Workshop on Internet Economics (WIE2011) Report

kc claffy
CAIDA/UC, San Diego
La Jolla, CA, USA
kc@caida.org

ABSTRACT

The second Workshop on Internet Economics [2], hosted by CAIDA and Georgia Institute of Technology on December 1-2, 2011, brought together network technology and policy researchers with providers of commercial Internet facilities and services (network operators) to further explore the common objective of framing an agenda for the emerging but empirically stunted field of Internet infrastructure economics. This report describes the workshop discussions and presents relevant open research questions identified by its participants.

Categories and Subject Descriptors

C.2.3 [**Network Operations**]: Network Management Public Networks; C.2.5 [**Local and Wide-Area Networks**]: Internet; J.4 [**Social and Behavioral Sciences**]: Economics; K.4.1 [**Public Policy Issues**]: Transborder data flow

General Terms

Economics, Legal Aspects, Management

Keywords

Economics, Internet, Network management

1. BACKGROUND

Building on the success of our first (virtual) Workshop on Internet Economics (WIE09) [5], we expanded the scope and depth of the second workshop in this series, inviting experts in the following topics of interest: peering strategies and conflicts; content delivery; traffic and topology dynamics of the peering ecosystem; sustainable business models and industry structure. The workshop format was structured to promote constructive, focused engagement. Attendees presented research results, offered updates on data sources, moderated topic discussions, served as formal respondents to other speakers, and critiqued the relevance and potential impact of the presented results An intended goal was to establish a set of open questions that can frame an Internet economics research agenda, and more specifically to improve the realism, utility, and predictive power of economic models of Internet topology and dynamics. Even this expanded scope falls short of the recognized breadth of the emerging discipline of Internet economics; we plan to include other recurring issues and questions at the next workshop, including the economics of privacy, advertising, censorship, and intellectual property.

2. PEERING STRATEGIES AND ISSUES

David Clark opened the session by presenting some analysis of the economic policy challenges of interconnection in the Internet [6], providing relevant background for several sessions of the workshop. The Internet industry has a long tradition of using traffic ratios, i.e., balance of flows between two peers, as a key metric in peering decisions [4]. In the first decade of the evolution of the commercial Internet, beginning in the early-90s, networks typically peered with each other without financial exchange, a departure from the settlement model familiar to the previous telecom paradigm. Within a few years, many networks became selective about peering partners, preferring to peer with networks similar in size or topological importance, and using the ratio of traffic flows between two peering networks as a proxy measure of this similarity.

This traffic ratio metric embeds a never-validated assumption that these metrics are correlated to the cost of delivering the traffic. Some operators in the room believed that although it may have never been formally validated, the traffic ratio rule does have a cost-related basis when coupled with the de facto industry standard *hot potato routing* policy. Hot-potato routing means that an ISP by default will choose the shortest internal path to an adjacent next-hop network, minimizing the number of hops a packet travels on its own network, and increasing the cost for the receiving network. Using this reasoning, if both sides use hot-potato routing, the cost for each is proportional to the traffic received, and the balance of flow required by the traffic ratio rule implies balance of cost.

Regardless of its applicability to wide-area network peering, the above model does not cover the case of CDNs connecting to access networks. CDNs can source their content from multiple locations, and they normally choose a source close to the destination to reduce latency. In doing so, they tend to minimize the distance the traffic travels over the receiving access network, which is the opposite of what happens with hot potato routing. Although some CDNs do not have much wide-area network infrastructure anyway, leaving access networks with the burden of transmitting all of the CDN's content to the end user.

Furthermore, even in cases where this traffic-ratio model has traditionally made some sense, divergence of traffic ratios [2] inevitably led to disputes about the role of settlements, or *paid peering*. According to many participants, both relationships were just forms of customer-provider relationships, as opposed to a relationship among peers or partners. But participants generally agreed that settlement-free peering could no longer be sustained when one peer doubles or triples its traffic sent to the other peer. Operators also confirmed that virtually all negotiated peering agreements with any recourse (SLAs, notice of cancellation) involve financial exchange; such terms do not exist for settlement-free peering agreements by default.

Ren Provo of Comcast pointed out that there are alternatives to charging for peering, such as the use of MEDs (Multi-Exit Discriminators) to mitigate traffic ratio imbalances, to keep traffic on one's own network longer. (MEDs allow an AS to notify a neighbor AS of its preference as to which of several links are preferred for inbound traffic.) Susan Martens of AT&T confirmed that they have gained considerable flexibility with the use of MEDs too, to keep the 'bit miles' traveled on a network under a cap for some set of traffic exchanged with a peer. In other words, as ratios get out of balance, they require that the other network shift from hot potato routing to a cold-potato routing in which they keep the traffic internally as long as possible. She noted that there is a limit to the capability of MEDs to mitigate peering imbalances – although MEDs work reasonably well for moderate traffic imbalances, not even MEDs can re-balance network costs proportionally with extreme traffic imbalances. Level 3 (not represented at this workshop, although presented on this topic at WIE09 [1]) recently proposed an alternative method to equitably share the cost of moving the traffic across the backbone links [8], but the proposal ignores the substantially larger last-mile infrastructure cost component that Level 3 itself does not typically have to support since it is not an access provider.

In a related discussion, Maurice Dean from Google explained how the current peering fabric in the U.S., characterized by relatively few large peering points to optimize economic constraints of those investing in and maintaining infrastructure, imposes some performance cost on end users, at least in terms of variation in latency, perhaps also in terms of consumer choice. However, an open research challenge is to empirically validate and quantify this hypothesis. Most large content providers gather and analyze a tremendous amount of performance data – latency, jitter, packet loss, and/or retransmission rates – often categorized, at least roughly, across ISPs or geographic regions, and sharing this data with researchers could support investigation of the relationship between peering behavior (geography, interconnection bandwidth) and suboptimality of the user experience.

Continuing the theme of empirically observable effects of industry structure, Bill Norton provided a colorful historical perspective on the evolution of Internet peering. Contrasting diagrams reflecting the peering ecosystem ten years ago and today, he observed an astonishing inversion in positional power over the last decade. In parallel with transit prices dropping another order of magnitude (having dropped an order of magnitude in the 1990s), cable companies and other access providers who used to pay *Tier1 backbones* for transit are now commanding payments from content providers and CDNs who need to reach eyeballs, i.e., customers of the access providers. As this Tier1 backbone component of the ecosystem loses traffic to direct interconnection between content networks (CDNs) and access networks, their pricing and profitability continue to drop, even if revenues increase due to higher overall traffic volumes. Not suprisingly, every year fewer backbones operate as their own business.

The costs of moving traffic are influenced by, and in turn influence, network business relationships, internetwork topology, routing policies, and the resulting interdomain traffic flow. The morning conversation repeatedly revealed a wide range of opinions on cost models for peering, effectively illustrating the gaps in our understanding of these interactions, and perfectly motivating the following discussions on scientific approaches to modeling industry peering behavior.

2.1 Modeling of peering strategies and incentives

We followed a lively morning session with an even livelier afternoon, inspired (and provoked) by two presentations from academic researchers on approaches to computational modeling of the Internet peering ecosystem. Aemen Lodhi (with advisor Constantine Dovrolis) of Georgia Tech presented an agent-based model of peering, specifically a dynamic and decentralized ecosystem of independent agents that try to optimize a cost-related fitness function based on limited information. The model allows exploration of the presence or lack of equilibria, or profit, in the face of various network peering and infrastructure expansion strategies, traffic flows, geographic constraints, and costs. Network operators in the audience gave pointed feedback on how to make the model more realistic, including how transit pricing tends to work between providers, e.g., most providers bill on the 95th percentile traffic volumes, the greater of the in and out directions. Importantly, peering decisions are influenced by many other factors that are not captured by this model, and (some operators thought) may not be amenable to modeling at all. The fitness metric used is far more narrowly defined than in the real world, incorporating no notion of the significant investments that may occur to make peering happen, or different behavior in different locations based purely on market-related constraints of a given geographic region.

Amogh Dhamdhere of CAIDA/UCSD presented two related recent results [9, 3]: formalizing a cost model for network traffic, and applying this cost model to a value-based framework for peering agreements. Several operators challenged Amogh that value could not be determined without some external entity and cost model, but Amogh clarified that each peer assigns their own value to the potential peering session, and it will only be created if each side considers themselves better off (higher value) in doing so. For example, if the value of the peering session is dwarfed by the value of protecting one's home market from any new entrants, the peering is unlikely to happen. The inversion in positional power that Bill Norton earlier described in his talk provides some empirical validation of this model. Specifically, in the 1990s, the access providers tended to pay upstream Tier1 backbones for connectivity. The economics have shifted over the last two decades, and now many content providers and CDNs can afford to build and maintain their own backbones, due to their higher revenue margins and the plummeting prices of long-distance transit. Cable and access providers in turn have developed business models for monetizing their customer eyeballs, which content providers and CDNs try to reach as efficiently as possible. These economic forces have reshaped the value proposition for peering in favor of the access providers, many of whom now find more value in peering directly with major content providers.

Patrick Gilmore repeated his warning, and presented slides to underscore his point, that the models were not capturing most of the dynamics in the reality of peering. A philosophical debate about goals and expectations of scientific research followed. Do we expect research to exactly match the real world, or should it allow us to investigate how things could or should be? Constantine confirmed from experience that talking to ten peering coordinators might give you some understanding of how things are today, but it promises no understanding of how things *should* be. Modeling in particular is not about achieving a complete understanding of the phenomena being modeled – on the contrary, sometimes the greatest benefit to modeling is to reveal explicitly where control or understanding is lacking.

Connecting this discussion to the earlier morning threads, John Chuang from UC Berkeley found it interesting that despite the recognized performance issues and critical nature of the infrastructure, industry players are still not interested in negotiating agreements with teeth (SLAs) for agreements with zero payment. John reminded us, "Zero is just a number", and its special treatment by the industry implies that they have not addressed the issue of what constitutes respective value, a gap that represents a daunting challenge for academic research attempting to formalize it.

Nikolaos Laoutaris from Telefonica contributed related thoughts on the missing link between academic and operational network economics, having now spent time in both fields. Theory and practice are especially disconnected in network economics, compared to other areas, since operators do not use theory or models; they just use rules of thumb such as the traffic-ratio rule discussed earlier. Operators do not even know how close or far they are from optimal efficiency in terms of performance, cost, or (especially) security. Many open debates in economics lack empirical quantification – not just from a lack of data but a lack of capability to extract knowledge from data that does exist. Nikolaos repeated a theme: the most needed contribution is in solid quantification of what we know. For example, in his experience traffic volumes are growing at 40% per year, creating an increasing decoupling between costs and revenues, and inspiring the network neutrality debate, or more specifically from his perspective, the need for providers to find a way to "discriminate positively, not negatively".

As another example, Nikolaos pointed out that an amazing amount of cost data is public (at least in Europe) because of regulatory requirements to publish local loop unbundling costs. Many years ago, the FCC went through its own cost benchmarking exercise both for domestic and ITU settlements. But a cost-plus compensation environment creates an incentive to gold plate investment patterns in a way that will not match those of carriers who do not have to reveal those numbers. So while such data provides a starting point (lower/upper bound), modeling efforts must take into account the fact that unbundling and associated data disclosure requirements actually do change the investments being made, since they change the risk/reward profile.

2.2 New pricing models for transit

This conversation served as a natural transition to talks on transit models. Sergey Gorinsky from Institute IMDEA Networks advocated CIPT (Cooperative IP Transit), a model for cooperative buying to reduce transit costs [10]. A cooperative buying club would jointly purchase large-capacity transit to benefit from subadditive pricing. Based on real traffic of numerous European operators and real pricing function, Sergey quantified significant aggregate and individual cost reductions when CIPT used Shapley values to share transit costs among members. In the subsequent discussion, a few network operators raised concerns about the overhead of maintaining CIPT, e.g., with respect to cost sharing and service/billing problem resolution. Also, transit prices for such cooperative buying models might need to be adjusted to be realistic. Nikolaos Laoutaris pointed out that large colocation providers employ CIPT-like practices, at least in Asia. The R&E community in the U.S. (thequilt.net) is another notable exception, which has had a successful cooperative buying club for years. Richard Steenbergen suggested that a more common and realistic model is the gearless reselling model, where the customer still has a direct relationship with the provider but uses an intermediary as the billing agent who takes responsibility for being the customer of record for credit purposes.

Vytautas Valancius presented results from a study of tiered pricing in the Internet transit market, where he found that a small number of tiers (two) would capture most of the benefit of tiered pricing, assuming one could get accounting software implementations fixed to support it. The proposal triggered strong objections from several providers who have tried but failed to get industry traffic flow accounting software to provide any reasonable software support for tiered qualities of services. Valas cheerfully promised us that if the tiering promised a sufficiently significant improvement in efficiency, he would fix the software himself! Regardless of software complexity, their study concluded that more than two tiers yields diminishing returns. Richard Steenbergen also pointed out that in the real world tiering is more often in terms of on-net vs. off-net or even in terms of desirable vs. undesirable peers.

3. IPV6 DEPLOYMENT

As a timely interlude characterized by similar themes, Geoff Huston led a discussion on IPv6 on the second day of the workshop, presenting results from extensive analysis of edge support for IPv6. Although almost all host operating systems (Windows, OSX, Linux) and mobile phones support IPv6, and almost half of IPv4 transit networks also announce an IPv6 address prefix, the reality remains that only a small fraction of a percent (0.1-0.3%) of clients accessing Google's services are doing so using IPv6. It can no longer be blamed on ignorance of the need. Industry is well aware of IPv6, but disinterested, for some obvious reasons: IPv6 offers no comparative advantage against competitors; it's not cheaper; it offers no new functionality, there is no demand for it; and there is no observable difference for most consumers. The only benefit of IPv6 deployment is to help manage future risk. But people are inherently poor at valuing future risk and certain organizations, including access monopolies, have other ways of controlling risk since they control the (customer) eyeballs. Geoff mentioned the example of earthquakes, where risk management does include government regulation of building codes, presumably since builders are also unlikely or unable to value or manage risk appropriately.

Geoff concluded that IPv6 is a public good, and that there is no way market forces will bring about the transition on their own. He described several conventional approaches to the distribution of a public good, all involving some type of government intervention: (1) regulatory mandate; (2) government purchase contracts, such as those currently being undertaken in the U.S. and other countries; (3) subsidies and incentives (his personal favorite), either giving consumers or providers money to consume or provide IPv6 services, and accelerate schedules for IPv4 equipment amortization; (4) public provision of IPv6 connectivity (funded by bonds or taxes) which is how telephone service was provided for a big chunk of the twentieth century so is hardly a foreign concept for the industry.

Geoff warned, as many others have, that the alternative is broad-based carrier-grade NAT (CGN), with which we have no experience, no standards, and no idea how it will actually work. But without external intervention, he believed that economic forces would push the industry in the direction of CGNs, for another (related) obvious reason: content is currently operating at much higher margins and profitability than the carriage industry. In Geoff's view, this differential in profit margins implied that the latter must look at IP address exhaustion and CGN as an opportunity to renegotiate terms and conditions with "rich" content providers, e.g., Google, Netflix. In this view, IPv6 is just a collateral victim, quite beside the point.

In a related thread, Patrick Gilmore asked whether there are other examples of related but distinct industries that are allowed to extract rent from each other because one is not so profitable as the other, e.g., Verizon vs. Google, Exxon vs. GM, Comcast vs. Akamai.

4. INDUSTRY STRUCTURE

The rate of change in the market actors (the rise of content, the rise of the access network, the decline of the transit provider) means that the market lacks a stable and well-understood structure. These dynamics (and lack of clear structure) will continue for some time. Unfortunately, we lack a rigorous test of whether there is *abuse* of market power, which is typically considered a function of two things: competition and customer lock-in. Geoff questioned whether competition should even be a public policy goal for an industry with massive economies of scale rooted in extraordinarily high capital costs that inhibit entrance. Using similar reasoning, regulators in countries including Japan and the U.K. separated the pipe from the services sold on top (unbundling), resulting in a thriving, competitive retail ISP market atop a regulated monopoly pipe [11]. In the U.S. the opposite regulatory direction has resulted in what some call *facilities-based competition*, typically between cable and DSL/FTTX technologies, and characterized by higher retail costs.

John Chuang introduced us to the concept of a two-sided bilateral oligopoly, i.e., where a set of access networks, each representing a terminating access monopoly with respect to their customers, are connected to a set of powerful content providers, each with monopoly control over their respective sets of content [7]. He proposed this model as an improvement to the recently popular two-sided market model. The problem with the two-sided market model in describing the current Internet is that since (today) both the access network and the CDN sit between the consumer and the content, some sort of vertical integration must exist between access and CDN providers so as to establish a two-sided platform between the two ends. Such vertical integration is still the exception today rather than the norm, so a model that can accommodate industry structure with or without such vertical integration is likely to be more informative. Important characteristics of a bilateral oligopoly include: (1) negotiation and bargaining leading to long-term contracts that create entry barriers for both fields of business; (2) an appropriate balance of power across the two sides. These characteristics can result in lower consumer prices and increased consumer surplus. The proposition was thought-provoking: can the presence of a bilateral oligopoly help to counteract the market power of terminating access monopolies? He offered two other example industries where such a dynamic occurred: the hospital market and the health insurer markets, and the music distribution and record label markets.

Dave Clark questioned whether the CDN market was really one of the two sides of the oligopoly in the Internet market, since both access providers and content owners seem to have even more power than the CDNs. In Europe many access providers also maintain their own caching infrastructure to support efficient content distribution, which reduces the power of the CDNs. On the content side it is even more stark. Google does not even use a CDN; it builds its own and bargains with access networks. In contrast, Netflix uses a CDN intermediary, indeed to play them off against each other, so that the CDN itself loses market power and instead gets squeezed in the middle. Like transit providers, they are structurally disadvantaged in the evolving ecosystem.

While content owners are only starting to explore how to value their content on the Internet (and we never got to the issue that content comes with all sorts of value), there are cascades of power delegation in this ecosystem. Netflix is not even a content owner (although trying to change that), but just a tenuous holder of some licensing rights. The real owners are licensing through multiple channels, so even Netflix is limited in the bargains they can negotiate. Content and license owners, not to mention advertisers, would likely consider our interest in peering cost models to be rather beside the point – they just want to reach the customer in as many (revenue-producing) ways as possible, and what would interest them is the relative ability to extract payments from Netflix, Apple, Google, etc.

Several other talks covered dimensions of industry structure in the face of technology transition. Dah Ming Chiu from the Chinese University of Hong Kong described how the routing scalability issue is worse if there is a lot of demand for multihoming, but various providers argued there is little demand for multi-homing right now. Heikki Hammainen from Aalto University in Finland proposed techniques for congestion management in competitive mobile access markets, including per-session splitting across providers, which intrigued both researchers and operators but also made them both wonder how disputing a bill would work. Richard Ma of the National University of Singapore proposed a "public option" Internet service provider to address network neutrality concerns, i.e., to provide competition where the market cannot do so. A network-neutral public option could free commercial providers from such neutrality obligations, allowing them to experiment with innovations in service differentiation. Although municipal networking has received strong legal resistance from commercial providers in the U.S. (who do not want the competition and have aggressively and successfully sued and lobbied to prevent or stop municipal networking projects), the U.S. FCC has endorsed the public option as a method of bringing broadband to underserved communities.

5. DISTILLING A RESEARCH AGENDA

Below, we collate a set of research questions and goals that can serve as a basis for a future research agenda, regardless of our awareness of how to approach them at this time.

5.1 Modeling market dynamics

1. How can we build models to capture, explain, and predict macroscopic shifts in market power dynamics, e.g., transitions from settlement-free to paid peering, and the rise and fall of "hegemonic Internet powers", based on publicly available data such as ex parte FCC filings and information on who has arrived and departed from major exchange points in the last ten years?
2. How can we better model the notion of value derived from a peering session?
3. Can we develop more sophisticated models and metrics, such as Shapley-value approximations to distribute costs among different entities?
4. Can we determine realistic demand elasticity for access to the Internet, and how it differs across customer classes?

5.2 Consumer welfare

1. Can we use existing data collected by content providers on the quality of the user experience (packet rates, loss, retransmits) to study the impact of competition of last mile on user welfare? Are concerns about competition accurate, or is the market working?
2. What is consumer willingness to pay for various services? How much do people value different applications: search vs. online social networking vs. email vs. gaming?
3. Can we characterize customer churn? How loyal are users to a given provider?
4. Is there a connection between market power and churn?
5. Can we quantify the factors that influence coalition formation to purchase transit cooperatively?
6. Is content churn a relevant issue on any level of Internet economics, e.g., one movie vs. another on the content level, among marketers of the same movie on the portal level, Akamai vs. Level 3 on the channel level?

5.3 Industry structure and cost models

1. Will the Internet industry undergo a disruptive structural reorganization or an evolutionary adaptation with the rapid rise of large-scale video distribution? What economic models and empirical data are most suited for studying the shifts in market power, the charging of services, the recovery of costs, and the implications to social welfare?
2. How do the economics of advertising and content production relate to the economics of bit transport? How do we quantify the economic efficiencies that results from various proposed structures of these increasingly coupled, complex industries, e.g., what is the welfare implication of content providers subsidizing bandwidth metering fees?
3. How do we study what would happen if we only allowed ISPs to charge one side of their two-sided market: consumers or peers?
4. Why are there fewer IXPs in the US than Europe? How many IXPs are sustainable in a given region, and IXP consolidation be observed, e.g., via PeeringDB? Do networks present in multi-IXP cities choose one or the other IXP, or both?
5. Is there a strong motivation or demand for multi-homing? What is a good pricing model for multi-homed connections, especially in residential settings? How can we comprehensively treat the reliability, competition, fixed vs. mobile, and other diverse aspects of multi-homing?
6. What is the impact of mobile access evolution on interconnection?
7. How would per-session wireless (mobile/cellular) competition work (selecting one's carrier on a per-call basis)?

5.4 Economic policy

1. How might a regulator (or anyone else) assess (measure, determine, etc.) the evidence of significant market power in the interconnection market? What would be evidence of the abuse of that market power?
2. Is it possible to create an economic-theoretic model to help decide whether Level 3 should go to the FCC over a peering dispute? For example, should peering disputes even be adjudicated by the FCC, or is it a role for the courts?
3. How do different regulatory environments encourage or discourage a shift in market power from content to eyeballs?
4. How do other examples of liquidity crises in history impact IPv4 exhaustion and the IPv6 transition?

6. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, CAIDA, UC, San Diego
- Co-Host: Constantine Dovrolis, Georgia Tech
- Presenter: Dave Clark, MIT
- Presenters: Maurice Dean and Keith McCallion, Google
- Presenter: Dah Ming Chiu, Chinese University, Hong Kong
- Presenter: Barry Tishgart, Comcast
- Presenter: Ren Provo, Comcast
- Presenter: Richard Steenbergen, Peering DB
- Presenter: Bill Norton, Access Power Peering
- Presenter: Aemen Lodhi, Georgia Tech
- Presenter: Amogh Dhamdhere, CAIDA, UC, San Diego
- Presenter: Patrick Gilmore, Akamai
- Presenter: Geoff Huston, APNIC
- Presenter: Tom Vest, RIPE
- Presenter: Nikolaos Laoutaris, Telefonica
- Presenter: Sergey Gorinsky, Institute IMDEA Networks
- Presenter: Vytautas Valancius, Georgia Tech
- Presenter: John Chuang, UC Berkeley
- Presenter: Heikki Hammainen, Aalto University
- Presenter: Richard Ma, Advanced Digital Science Center, Singapore
- Other Attendees: Susan Martens (AT&T), Emile Aben (RIPE), Brad Huffaker (CAIDA), Joshua Polterock (CAIDA)

ACKNOWLEDGMENTS. Thanks to David Clark, John Chuang, and Sergey Gorinsky for excellent feedback and contributions of text for this workshop report, and thanks to all participants for their insights. The workshop was supported by the sponsors and members of CAIDA, and made possible by the National Science Foundation's Office of Cyberinfrastructure (NSF OCI-0963073) and NSF CISE's Division of Computer and Network Systems (NSF CNS-1017064).

7. REFERENCES

[1] The Workshop on Internet Economics (WIE), 2009. http://www.caida.org/workshops/wie/0909/.

[2] Amogh Dhamdhere and Constantine Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, December 2010. http://www.caida.org/~amogh/papers/iter2-CONEXT10.pdf.

[3] Amogh Dhamdhere, Pierre Francois, and Constantine Dovrolis. A Value-based Framework for Internet Peering Agreements. In *International Teletraffic Congress (ITC)*, September 2010. http://www.caida.org/~amogh/papers/vbp_ITC10.pdf (Also presented at NANOG 49, June 2010).

[4] Ignacio Castro and Sergey Gorinsky. T4P: Hybrid Interconnection for Cost Reduction. In *Proceedings of NetEcon*, 2012.

[5] k. claffy. Workshop on Internet Economics (WIE2009) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, (2), Apr 2010.

[6] David D. Clark, William Lehr, and Steven Bauer. Interconnection in the Internet: The Policy Challenge. In *Telecommunications Policy Research Conference*, September 2011.

[7] John Chuang. Loci of Competition for Future Internet Architectures. *IEEE Communications Magazine*, July 2011.

[8] Level 3 Communications. Level 3 Proposal for Equitable Sharing of Backbone Costs between Peering Networks, August 2011. http://fjallfoss.fcc.gov/ecfs/document/view?id=7021703819.

[9] Murtaza Motiwala, Amogh Dhamdhere, Nick Feamster, and Anukool Lakhina. Towards a Cost Model for Network Traffic. *ACM SIGCOMM Computer Communication Review (CCR)*, January 2012.

[10] Rade Stanojevic, Ignacio Castro, and Sergey Gorinsky. CIPT: Using Tuangou to Reduce IP Transit Costs. In *Proceedings of ACM CoNext*, 2011.

[11] Yochai Benkler, et al. Next Generation Connectivity: A review of broadband Internet transitions and policy from around the world, October 2012. http://www.fcc.gov/stage/pdf/Berkman_Center_Broadband_Study_13Oct09.pdf.