Incentive-Compatible Caching and Peering in Data-Oriented Networks

Jarno Rajahalme† Mikko Särelä‡ Pekka Nikander‡§ Sasu Tarkoma§¶

ABSTRACT

Several new, data-oriented internetworking architectures have been proposed recently. However, the practical deployability of such designs is an open question. In this paper, we consider data-oriented network designs in the light of the policy and incentive structures of the present internetworking economy. A main observation of our work is that none of the present proposals is both policy-compliant and incentive-compatible with the current internetworking market, which makes their deployment very challenging if not impossible. This difficulty stems from the unfounded implicit assumption that data-oriented routing policies directly reflect the underlying packet-level inter-domain policies. We find that to enable the more effective network utilization promised by data-oriented networking, essential caching incentives need to exist, and that data-oriented peering needs be considered separately from peering for packet forwarding.

KEYWORDS: Inter-domain policy, Peering, Data-oriented Networking, Internetworking, Network Architecture

1 INTRODUCTION

The Internet architecture was originally developed for the needs of distributed computing, such as telnet, FTP, and RJE [4]. However, the majority of the present Internet traffic, especially between network domains, is of content delivery nature, such as file sharing, static web content, Internet radio, or other recorded voice or video, or control traffic needed to locate or deliver any wanted pieces of content. Hence, the current situation can be characterized by stating that while the network use has shifted towards content-centric usage patterns (see e.g. [13]), the original host-oriented architecture, which we still use, is optimized for interactive communication between topologically addressed end-points.

To address the shift in the network usage, several content, information, or data-oriented models have been proposed in the literature (see e.g. [11], [14]). We will briefly explore the relevant features of some of these in Section 3 below; however, in general the data-oriented architectures base their routing decisions on information or content-related identifiers instead of topological prefixes or host addresses (locators).

In effect, due to their location independent identifiers data-oriented networking architectures enable a wider choice of internetworking control for the ASes. In the current BGP model, an AS can merely advertise reachability of a certain IP address block to a neighbor. It has no control over how that reachability will be propagated or used. Consequently, due to the advertisement, all kinds of traffic not envisioned by the originating AS may take place. As an example, consider the unexpected rise of the traffic between home users.

However, if the ASes advertise the availability of data, as in DONA [14], they would know that such an advertisement may only produce requests for that specific piece of data. This level of control comes with a cost: there are orders of magnitude more data items than network routes to keep track of. Therefore, the inter-domain scalability is a major challenge to any data-oriented network design. The wider implications of such added control are not well understood.

All *clean-slate* designs are bound to have unsolved deployment issues.[1] In this paper we address the policy-compliancy and incentive-compatibility of the data-oriented architectures with the current internetworking market, as well as possible implications for data-oriented peering.

Assuming a suitable relative balance between communication and storage costs, even a superficial analysis shows that in a data-oriented architecture it makes sense to cache and disperse data through the network in various ways, thereby reducing the amount of network traffic. At the practical level, this is clearly shown to be the case by the proliferation of content delivery networks, such as Akamai, and the widespread practice of caching web pages.

However, we argue that the present internetworking market is structured in a way that creates disincentives for some of the ISPs to deploy content caching, and this presumably acts as a disincentive against data-oriented networking in

This work was supported by the EU FP7 PSIRP project under grant INFSO-ICT-216173.

†Nokia Siemens Networks
‡Ericsson Research Nomadic Lab
§Helsinki Institute for Information Technology HIIT
¶Helsinki University of Technology, Computer Science and Engineering Dept.

ReArch'08, December 9, 2008, Madrid, SPAIN.

[1] For example, see the problems with the IP multicast deployment [5].

general. Therefore, in order to allow the network architecture to evolve towards a data-oriented one, the caching policies (e.g. what to cache and when) will need to be separated from actual caching mechanisms (e.g. packet inspection, storage) [21], thereby allowing them both to evolve independently and free from policy-related disincentives.

Today, the Internet topology is defined by inter-domain policies, established between autonomous systems (ASes) [9], or domains for short. These policies reflect the business relationships between the domains [12]. In a way, these relationships are more fundamental in nature than the deployed network architecture itself. The policies define who carries whose traffic and where, not the architecture. Hence, also the proposed data-oriented network architectures must take this economic aspect into account, as there is no chance for an abrupt change in the inter-domain business relationships [18].

For example, some of the data-oriented architecture proposals make the following simple assumption:

> "ISPs who already peer at the IP routing level are [expected to be] motivated to peer at the content routing level to provide their customers faster access to nearby content servers – and increase the benefit of placing content servers in their network." [11, p. 9]

Unfortunately, as our analysis in Section 4 shows, this assumption is only partially valid.

Hence, while we contend that it may well be possible to evolve the data-oriented architectures so that they can be bootstrapped off from the current policy structure, we also note that the present policy-constrained inter-domain paths of the Internet are often not the shortest ones [22], or the ones with the best performance [20]. That is, we surmise that the richer network service model, provided by the data-oriented network architectures, opens up the possibility of using more efficient policies, leading to better network utilization and lower costs; see Section 5.

In the rest of this paper, we first briefly describe how the inter-domain topology of the Internet arises from the inter-domain policies in Section 2. We then, in Section 3, take a brief look at the proposed data-oriented network architectures, and place them on a simple policy-constrained Internet scenario. In Section 4, we analyze the incentives of various network stakeholders to participate in the data-oriented networking models. Finally, in Section 5, we describe how data-oriented caching, if deployed, could enable new inter-domain peering policies. We outline our planned future work in Section 6 and give some preliminary conclusions in Section 7.

2 POLICY-CONSTRAINED INTER-DOMAIN TOPOLOGY

The Internet routing system divides to two well-defined regions. The intra-domain routing protocols (OSPF [15], IS-IS [17]) are used within administrative domains (ASes), while the inter-domain routing is based on the Border Gateway Protocol (BGP, [19]). The contractual relationships between domains (the inter-domain policies) define the global Internet topology, not the underlying physical connectivity [9], [22]. Since the policies are different for each domain, the inter-domain topology is effectively different for each pair of source and destination [7]. While this complicates the overall inter-domain topology picture, it also narrows down the choices for inter-domain routing, and in some sense makes the inter-domain routing space more tractable than without any such limitation [23].

The two main types of inter-domain relationships in the Internet are the *transit* and *peering* relationships [12], while a third type, the *sibling* relationship, is used between the multiple domains under the same administration [9].

A transit relationship is a customer-provider relationship, where the provider agrees to advertise the customer's routes and thus make the customer's network reachable from the rest of the Internet (as seen by the provider). The customer also gets routes to all Internet destinations, or uses the provider as a default route, and thus can access the rest of the Internet. The customer pays the provider for these services (reachability and access) [6]. In a *partial transit* relationship only a subset of the Internet is made available.

In a peering relationship the ASes agree to exchange traffic only between themselves and their customers by advertising the corresponding routes to each other. The typical incentives for peering are, for large ASes, reciprocal reachability and robustness benefits, and, for smaller ASes, savings on the upstream transit costs [2]. This is why also big content providers, such as YouTube (AS36561), actively seek peering relationships with other networks [3]. While peering is typically settlement-free, also *paid peering* is possible.

This structure leads to a tiered Internet model, where the ASes at the topmost tier (the Tier-1) form, in practice, an oligopoly that does not buy transit from anyone else; all are peering with each other. Of course, this makes the whole Internet connected.[2]

It is important to understand here that since the non-Tier-1 ISPs actively seek to reduce their transit costs, *any architectural change that reduces the amount of traffic over the inter-domain transit links may lead to reduced Tier-1 income*. For example, if an architectural change allows Tier-2 ISPs to avoid sending or receiving some traffic that they would otherwise need to transmit through the Tier-1 ISPs, the economic effect is the same as in the case of the Tier-2 ISPs peering directly: the Tier-2 ISPs manage to reduce their costs and therefore the Tier-1 oligopoly loses some of its income.

In [9], Gao derives an empirical *valley-free* inter-domain path model spanning from the bilateral relationships between ISP, as deducible from the BGP routing data. In this model, all Internet paths can be divided into three segments, out of which any segment may be omitted for a given path: 1. The

[2] We are ignoring any possible *local tier-1s* not purchasing transit, but not peering with all other Tier-1 networks.

uphill path is the sequence of domains from the source domain up to the first provider-to-customer or peer-to-peer link. All the links in this segment are either customer-to-provider, or sibling-to-sibling links. 2. A peer-to-peer link in the middle. This may be e.g. a link between two Tier-1 providers, or any peering link between ISPs. 3. The *downhill path* is the sequence of domains from the first provider-to-customer link to the destination domain. All the links in this segment are either provider-to-customer, or sibling-to-sibling links.

Due to the existing inter-domain policies, practically all Internet paths observe this valley-free model. Gao [10] finds that the maximum AS path length yielding from this structure is about 13, while the maximum shortest AS path length[3] would be around 6. The reason for this increased stretch is that the possible "detours" on the shorter paths [22] have no incentive to provide the transit service between the end-points. Later we examine whether data-oriented networking might enable such incentives (Section 5).

3 DATA-ORIENTED NETWORKING

In this section, we briefly consider three architectures that we consider leading the way towards data-oriented networking, each with different namespace design and scalability properties: TRIAD [11], ROFL [1], and DONA [14].

TRIAD uses a BGP-like inter-domain routing protocol to distribute routes to servers identified with DNS names ("BGP with names"). TRIAD assumes that ISPs would be naturally willing to peer on the name level, if they are already peering on the IP level. Furthermore, TRIAD assumes also that the ISPs would willingly advertise their customers' name-based routes to the rest of the Internet. This model leads to the default-free domains[4] carrying routes to all named services. TRIAD addresses the resulting scalability challenge mainly by restricting the managed namespace to service names, individual hosts or data items are not named separately.

ROFL explores the possibility of routing on flat labels in the Internet scale without the underlying IP forwarding assumed by TRIAD and DONA. ROFL does not limit the number of labels assigned to each host, which makes it possible to extend the ROFL anycast design to address individual data items. However, this increases the number of labels in the system by several orders of magnitude.

In ROFL, inter-domain routing is based on a hierarchical distributed hash table structure (Canon [8]). The ROFL use of hierarchical DHT design avoids aggregating all the routing state to any single ISP, but with the penalty that the ROFL routes are 2-3 times longer than BGP-routes. More significantly, the routes may require traversing domains in a manner clearly violating the inter-domain routing policies. The latter is a feature of the simplistic application of the Canon

[3]Maximal shortest path between any two ASes in the global AS graph *ignoring the inter-domain policies*.

[4]The domains that do not have default routes and thus must carry the full Internet routing table, e.g. the Tier-1 domains.

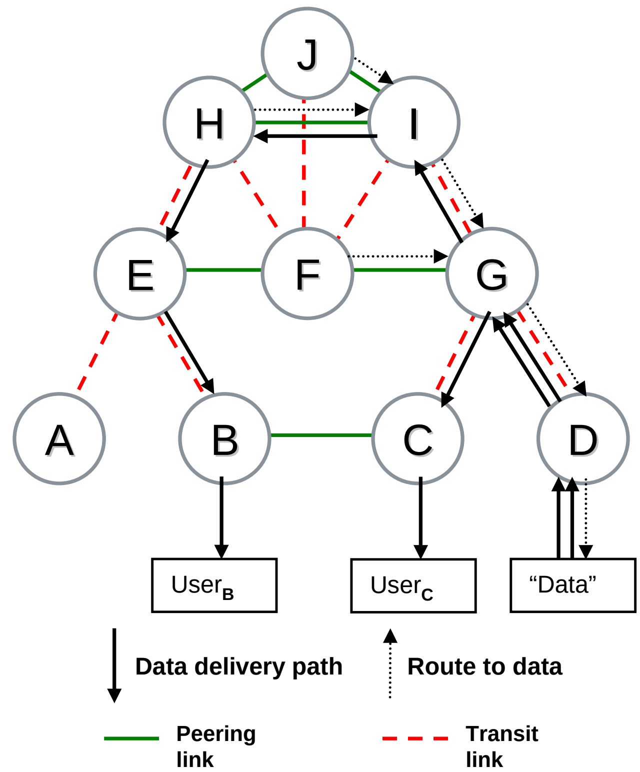

Figure 1: Policy-constrained data-oriented networking. The valley-free path segments for the User$_B$ are: (D-G-I) (uphill), (I-H) (peering), and (H-E-B) (downhill).

method to the inter-domain hierarchy, making traffic from one customer domain to be potentially routed via arbitrary other (possibly competing) customer domains of a given ISP; supposedly, this might be avoidable through a different inter-domain routing design.

DONA takes the TRIAD design further by replacing the DNS names with self-certifying identifiers. DONA also takes the presented peering assumptions literally, and explicitly builds the architecture on the valley-free policy-constrained AS topology. In DONA, each domain maintains data-level routing state for the data items hosted by itself, or by any customer or peering domain. This model leads to similar scaling behavior to TRIAD, but on a bigger scale: Routing state for all advertised data items must be maintained by all the Tier-1 domains. As our analysis later shows, these domains may have little interest in participating in such a data routing architecture.

Now consider the scenario in Figure 1. Here the dashed lines represent transit relationships between the customer (below) and the provider domains (above), and the solid lines represent peering relationships. This example follows the DONA model, where the data-level routing state pointing to the closest copy of a specific data item is distributed to all peers, providers, and their peers and providers (dotted arrows). The data requests are forwarded up along the provider hierarchy and then following the specific data routing state when found. Here we assume symmetrical routing, making the resulting data path (solid arrows) to be the reverse of the domain-level path followed by the data request. This ensures that the data paths are always policy-compliant. The *domain-level stretch* for the User_B in this case is 1.7: The policy-compliant domain-level path (D-G-I-H-E-B) has 5/3 times the hops than the shortest possible domain-level path (D-G-C-B).

The possibility for transparent, architecturally integrated route and/or data caching is a central piece in the data-oriented designs considered above. However, so far the architecture proposals have left a crucial policy question unanswered: When, exactly, should a domain cache which piece of data? For example, in the Figure 1 above, without caching or data specific forwarding state in either of the domains D or G, the data is sent twice from the server hosting the data. To understand why this is the case we will need to examine the incentives for caching.

4 INCENTIVE-COMPATIBILITY OF DATA-ORIENTED NETWORKING

In this section, we take a look at the potential incentives for data-oriented networking, analyzing the economically motivated stances of the different domains forming the Internet. To begin with, we adopt the present, underlying commercial structure for packet routing and forwarding. Starting from there, we explore the potential changes or challenges brought forward by the data-oriented model, especially regarding caching, when employed on the top of the basic model.

One motivational factor for data-oriented networking is the expected more efficient and timely use of network resources. This is possible due to the network routing state being created by the data requests, which enables sharing the communication and storage resources between multiple recipients. In the case of synchronous transmission, this kind of state and resource sharing essentially results in a multicast service. Caching, in turn, enables asynchronous sharing. In general, this kind of sharing has similar incentive structures to peering for packet forwarding: savings on transit costs and reduced latency.

As communication and storage sharing enables higher utilization of the existing peering relationships, some of the weight of the overall traffic moves down from the Tier-1s to the lower tiers. This forms the first challenge for the so-far presented data-oriented networking proposals: If the existing Tier-1 domains are faced with less growth, or even a small decline in their transit traffic, and therefore less income, why should they participate in the data-oriented networking in the first place? Why should they invest in new data-centers, such as ones needed for DONA, if that will cut their profits? Apparently, we need a data-oriented architecture that is not critically dependent on the Tier-1 networks for other than the present packet-forwarding transit service.

The next question is cache placement: Who has an incentive to cache the data? It seems clear that Tier-1 domains, in their transit role, have *no incentive to cache data, if the data served from their caches would be away from their customer links*, thereby reducing their transit-fee income. The same applies to *all uphill domains*: data served from caches will, in general, be away from their customer links, and therefore cuts their revenues. This answers our question at the end of Section 3: *There is no point serving data on your customer's behalf, unless the customer pays for the caching service*. Therefore there is no caching in either of the domains D or G in the Figure 1 above. However, on the *downhill* paths all domains seem to have an incentive to cache the data, as data served from caches is away from their internal or external transit links. The customer network domains, being at the bottom of the downhill paths, have the most obvious incentive for caching, as they can also directly benefit from the reduced latency.

Note, however, that most domains will act in both roles (uphill/downhill) at the same time for different data streams and may have considerable internal transit costs. Therefore, the caching policy needs to be based on the traffic direction in each case. Furthermore, the actual caching decision needs to be based on factors such as the pricing model and traffic situation on the intra- and inter-domain links, the cost of caching itself, and the popularity of the data item in question, among other considerations.

5 DATA-ORIENTED INTER-DOMAIN PEERING

Earlier we saw that the so-far proposed data-oriented architectures have assumed that peering on the IP forwarding level would automatically lead to corresponding peering at the data routing level. Then we observed that the resulting policy-compliant paths are sometimes longer than would be possible based on connectivity alone, but the current inter-domain policy structure does not allow the shorter paths to be taken.

The discussion so far points to an interesting question, namely are there cases in which the *valley-free* model could be applied to the data-oriented networking in a relaxed form so that some of the shorter inter-domain paths could be utilized? Are there possibilities for *data-oriented peering*?

Consider the scenario in Figure 2. The domain C is caching the data the User_C has requested, so that other users in domain C interested in the same data might get it faster. Now it becomes possible for the domain C to further advertise the same data over its peering links. The rationale for doing this may be that the settlement-free peering link (C-B) may be underutilized, and that providing the data over the peering link is not causing any new external transit costs for the domain C. Now, comparing to the situation in Figure 1,

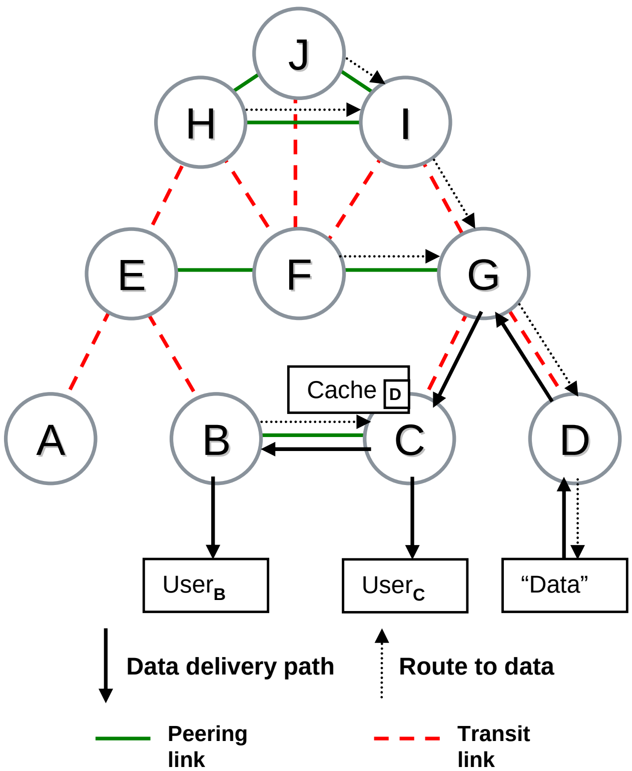

Figure 2: Data-oriented peering.

the domain-level path from D to B becomes shorter, the end-to-end latency likely smaller, and domain B saves the cost of transit through E. In effect, the domain C is in a transit role through its peering and transit relationships with the domains B and G, respectively. Trying to do this with the existing packet forwarding level inter-domain policies makes no sense, since the downhill path (G-C-B) would contain a peering link (C-B), which is excluded in the currently observed valley-free model.[5]

Each domain needs to apply local policy and the current intra- and inter-domain traffic situation to decide when to extend the data advertisements to any of its peering domains. These decisions need not be coordinated with the other domains. As the network situation fluctuates, the domain may vary the data set it is advertising over the peering links. In a data-oriented network, more dynamic inter-domain peering policies become possible also due to the looser coupling between the data and packet-level routing and forwarding functions. For example, in Figure 2, when User_C ceases to request the data, the domain C can decide not to advertise the specific data to its peers anymore, and the domain B will need to revert to the valley-free route presented in Figure 1. These dynamic changes are not visible at the packet forwarding level routing policies.

[5]The reason for the exclusion is that the transit accounting does not transfer over the peering link.

In the current BGP4 Internet, each AS is inviting traffic with BGP route updates over the inter-domain links, but each AS is in direct control of only the traffic they send out. Consequently, a peering link is considered imbalanced if one peer is sending significantly more traffic than the other, the limit reported in the literature being 4:1 [16] or 2:1 [6]. In the data-oriented model the sender of the data is only in indirect control of the data being sent, as data availability is being advertised. The receiver initiates the data transfer explicitly, and thus can be considered the benefactor of the traffic instead of the sender. Thus we believe that the accounting model for peering balance needs to evolve to signify the benefactor of the data-oriented traffic. This would enable avoiding settlement fees by advertising and providing more data over imbalanced peering links. While this further increases the imbalance on the packet-forwarding level, the overall benefits between the peers would be balanced.

Hence, data-oriented peering provides a new application independent means for optimizing network resources, in ways that are not possible at the packet forwarding level only. Architecturally, the peering and caching mechanisms need to exist in each domain, but the peering and caching policies need to be separated from the mechanisms.

Depending on the inter-domain relationships the peering model may combine the packet and data level considerations, or they can be managed separately. The latter case is needed in overlay architectures, where some domains only provide packet forwarding service.

6 FUTURE WORK

As we have observed in this paper, many current data-oriented network architectures, for example DONA, appear to have deployment related incentive challenges. We believe that data-oriented network architectures require further examination for them to become both policy-compliant with the existing network structure and incentive-compatible with the critical stakeholders of these architectures.

Based on the observations made in this paper, we believe that the following areas offer the possibility for new insights and possible solutions regarding data-oriented internetworking: 1. Game theoretic and economic analysis of the caching and peering policy cases summarized in this paper, 2. analysis of the data-oriented model on a realistic, policy augmented inter-domain model, and 3. exploring the possibility of evolving the proposed data-oriented architectures to be more incentive-compatible with the business benefits of each domain in their role in the global Internet.

7 CONCLUSIONS

We have shown that the proposed data-oriented internetworking architectures have inter-domain policy and deployment related incentive problems.

It seems prudent not to expect Tier-1 ISPs to actively participate in creation of the data-oriented Internet, as the increased efficiency of the network use will limit the growth of their transit traffic and thus revenue, at least in the short term. Instead, the initial deployment burden of the architecture should be placed where the immediate benefits exist: the end customers and access network providers, as well as content service providers.

The envisaged increased network utility is based on the more effective use of the existing network resources and peering relationships, as the data-oriented model allows provision of cached content to peers without incurring additional (external) transit costs. In some cases this will lead to shorter paths and thus lower end-to-end latency, as well as to reduced transit traffic, as the number of copies being sent over the transit networks decreases.

The higher-level service model of the data-oriented architectures makes the caching and peering policy a central tool for inter-domain traffic engineering and management. However, these policies are non-trivial. The different roles the domain may have in the valley-free model determine the effect of caching on its revenue. The overall picture contains also the various internal and external costs, such as those related to storage and bandwidth. Also, the looser coupling between the packet forwarding and data-oriented routing functions allows better control over peering traffic than what is currently possible with BGP.

Finally, we believe that the accounting model for peering balance needs to evolve to signify the benefactor of the data-oriented traffic, instead of just assuming that the sender of the packets should pay.

ACKNOWLEDGMENTS

We thank our colleagues in the PSIRP project for encouragement and support, and the anonymous reviewers for their insightful comments and suggestions.

REFERENCES

[1] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, and S. Shenker. ROFL: Routing on Flat Labels. In *ACM SIGCOMM'06. Proceedings*, pages 363–374, September 2006.

[2] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the internet's as-level topology. *IEEE INFOCOM'06. Proceedings*, pages 1–12, April 2006.

[3] C. Corbett. Peering of Video. Presentation at NANOG 37, available online at `http://www.nanog.org/mtg-0606/pdf/bill.norton.3.pdf`, June 2006.

[4] J. Day. *Patterns in Network Architecture: A Return to Fundamentals*. Prentice-Hall, 2008.

[5] C. Diot, B. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the IP multicast service and architecture. *Network, IEEE*, 14(1):78–88, 2000.

[6] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. Complexity of internet interconnections: Technology, incentives and implications for policy. In *The 35th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2007.

[7] J. Feigenbaum, V. Ramachandran, and M. Schapira. Incentive-Compatible Interdomain Routing. Tech. Report 1342, Yale Univ., May 2006.

[8] P. Ganesan, K. Gummadi, and H. Garcia-Molina. Canon in G major: designing DHTs with hierarchical structure. *Distributed Computing Systems. Proceedings*, pages 263–272, 2004.

[9] L. Gao. On inferring autonomous system relationships in the Internet. *Networking, IEEE/ACM Transactions on*, 9(6):733–745, Dec 2001.

[10] L. Gao and F. Wang. The extent of AS path inflation by routing policies. *IEEE GLOBECOM'02. Proceedings*, 3:2180–2184, Nov. 2002.

[11] M. Gritter and D. R. Cheriton. An Architecture for Content Routing Support in the Internet. In *USENIX USITS'01. Proceedings*, Berkeley, CA, USA, 2001.

[12] G. Huston. Interconnection, Peering, and Settlements. In *Proc. Internet Global Summit (INET)*, Jun. 1999.

[13] V. Jacobson. A New Way to look at Networking. Talk at Google, available online at `http://video.google.com/videoplay?docid=-6972678839686672840`, 2006.

[14] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica. A Data-Oriented (and Beyond) Network Architecture. In *ACM SIGCOMM'07. Proceedings*, pages 181–192, 2007.

[15] J. Moy. Open Shortest Path First (OSPF) version 2. RFC 2328, The Internet Engineering Task Force, Apr. 1998.

[16] W. B. Norton. Internet Service Providers and Peering. Technical report, Unpublished draft, 2001. `http://www.scs.stanford.edu/nyu/04sp/sched/peering.pdf`.

[17] J. Parker, Ed. Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS). RFC 3787, The Internet Engineering Task Force, May 2004.

[18] S. Ratnasamy, S. Shenker, and S. McCanne. Towards an Evolvable Internet Architecture. In *ACM SIGCOMM'05. Proceedings*, pages 313–324, 2005.

[19] Y. Rekhter and T. Li. A Border Gateway Protocol (BGP-4). RFC 1771, The Internet Engineering Task Force, Mar. 1995.

[20] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: Informed Internet Routing and Transport. *Micro, IEEE*, 19(1):50–59, Jan/Feb 1999.

[21] A. Snoeren and B. Raghavan. Decoupling policy from mechanism in Internet routing. *ACM SIGCOMM Computer Communication Review*, 34(1):81–86, 2004.

[22] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin. The Impact of Routing Policy on Internet Paths. *IEEE INFOCOM'01. Proceedings*, 2:736–742 vol.2, 2001.

[23] X. Yang, D. Clark, and A. Berger. NIRA: A New Inter-Domain Routing Architecture. *Networking, IEEE/ACM Transactions on*, 15(4):775–788, Aug. 2007.