

[31] Y.-K. Che and D. B. Hausch. Cooperative investments and the value of contracting. *American Economic Review*, pages 125–147, 1999.

[32] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes from P2P Users. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies (CoNEXT)*, pages 217–228, 2009.

[33] H. K. Cheng, S. Bandyopadhyay, and H. Guo. The debate on net neutrality: A policy perspective. *Information Systems Research*, 22(1):60–82, 2011.

[34] T.-Y. Chung. Incomplete contracts, specific investments, and risk sharing. *The Review of Economic Studies*, 58(5):1031–1042, 1991.

[35] D. Clark, S. Bauer, W. Lehr, K. Claffy, A. Dhamdhere, B. Huffaker, and M. Luckie. Measurement and Analysis of Internet Interconnection and Congestion. In *Proceedings of the Research Conference on Communication, Information and Internet Technology Policy*, 2014.

[36] S. Cowley. ISP spat blacks out Net connections, 2005. `http://www.infoworld.com/t/networking/isp-spat-blacks-out-net-connections-492`.

[37] W. Dai and S. Jordan. Modeling ISP tier design. In *Teletraffic Congress (ITC), 2013 25th International*, pages 1–9, 2013.

[38] A. Dhamdhere, H. Cherukuru, C. Dovrolis, and K. Claffy. Measuring the Evolution of Internet Peering Agreements. In *NETWORKING*, volume 7290 of *Lecture Notes in Computer Science*, pages 136–148. Springer Berlin Heidelberg, 2012.

[39] A. Dhamdhere and C. Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *Proceedings of the 6th International COnference*, Co-NEXT '10, pages 21:1–21:12, New York, NY, USA, 2010. ACM.

[40] A. Dhamdhere and C. Dovrolis. Twelve Years in the Evolution of the Internet Ecosystem. *IEEE/ACM Trans. Netw.*, 19(5):1420–1433, Oct. 2011.

[41] A. Dhamdhere, C. Dovrolis, and P. Francois. A Value-based Framework for Internet Peering Agreements. In *International Teletraffic Congress (ITC)*, Amsterdam, The Netherlands, Oct 2010.

[42] N. Economides and B. E. Hermalin. The economics of network neutrality. *The RAND Journal of Economics*, 43(4):602–629, 2012.

[43] N. Economides and J. Tåg. Network neutrality on the internet: A two-sided market analysis. *Information Economics and Policy*, 24(2):91–104, 2012.

[44] A. S. Edlin and S. Reichelstein. Holdups, standard breach remedies, and optimal investment. *American Economic Review*, 86:478–501, 1996.

[45] J. Engebretson. Level 3/Comcast dispute revives eyeball vs. content debate, Nov. 2010. `http://www.telecompetitor.com/level-3comcast-dispute-revives-eyeball-vs-content-debate/`.

[46] J. Engebretson. Behind the Level 3-Comcast peering settlement, July 2013. `http://www.telecompetitor.com/behind-the-level-3-comcast-peering-settlement/`.

[47] E. Farhi and A. Hagiu. *Strategic interactions in two-sided market oligopolies*. Harvard Business School, 2008.

[48] Federal Communications Commission. Eighth Broadband Progress Report. `http://www.fcc.gov/reports/eighth-broadband-progress-report`.

[49] V. Gehlen, A. Finamore, M. Mellia, and M. M. Munafò. Uncovering the Big Players of the Web. In *Proceedings of the 4th International Conference on Traffic Monitoring and Analysis*, TMA'12, pages 15–28, Berlin, Heidelberg, 2012. Springer-Verlag.

[50] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of the 9th International Conference on Passive and Active Network Measurement (PAM)*, 2008.

[51] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: A strategy for transitioning to bgp security. *SIGCOMM Comput. Commun. Rev.*, 41(4):14–25, Aug. 2011.

[52] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy. Inferring Complex AS Relationships. In *Internet Measurement Conference (IMC)*, Nov 2014.

[53] V. Giotsas, S. Zhou, M. Luckie, and K. Claffy. Inferring Multilateral Peering. In *Proceedings of ACM CoNEXT*, 2013.

[54] S. J. Grossman and O. D. Hart. The costs and benefits of ownership: A theory of vertical and lateral integration. *The Journal of Political Economy*, pages 691–719, 1986.

[55] S. Guriev. Incomplete contracts with cross-investments. *Contributions in Theoretical Economics*, 3(1), 2003.

[56] G. Gürsun and M. Crovella. On Traffic Matrix Completion in the Internet. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, New York, NY, USA, 2012. ACM.

[57] A. Hagiu. Pricing and commitment by two-sided platforms. *The RAND Journal of Economics*, 37(3):720–737, 2006.

[58] S. Hasan, S. Gorinsky, C. Dovrolis, and R. K. Sitaraman. Trade-offs in Optimizing the Cache Deployments of CDNs. In *Proceedings of IEEE INFOCOM*, 2014.

[59] L. He and J. Walrand. Pricing and revenue sharing strategies for internet service providers. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 1, pages 205–216. IEEE, 2005.

[60] P. Holme, J. Karlin, and S. Forrest. An Integrated Model of Traffic, Geography and Economy in the Internet. *SIGCOMM Comput. Commun. Rev.*, 38(3):5–16, July 2008.

[61] IETF PIARA Working Group. Incentives for Route Aggregation and Efficient Address Utilization in the Internet. `https://datatracker.ietf.org/doc/draft-bellovin-piara-framework/`.

[62] R. Johari, S. Mannor, and J. N. Tsitsiklis. A contract-based model for directed network formation. *Games and Economic Behavior*, 56(2):201–224, 2006.

[63] R. Johari and J. N. Tsitsiklis. Routing and peering in a competitive internet. *LIDS Publication*, 2570, 2003.

[64] Jon Brodkin. Time Warner, net neutrality foes cry foul over Netflix Super HD demands, 2013. `http://arstechnica.com/business/2013/01/timewarner-net-neutrality-foes-cry-foul-netflix-requirements-for-super-hd/`.

[65] B. Klein, R. G. Crawford, and A. A. Alchian. Vertical integration, appropriable rents, and the competitive contracting process. *Journal of law and economics*, pages 297–326, 1978.

[66] J. Krämer and L. Wiewiorra. Network neutrality and congestion sensitive content providers: Implications for service innovation, broadband investment and regulation, 2009.

[67] J. Krämer, L. Wiewiorra, and C. Weinhardt. Net neutrality: A progress report. *Telecommunications Policy*, 37(9):794–813, 2013.

[68] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-domain Traffic. In *Proceedings of the ACM SIGCOMM*, New York, NY, USA, 2010. ACM.

[69] J.-J. Laffont, S. Marcus, P. Rey, and J. Tirole. Internet interconnection and the off-net-cost pricing principle. *RAND Journal of Economics*, pages 370–390, 2003.

[70] J.-J. Laffont, P. Rey, and J. Tirole. Network competition: I. overview and nondiscriminatory pricing. *The RAND Journal of Economics*, pages 1–37, 1998.

[71] J.-J. Laffont, P. Rey, and J. Tirole. Network competition: Ii. price discrimination. *The RAND Journal of Economics*, pages 38–56, 1998.

[72] J.-J. Laffont and J. Tirole. *Competition in telecommunications.* MIT press, 2001.

[73] S. Lippert and G. Spagnolo. Internet Peering as a Network of Relations. Discussion Paper Series of SFB/TR 15 Governance and the Efficiency of Economic Systems 191, Free University of Berlin, Humboldt University of Berlin, University of Bonn, University of Mannheim, University of Munich, Nov. 2006.

[74] A. Lodhi, A. Dhamdhere, and C. Dovrolis. GENESIS: An Agent-based Model of Interdomain Network Formation, Traffic Flow and Economics. In *INFOCOM, Proceedings IEEE*, 2012.

[75] A. Lodhi, A. Dhamdhere, and C. Dovrolis. GENESIS-CBA: An Agent-based Model of Peer Evaluation and Selection in the Internet Interdomain Network. *Complex Adaptive Systems Modeling*, 2013.

[76] A. Lodhi, A. Dhamdhere, and C. Dovrolis. Open peering by internet transit providers: Peer preference or peer pressure? In *INFOCOM, 2014 Proceedings IEEE*, 2014.

[77] A. Lodhi and C. Dovrolis. A Network Formation Model for Internet Transit Relations. In *Proceedings of the 2010 Workshop on Economics of Networks, Systems, and Computation*, New York, NY, USA, 2010. ACM.

[78] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and k. claffy. Using peeringdb to understand the peering ecosystem. *SIGCOMM Comput. Commun. Rev.*, 44(2):20–27, Apr. 2014.

[79] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, and K. Claffy. Challenges in Measuring Internet Interdomain Congestion. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2014.

[80] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k. claffy. AS Relationships, Customer Cones, and Validation. In *ACM SIGCOMM Internet Measurement Conference (IMC)*, Oct 2013.

[81] R. Lychev, S. Goldberg, and M. Schapira. Bgp security in partial deployment: Is the juice worth the squeeze? *SIGCOMM Comput. Commun. Rev.*, 43(4):171–182, Aug. 2013.

[82] R. T. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. In *Proceedings of 2007 ACM Conference on Emerging Networking Experiments and Technology*, Columbia University, New York, December 2007.

[83] R. T. Ma, J. C. Lui, and V. Misra. On the Evolution of the Internet Economic Ecosystem. In *Proceedings of the International Conference on World Wide Web*, Republic and Canton of Geneva, Switzerland, 2013. International World Wide Web Conferences Steering Committee.

[84] R. T. B. Ma, D. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers. In *Proceedings of the ACM CoNEXT Conference*, New York, NY, USA, 2008. ACM.

[85] Measurement Lab Consortium. ISP Interconnection and its Impact on Consumer Internet Performance, 2014. `http://www.measurementlab.net/static/observatory/M-Lab_Interconnection_Study_US.pdf`.

[86] E. A. Meirom, S. Mannor, and A. Orda. Network Formation Games and the Internet Structure. *ArXiv e-prints*, July 2013.

[87] J. Mikians, A. Dhamdhere, C. Dovrolis, P. Barlet-Ros, and J. Sole-Pareta. Towards a Statistical Characterization of the Interdomain Traffic Matrix. In *IFIP Networking 2012*, pages 111–123, May 2012.

[88] M. Motiwala, A. Dhamdhere, N. Feamster, and A. Lakhina. Towards a cost model for network traffic. *SIGCOMM Comput. Commun. Rev.*, 42(1):54–60, Jan. 2012.

[89] P. Njoroge, A. Ozdaglar, N. E. Stier-Moses, and G. Y. Weintraub. Investment in two-sided markets and the net neutrality debate. *Review of Network Economics*, 12(4):355–402, 2010.

[90] G. Nöldeke and K. M. Schmidt. Option contracts and renegotiation: a solution to the hold-up problem. *The RAND Journal of Economics*, pages 163–179, 1995.

[91] W. B. Norton. *The Internet Peering Playbook: Connecting to the Core of the Internet.* DrPeering Press, 2011.

[92] OECD. OECD Broadband Portal. `http://www.oecd.org/sti/broadband/oecdbroadbandportal.htm`.

[93] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure. In *Proceedings of ACM SIGMETRICS*, 2008.

[94] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (In)completeness of the Observed Internet AS-level Structure. *IEEE/ACM Transactions on Networking*, 18(1):109–122, Feb. 2010.

[95] R. V. Oliveira, B. Zhang, and L. Zhang. Observing the Evolution of Internet AS Topology. In *Proceedings of ACM SIGCOMM*, 2007.

[96] J. Pil Choi and B.-C. Kim. Net neutrality and investment incentives. *The RAND Journal of Economics*, 41(3):446–471, 2010.

[97] V. Reddyvari Raja, A. Dhamdhere, A. Scicchitano, S. Shakkottai, k. claffy, and S. Leinen. Volume-based transit pricing: Is 95 the right percentile? In *Passive and Active Measurement*, 2014.

[98] C. Reggiani and T. Valletti. Net neutrality and innovation at the core and at the edge. *The School of Economics Discussion Paper Series*, 1202, 2011.

[99] M. Ricknas. Sprint-Cogent Dispute Puts Small Rip in Fabric of Internet, 2008. `http://www.pcworld.com/article/153123/sprint_cogent_dispute.html`.

[100] J.-C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

[101] J.-C. Rochet and J. Tirole. Two-sided markets: A progress report. *The RAND Journal of Economics*, 37(3):645–667, 2006.

[102] M. Roughan, S. J. Tuke, and O. Maennel. Bigfoot, Sasquatch, the Yeti and Other Missing Links: What We Don't Know About the AS Graph. In *Proceedings of ACM SIGCOMM IMC*, pages 325–330, 2008.

[103] Sandvine. Global Internet Phenomenon Report, 1H13, 2013.

[104] F. Schuett. Network neutrality: A survey of the economic literature. *Review of Network Economics*, 9(2), 2010.

[105] Sean Buckley. Cogent and Orange France fight over interconnection issues, 2011. `http://www.fiercetelecom.com/story/cogent-and-orange-france-fight-over-interconnection-issues/2011-08-31`.

[106] S. Shakkottai, R. Srikant, A. Ozdaglar, and D. Acemoglu. The Price of Simplicity. *IEEE J.Sel. A. Commun.*, 26(7):1269–1276, Sept. 2008.

[107] G. Shrimali and S. Kumar. Can Bill-and-Keep Peering Be Mutually Beneficial? In *Internet and Network Economics*, pages 738–747. Springer, 2005.

[108] G. Shrimali and S. Kumar. Paid Peering Among Internet Service Providers. In *Proceeding from the 2006 Workshop on Game Theory for Communications and Networks*, GameNets '06, New York, NY, USA, 2006. ACM.

[109] G. Shrimali and S. Kumar. Bill-and-keep peering. *Telecommunications Policy*, 32(1):19–32, 2008.

[110] G. Siganos, M. Faloutsos, and C. Faloutsos. The Evolution of the Internet: Topology and Routing. *University of California, Riverside technical report*, 2002.

[111] R. Stanojevic, N. Laoutaris, and P. Rodriguez. On economic heavy hitters: Shapley value analysis of 95th-percentile pricing. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 75–80. ACM, 2010.

[112] V. Valancius, N. Feamster, R. Johari, and V. Vazirani. Mint: A market for internet transit. In *Proceedings of the 2008 ACM CoNEXT Conference*, CoNEXT '08, pages 70:1–70:6, New York, NY, USA, 2008. ACM.

[113] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani. How Many Tiers?: Pricing in the Internet Transit Market. In *Proceedings of the ACM SIGCOMM 2011 Conference*, SIGCOMM '11, pages 194–205, New York, NY, USA, 2011. ACM.

[114] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani. How many tiers?: Pricing in the internet transit market. In *Proceedings of the ACM SIGCOMM 2011 Conference*, SIGCOMM '11, pages 194–205, New York, NY, USA, 2011. ACM.

[115] Verizon. Unbalanced peering, and the real story behind the verizon/cogent dispute, June 2013. `http://publicpolicy.verizon.com/blog/entry/unbalanced-peering-and-the-real-story-behind-the-verizon-cogent-dispute`.

[116] D. Weller and B. Woodcock. Internet Traffic Exchange: Market Developments and Policy Challenges. *OECD Digital Economy Papers*, 2006.

[117] E. G. Weyl. A price theory of multi-sided platforms. *The American Economic Review*, pages 1642–1672, 2010.

[118] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level Topology. *SIGCOMM Computer Communication Review*, 35(1):53–61, Jan 2005.