

Workshop on Internet Economics (WIE2009) Report

kc claffy
CAIDA/UC, San Diego
kc@caida.org

ABSTRACT

On September 23, 2009, CAIDA hosted a virtual Workshop on Internet Economics [3] to bring together network technology and policy researchers, commercial Internet facilities and service providers, and communications regulators to explore a common goal: *framing a concrete agenda for the emerging but empirically stunted field of Internet infrastructure economics*. With participants stretching from Washington D.C. to Queensland, Australia, we used the electronic conference hosting facilities supported by the California Institute of Technology (CalTech) EVO Collaboration Network [2]. This report describes the workshop discussions and identifies relevant open research questions identified by 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 Behavioural Sciences]: Economics; K.4.1 [Public Policy Issues]: Transborder data flow

General Terms

Economics, Legal Aspects, Management

Keywords

Economics, Internet, Network management

1. BACKGROUND

Internet economics research exists in a nascent state. Like other sub-disciplines of Internet research, the two biggest impediments to progress are the lack of available empirical data about realistic operational Internet infrastructure, and the highly divergent missions of stakeholders in a position to guide a discipline forward. In 1995, when the U.S. National Science Foundation ended its support of the NSFnet backbone and coordinated a transition of backbone operations to the private sector, few foresaw that the Internet would become the world's primary critical communications substrate. The lightest possible regulatory constraints on Internet evolution was the prevailing philosophy of the U.S. government during the 1990's. Such an approach was arguably consistent with the recent (1984) landmark decision concluding a decades-long effort to deregulate the long-distance telephony industry. We are still struggling to navigate – and even measure – four dimensions of the provisioned infrastructure and its underlying architecture that received little attention during and following the transition: security, scalability, sustainability, and stewardship. Against this backdrop we hosted our first “Workshop on Internet Economics” (WIE09),

seeking to help the research community circumscribe the emerging discipline of Internet economics. Notably, none of the participants would classify themselves as economists, much less “Internet economists”, but all had a growing interest in advancing our understanding of how the underlying economic dynamics of the Internet ecosystem affect the accuracy and interpretability of Internet research and data analysis.

2. TOPICS PRESENTED

A one-day workshop prohibited coverage of the vast range of topics relevant to the economics of Internet service provisioning, so we started with two primary themes reflecting recent research directions and interests of workshop participants. The first session included three talks about the economics of Internet identifiers (IP addresses and AS numbers). The second session focused on the economics of ISP peering and interconnection. Both sessions were followed by moderated roundtables that sought to interpret presented ideas, synthesize results, and identify and evaluate future research directions. Two additional talks hinted at the breadth of the discipline of “Internet economics” and the acute need for empirical data to support scientific study of Internet infrastructure economics. Participants expressed interest in having future workshops in this series discuss related topics in greater depth, e.g., sustainable business models for long-haul backbone networks; the impact of architecture and design on costs; economics of privacy, censorship, and intellectual property.

3. ECONOMICS OF INTERNET PROTOCOL NUMBERING RESOURCES

Perhaps the most acute crisis threatening the Internet today is the exhaustion of the IPv4 address space and the ubiquitous lack of preparation for the transition to a network architecture with a much larger address space. The Internet addressing authority will be out of its pool of available addresses in 2-3 years, but the designed and recommended solution, IPv6, requires infrastructure upgrades that will pose an economic burden to networked organizations in an already challenging fiscal environment. The most indisputably essential characteristic of IPv6 is that it was designed to provide orders of magnitude more address space than the world's foreseeable IP connectivity needs (2^{128} or about 3.4×10^{38} addresses vs. 4.3×10^9 in IPv4). Yet IPv6 is not backward-compatible, i.e., IPv6-capable devices cannot communicate with IPv4-capable devices already in the field, so transition will include the cost of upgrading every IPv4 device to handle the IPv6 protocol, a cost which most IP service providers have no incentive – and in some cases insufficient capital – to cover. So despite appealing features of IPv6, without some exogenous influence as well as available capital for infrastructure

upgrades, IPv6 will likely fail, with significant implications for not only the service provisioning industry, but also for how people use the Internet. An obvious short-term solution to the IPv4 address exhaustion problem is to exercise the clause in RFC2050 to reclaim unused address resources¹, but no one is seriously considering that option for fear of endless lawsuits, and it would only buy time, not solve the fundamental address shortage problem.

The leading alternative strategy, which the three largest RIRs have already established policy to support, is to allow IPv4 address holders to transfer them as if they were property [4]. This policy shift contradicts the most fundamental premise of the current allocation architecture – that IPv4 addresses are not property [1] – considered essential to maintaining scalability of the routing system as well as conservation and fair allocation of the finite IPv4 address space. The prevailing argument in favor of IPv4 address market formation is that allowing people to sell IP addresses will release a lot of otherwise tied up space, and it is more politically palatable than reclamation, especially to holders of IPv4 space, which includes most people involved in address policy development.

The IP address market scenario will also require resolving issues such as who maintains the authoritative database(s) for address ownership, what compels address holders to keep those records current, and what identifying data should be available about address owners. IP address markets will also create a windfall for the U.S. federal government, and accompanying disapproval from the rest of the world. And, like reclamation, allowing an IPv4 address market will only buy us time. Even advocates of the market solution acknowledge that it is technologically [9] as well as economically and socially inferior to a solution that provides publicly recognized IP addresses to anyone who needs them, which IPv4 will never be able to do. If incumbent providers (IPv4 address holders) forgo IPv6 in favor of other IPv4-dependent technologies, it will effectively close the industry to new entrants; indeed, for incumbents, the ability to profit from IPv4 scarcity is a counter-incentive to deploy IPv6. Not only must we expect address monetization to induce consolidation of the address space into fewer hands, but given the financial industry's recent trajectory, we must also consider possible threats from de-coupling the address liquidity mechanism from the productive, innovative elements of the Internet economy [13].

3.1 Current status of IPv6 deployment

Many attempts have been made to evaluate the status of IPv6 adoption and penetration, but none have found significant activity, even though IPv6 has been implemented on all major network and host operating systems. Challenging economic conditions have further lowered the chance that any ISPs will voluntarily invest capital in creating and operating the parallel networks that will be required while the world transitions to IPv6. Worse, most IPv6 traffic currently observed on the Internet is tunneled over IPv4, which delivers potentially two layers of performance problems to the user. Measuring the diffusion of IPv6 technologies, and the influence of fiscal and technological factors on the transition, e.g., IPv4-IPv6 gateways, continues to pose a challenge due to lack of both traffic and financial (cost) data.

3.2 Economics of Routing Table Growth

Regardless of the chosen IPv6 transition strategy (or lack there

¹“The IANA reserves the right to invalidate any IP assignments once it is determined the the requirement for the address space no longer exists. In the event of address invalidation, reasonable efforts will be made by the appropriate registry to inform the organization that the addresses have been returned to the free pool of IPv4 address space.” [8]

of) Internet growth and evolution of the current architecture will inevitably lead to core Internet routing table growth, the underlying economics of which is unclear. Routing table de-aggregation, i.e., splintering of the address space into smaller fractions each represented with its own entry in the routing table, is a natural outcome of inter-domain traffic engineering and other operational ISP business objectives. A market-based approach to Internet address allocation will necessarily induce additional, perhaps dramatically more, deaggregation of the routing table, as address space holders divide up their space to sell (or more likely, rent) in smaller portions to maximize their own income from the resource. Deaggregation will cause two related and costly trends in router architecture design: the need for larger memory to handle the greater number of routing table entries; and more powerful computing technology capable of the increased computational load to execute the routing algorithm on a larger set of more dynamic entries. These two factors of routing table growth will increase the cost of building routers that can handle them, rendering fewer players in the routing ecosystem capable of participating.

Workshop participant Tom Vest explored the growing symmetries between the system of IP addressing and routing and systems of monetary and financial flows. IP addresses share some “liquidity mechanism” characteristics with currency, suggesting potential insight about the IPv6 transition from examining national currency transitions in history [13]. Public deliberations in the RIR communities over the last five years have struggled to develop rules that allow for continued Internet routing scalability as well as fair allocation of resources in the face of IPv4 depletion and stunted IPv6 deployment. Workshop participants did not have solutions to these problems; they only recognized that the technological and economic factors interacting in address allocation/transfer and routing table dynamics – and how to monitor these factors and dynamics – are central questions in the most significant architectural and policy decisions about the Internet today.

3.3 Implications of addressing economics for future designs/architectures

Although predicting the future is the ultimate goal of scientific endeavors, we must acknowledge that the Internet research community still struggles to predict even the past, since there have been so few incentives or resources dedicated to recording history. But we can safely infer that the technical obstacles to IPv6 deployment are lower than the economic and policy obstacles. Profit margins on IP packet transport are insufficient to sustain a competitive IPv4 infrastructure, much less to innovate to a mostly similar new architecture, or to support long-term architectural research. Furthermore, in the last 10 years the Internet has reached a scale so deeply embedded in the political and economic systems of the world, that no existing organization can have legitimate jurisdiction over finishing the important unfinished or outgrown parts of the architecture. Future Internet architecture research initiatives, such as NSF's FIND program, recognize the importance of network economics including addressing economics, but are blocked by policy and economic factors from making substantial interdisciplinary breakthroughs.

3.4 Open questions

1. Who today uses the IPv6 protocol? How can we effectively track and model the diffusion process?
2. Which factors drive or inhibit the market for IPv6? Is IPv6 an example of market failure? (Participant Geoff Huston published his thoughts on this question after the workshop [5].)
3. In what way are IP addresses a liquidity mechanism? What lessons can we draw from transitions by other countries switch-

ing from one currency to another? [13].

4. How will varying address pricing models affect consumers?
5. Will there be a tipping point in deaggregation, where network transit providers will have to selectively route to a part of the address space to maintain stable routing tables? What impact would such a change have on reachability?
6. Would decoupling of location and addressing solve these looming potential problems?

4. ECONOMICS OF INTERCONNECTION

The second session of the workshop mainly focused on the economics of ISP interconnection and how to study it. The Internet consists of thousands of diverse, autonomous interconnected networks with routing and peering policies dictated by diverse economic and strategic objectives, and resulting traffic flow is sometimes only loosely coupled with financial exchanges. The continual emergence of new applications, pricing schemes, protocol and architectural innovations (to some, “hacks”) renders the Internet ecosystem far too intricate for simple models to capture the complexity of network interactions. Most of these interactions are local in nature, without coordinated control or regulation, but they often have global impact affecting the performance and reliability experienced by users, the financial viability of network and service providers, and to some degree the global economy.

One objective of this workshop series is to help the community develop a scientific basis for understanding the structure and dynamics of the Internet infrastructure from an economic perspective, capturing interactions between network business relations, internet-network topology, routing policies, and the resulting interdomain traffic flow. Though there is a large body of work on economic aspects of the Internet such as pricing and interconnection mechanisms, current modeling efforts face two problems. First, the models of a more analytical flavor do not capture most operational realities of the Internet. Second, few models are parameterized or validated using measurements from real networks. Peering agreements are typically treated as trade secrets, under contracts covered by Non-Disclosure Agreements (NDAs). There is thus a disconnect between operational realities and economic models of the Internet, because there has been no way to either corroborate or refute conclusions drawn from previous models. Economically relevant data about interdomain traffic characteristics, routing and peering policies and pricing/cost structures has simply not been available. Empirically parameterized models that capture most real-world intricacies of network interconnection are missing from the literature.

4.1 Analytic models of peering

We began the session with a presentation by Vishal Misra on the history, properties, and applicability of the Shapley value distribution on ISP peering and settlements [12, 15]. Taxonomizing ISP networks into those serving eyeballs, content, or transit, Misra used Shapley values to derive a “fair” split of revenue, assuming players will demand fairness in equilibrium, but acknowledging that operational reality does not necessarily match the assumptions of his model. Ike Elliott from Level3 provided additional insight on the flawed economics of current peering relationships, which often create a greater cost burden for one of the peering partners, and economic advantage for the other. Most early Internet routing peering agreements (1990s) were settlement-free, i.e., networks exchanged traffic without compensation, but primarily only with other networks considered approximate equals in terms of infrastructure, traffic, bandwidth, or importance of connectivity. When these overly simplistic models were recognized as inadequate prox-

ies for value exchanged, some coarse parameterization of peering arrangements emerged, still quite amenable to gaming by network operators seeking to avoid capital investment and instead leverage that of other networks [10]. Today, a wide variety of peering agreements occur between different networks with vastly different infrastructure investments, based on sophisticated and fine-grained measurement-based peering arrangements, including “paid peering”.

Elliott proposed a new framework bilateral peering rules that would help ensure fairness of settlement-free peering relationships. The framework consisted of seven principles aimed at bringing neutrality to peering interconnection agreements [7]:

1. Fair peering relationships place an equal cost burden on each peering partner with a mechanism for correcting imbalance.
2. Fair Peering partners prefer each other’s routes above all others (peers).
3. Fair Peering partners announce full on-net routes to each other.
4. Fair Peering partners commit to augment interconnection capacity to handle all unconstrained traffic demand.
5. Fair Peering partners commit to interconnection at diverse locations to achieve a more survivable Internet.
6. Fair Peering partners commit to providing a minimum level of service quality to each other (uptime, latency, packet loss, and route management), and provide each other a well-defined escalation path for use in the event of quality degradation.
7. Each Fair Peering partner announces their Fair Peers to the public. (Any financial exchange or imbalance would still be under NDA)

A lively discussion of the viability of the framework followed, with a digression into the history of multinational switched voice settlements under the ITU. In the 1980s and 1990s, the international telecommunications market was devastated by arbitrage technologies (callback and redial) that took advantage of a growing gap between ITU-imposed fixed settlement rates for calls terminated in a given country, and the dramatically decreasing cost of completing those calls as technology advanced. The relevant insight from this earlier phase of telecommunications history is that with price regulation, even in self-regulatory regimes, it is hard to avoid situations where price lags cost, encouraging arbitrage, just as happens today in settlement-free peering arrangements [11]. Although Elliott intended the fair peering framework to correct or minimize this type of arbitrage for Internet peering, he did further analysis shortly after the workshop and agreed with several other participants that it was impossible to structure the framework in a way that both sides would have an incentive to voluntarily participate.

Srinivas Shakkotai proposed a scheme that uses game-theoretic techniques to design a P2P overlay that would enable ISPs to cooperate “fairly”, specifically to tradeoff service quality and cost of carrying P2P traffic, by trying to localize traffic as much as possible [14, 6].

4.2 Computational modeling of peering dynamics

Acknowledging the limitations of analytic tools due to their many simplifying assumptions (such as rationality), Constantine Dovrolis proposed a different approach: a computational model that treats the Internet as a dynamic system of agents that make decisions based on limited information. The specific modeling goal is to allow exploration of the presence or lack of equilibria, or profit, in the face of various network peering and infrastructure expansion strategies, and using as input the best available (which is admittedly

not good) data on topology characteristics, routing algorithms and policies, interdomain traffic flow, and pricing. Given observable initial conditions and six different provider/peer selection strategies for each network, he used the model to iteratively compute equilibria states where no node has the incentive to unilaterally change its connectivity. Although still in early stages, Constantine was optimistic that computational powerful and empirically grounded models were more likely to lead to predictive power than prevailing analytic models.

4.3 Data needed to support science and policy

The lack of necessary data to validate scientific models and/or justify [de]regulation was a common theme of workshop discussions. Steven Bauer presented MIT's recent Internet Traffic Analysis Studies (MITAS), a proposed attempt to get providers to volunteer traffic data to MIT researchers trying to understand the health of the Internet. They have limited traffic data from a single provider, with no permission to share it with other researchers, and no results yet, but Steve hoped to have results to present next year.

4.4 Open questions

1. How do settlements occur between ISPs? What models can help yield insights about optimal peering agreements?
2. Could the ISP industry self-regulate through transparent (disclosed), "fair" peering arrangements?
3. How can we get access to more data about peering relationships in the face of the fact that they are mostly treated by participants as trade secrets?
4. What are the correct metrics for congestion, and what data on congestion or utilization should be gathered and/or published?
5. What are the underlying cost structures for carrying traffic and expanding capacity? What would be the effects of different traffic management policies?
6. What indicators should regulators examine to better understand Internet market structure and dynamics at a country (or other) granularity? How might we use IP address assignments across national boundaries to use as metric for economic analysis?
7. How can researchers connect their studies to relevant political units, country, state, local municipalities?

5. WORKSHOP PARTICIPANTS

- Chairperson: kc claffy, CAIDA, UC, San Diego
- Presenter: Tom Vest, Eyeconomics, "The Internet as a Liquidity Mechanism: From Analogy to Isomorphism"
- Presenter: Geoff Huston, APNIC, "Is the transition to IPv6 a market failure?"
- Presenter: Roch Guerin, U. Pennsylvania, "Competing Network Technologies: The Role of Gateways"
- Moderator: Paul Vixie, Internet Systems Consortium
- Presenter: Vishal Misra, Columbia U., "A Shapley Value Perspective on ISP Settlements"
- Presenter: Srinivas Shakkottai, Texas A&M, "Designing ISP-friendly Peer-to-Peer Networks Using Game-based Control"
- Presenter: Ike Elliott, Level 3, "Fair Peering"
- Presenter: Constantine Dovrolis, Georgia Institute of Technology, "ITER: A Computational Model to Evaluate Provider and Peer Selection in the Internet Ecosystem"
- Presenter: Mark Cooper, Stanford, "The Economics of Digital Content: How to Win Friends, Influence People and Make

a Little Money where the Sneaky Exponential Trumps the Long Tail in Cyberspace"

- Presenter: Steven Bauer, MIT, "Broadband Microfoundations: the Need for Traffic Data"
- Moderator: Irene Wu, FCC
- Kevin Werbach, U. Pennsylvania
- Jon Callahan, Level 3
- Amogh Dhamdhere, CAIDA, UC, San Diego
- Mia Ahang, CAIDA, Beijing Jiaotong U.
- Daniel McCartney, Michigan State Law School
- Andrew Odlyzko, U. of Minnesota
- Sara Wedeman, Behavioral Economics

ACKNOWLEDGMENTS. The workshop was supported by CAIDA sponsors and members, and in particular made possible by a gift from the American Registry for Internet Numbers (ARIN) to study IPv6 evolution. We thank all participants for their insights.

6. REFERENCES

- [1] ARIN Number Resource Policy Manual, 2008. <https://www.arin.net/policy/nrpm.html>.
- [2] Evo, 2009. <http://evo.caltech.edu/evoGate/>.
- [3] The workshop on internet economics (wie), 2009. <http://www.caida.org/workshops/wie/0909/>.
- [4] Benjamin Edelman. *Running Out of Numbers: The Impending Scarcity of IP Addresses and What To Do About It*. HBS, June 2008. <http://www.hbs.edu/research/pdf/09-091.pdf>.
- [5] Geoff Huston. Is the Transition to IPv6 a "Market Failure?", September 2009. <http://www.potaroo.net/ispcol/2009-09/v6trans.html>.
- [6] H. Xie and Y.R. Yang and A. Krishnamurthy and Y. Liu and A. Silberschatz. P4P: Provider portal for applications. In *Proceedings of ACM SIGCOMM*, 2008.
- [7] Ike Elliot. Fair Peering, 2009. WIE2009 workshop presentation, http://www.caida.org/workshops/wie/0909/slides/wie0909_ielliott.pdf.
- [8] K. Hubbard, M. Koster, D. Conrad, D. Karrenberg and J. Postel. RFC 2050. Internet Registry IP Allocation Guidelines, 1996. <http://www.ietf.org/rfc/rfc2050.txt>.
- [9] Olaf Maennel, Randy Bush, Luca Cittadini, and Steven M. Bellovin. A Better Approach than Carrier-Grade-NAT. Technical Report CUCS-041-08, Columbia University, 2008.
- [10] William B. Norton. The Art of Peering: The Peering Playbook. *Equinix white papers*, 2002.
- [11] William B. Norton. The Evolution of the U.S. Internet Peering Ecosystem. *Equinix white papers*, 2004.
- [12] R. Ma and D.-m. Chiu and J. Lui and V. Misra and D. Rubenstein. On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers. In *ACM CoNEXT*, December 2008.
- [13] T. Vest. Medium of Exchange, 2008. <http://eyeconomics.com/eyeconomics.com/MoEDdefined.html>.
- [14] V. Aggarwal and A. Feldmann and C. Scheideler. Can ISPs and P2P users cooperate for improved performance? *ACM SIGCOMM Computer Communications Review*, 37(3), 2007.
- [15] Wikipedia. Shapley value, 2009. http://en.wikipedia.org/wiki/Shapley_value.