

ViAggre: Making Routers Last Longer!

Hitesh Ballani
Cornell University

Paul Francis
Cornell University

Tuan Cao
Cornell University

Jia Wang
AT&T Labs – Research

Abstract

This paper presents ViAggre (Virtual Aggregation), a “configuration-only” approach to shrinking the routing table on routers. ViAggre applies to legacy routers and can be adopted independently and autonomously by any ISP. ViAggre is effectively a scalability technique that allows an ISP to modify its internal routing such that individual routers in the ISP’s network only maintain a part of the global routing table. We find that ViAggre can shrink the routing table on routers by more than an order of magnitude while imposing negligible traffic stretch.

1 Introduction

The Internet default-free zone (DFZ) routing table has been growing at a rapid rate for the past few years [1]. Looking ahead, there are concerns that as the IPv4 address space runs out, hierarchical aggregation of network prefixes will further deteriorate resulting in a substantial acceleration in the growth of the routing table [2]. A growing IPv6 deployment would worsen the situation even more [3].

The increase in the size of the DFZ routing table has several harmful implications for inter-domain routing. At a technical level, increasing routing table size may drive high-end router design into various engineering limits. For instance, while memory and processing speeds might just scale with a growing routing system, power and heat dissipation capabilities may not [4]. On the business side, it makes networks less cost-effective by increasing the cost of forwarding packets [5] and making it harder to provision networks, not to mention the cost of actually upgrading the routers to account for larger routing tables. As a matter of fact, instead of upgrading their routers, a few ISPs have resorted to filtering out some small prefixes (mostly /24s) which implies that parts of the Internet don’t have reachability to each other [6]. It is a combination of these possibilities that led a recent Internet Architecture Board workshop to conclude that scaling the routing system was one of the most critical challenges of near-term Internet design [4].

The severity of the routing scalability problem has also meant that a number of proposals have focussed on reducing the size of the DFZ routing table [3, 7–14]. However, all these proposals require changes in the routing and addressing architecture of the Internet and perhaps this has contributed to the fact that

none of them have seen deployment.

An alternative is to tackle the routing scalability problem through a series of incremental, cost-effective upgrades. Guided by this, we propose Virtual Aggregation or *ViAggre*, a “configuration-only” solution that shrinks the routing table on routers.¹ *ViAggre applies to legacy routers*. Further, it can be *adopted independently and autonomously by any ISP* and hence the bar for its deployment is much lower. In effect, ViAggre is a scalability technique that allows an ISP to modify its internal routing such that individual routers in the ISP’s network only maintain a part of the global routing table. In this paper, we briefly discuss two deployment options through which an ISP can adopt ViAggre.

Preliminary results show that ViAggre can reduce the size of routing tables on routers by more than an order of magnitude while imposing negligible stretch on traffic. However, several important questions remain unanswered. These include the impact of an ISP adopting ViAggre on router load, network complexity and network robustness. We discuss ongoing work that aims to answer these questions. In spite of these questions, we believe that its simplicity makes ViAggre an attractive short-term alternative that can be used by ISPs to cope with the growing routing table till more fundamental, long-term architectural changes can be agreed upon and deployed in the Internet.

2 ViAggre design

ViAggre allows individual ISPs in the Internet’s DFZ to do away with the need for their routers to maintain routes for all prefixes in the global routing table. An ISP adopting ViAggre divides the global address space into a set of *virtual prefixes* that are larger than any aggregatable prefix in use today. For instance, an ISP could divide the IPv4 address space into 128 parts with a /7 representing each part (0.0.0.0/7 to 254.0.0.0/7). Note that such a naïve allocation would yield an uneven distribution of real prefixes across the virtual prefixes. However, the virtual prefixes need not be of the same length and as long as the virtual prefixes together cover the complete address space, the ISP can choose them such that they contain a comparable number of real prefixes.

¹Specifically, we focus on the router Forwarding Information Base (FIB).

The virtual prefixes are not topologically valid aggregates, i.e. there is not a single point in the Internet topology that can hierarchically aggregate the encompassed prefixes. ViAggre makes the virtual prefixes aggregatable by organizing *virtual networks*, one for each virtual prefix. In other words, a virtual topology is configured that causes the virtual prefixes to be aggregatable, thus allowing for routing hierarchy that shrinks the routing table. To create such a virtual network, some of the ISP’s routers are assigned to be within the virtual network. These routers maintain routes for all prefixes in the virtual prefix corresponding to the virtual network and hence, are said to be *aggregation points* for the virtual prefix. A router can be an aggregation point for multiple virtual prefixes and is required to only maintain routes for prefixes in the virtual prefixes it is aggregating.

Given this, a packet entering the ISP’s network is routed to a close by aggregation point for the virtual prefix encompassing the actual destination prefix. This aggregation point has a route for the destination prefix and forwards the packet out of the ISP’s network. In figure 1 (figure details explained later), router C is an aggregation point for the virtual prefix encompassing the destination prefix and $B \rightarrow C \rightarrow D$ is one such path through the ISP’s network.

2.1 Design Goals

The discussion above describes ViAggre at a conceptual level. However, the design space for organizing an ISP’s network into virtual networks is characterized by several dimensions. For example, the flexibility to change the ISP’s topology or to change the routers themselves lead to very different architectures, all of which allow for virtual prefix based routing. However, this paper aims for deployability and hence is guided by two major design goals:

1. *No changes to router software and routing protocols*: The ISP should not need to deploy new data-plane or control-plane mechanisms.
2. *Transparent to external networks*: An ISP’s decision to adopt the ViAggre proposal should not impact its interaction with its neighbors (customers, peers and providers).

These goals, in turn, limit what can be achieved through the ViAggre designs presented here. Routers today have a Routing Information Base (RIB) generated by the routing protocols and a Forwarding Information Base (FIB) that is used for forwarding the packets. Consequently, the FIB is optimized for looking up destination addresses and is maintained on fast(er) memory, generally on the line cards themselves. All things being equal, it would be nice to

shrink both the RIB and the FIB for all ISP devices, as well as make other improvements such as speed up convergence time.

While the basic ViAggre idea can be used to achieve these benefits (section 5), we have not been able to reconcile them with the aforementioned design goals. This paper takes the position that given the performance and monetary implications of the FIB size for routers, an immediately deployable solution that reduces FIB size is useful. Actually, one of the presented designs also shrinks the RIB on routers; only components that are off the data path need to maintain the full RIB. The rest of this section abuses terminology and uses the term “ViAggre” to refer to the specific design being presented.

2.2 Design-I

This section details one way an ISP can deploy virtual prefix based routing while satisfying the goals specified in the previous section. The discussion below applies to IPv4 (and BGPv4) although the techniques detailed here work equally well for IPv6. The key concept behind this design is to operate the ISP’s routing untouched and in particular, to populate the RIB on routers with the full routing table but to suppress most prefixes from being loaded in the FIB of routers. A standard feature on routers today is to prevent routes for individual prefixes in the RIB from being loaded into the FIB. We have verified this as part of our ViAggre deployment on Cisco 7300 and 12000 routers. Documentation for Juniper [15] and Foundry [16] routers specify this feature too. We use this as described below.

The ISP does not modify its routing setup – the ISP’s routers participate in an intra-domain routing protocol that establishes internal routes through which the routers can reach other while BGP is used for inter-domain routing just as today. For each virtual prefix, the ISP designates some number of routers to serve as aggregation points for the prefix and hence, form a virtual network. Each router is configured to only load prefixes belonging to the virtual prefixes it is aggregating into its FIB while suppressing all other prefixes.

Given this, the ISP needs to ensure that packets to any prefix can flow through the network in spite of the fact that only a few routers have a route to the prefix. This is achieved as follows:

– *Connecting Virtual Networks*. Aggregation points for a virtual prefix originate a route to the virtual prefix that is distributed throughout the ISP’s network but not outside. Specifically, an aggregation point advertises the virtual prefix to its iBGP peers. A router that is not an aggregation point for the virtual prefix would choose the route advertised by the aggregation point closest to it and hence, forward

packets destined to any prefix in the virtual prefix to this aggregation point.²

– *Sending packets to external routers.* When a router receives a packet destined to a prefix in a virtual prefix it is aggregating, it can look up its FIB to determine the route for the packet. However, such a packet cannot be forwarded in the normal hop-by-hop fashion since a router that is not an aggregation point for the virtual prefix in question might forward the packet back to the aggregation point, resulting in a loop. Hence, the packet must be tunneled from the aggregation point to the external router that advertised the prefix. While the ISP can probably choose from many tunneling technologies, the description in the rest of this paper assumes the use of MPLS Label Switched Paths (LSPs) for such tunnels.

However, an LSP from the aggregation point to an external router would require cooperation from the neighboring ISP. To avoid this, every edge router of the ISP initiates a LSP for every external router it is connected to. Thus, all the ISP routers need to maintain LSP mappings equal to the number of external routers connected to the ISP, a number much smaller than the routes in the DFZ routing table. Note that even though the tunnel endpoint is the external router, the edge router can be configured to strip the MPLS label from the data packets before forwarding them onto the external router. This, in turn, has two implications. First, external routers don't need to be aware of the adoption of ViAggre by the ISP. Second, even the edge router does not need a FIB entry for the destination prefix, instead it chooses the external router to forward the packets to based on the MPLS label of the packet. The behavior of the edge router here is similar to the penultimate hop in a VPN scenario and is achieved through standard configuration.

We now use a concrete example to illustrate the flow of packets through an ISP network that is using ViAggre. Figure 1 shows the relevant routers. The ISP is using /7s as virtual prefixes and router C is an aggregation point for one such virtual prefix 4.0.0.0/7. Edge router D initiates a LSP to external router E with label l and hence, the ISP's routers can get to E through MPLS tunneling. The figure shows the path of a packet destined to prefix 4.0.0.0/24, which is encompassed by 4.0.0.0/7, through the ISP's network. The path from the ingress router B to the external router E comprises of three segments:

1. VP-routed: Ingress router B is not an aggregation

²All other attributes for the routes to a virtual prefix are the same and hence, the decision is based on the IGP metric to the aggregation points. Hence, "closest" means closest in terms of IGP metric.

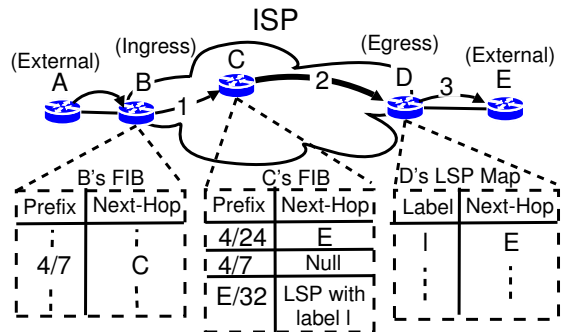


Figure 1: Path of packets destined to prefix 4.0.0.0/24 (or, 4/24) between external routers A and E through an ISP with ViAggre. Router C is an aggregation point for virtual prefix 4.0.0.0/7 (or, 4/7).

point for 4.0.0.0/7 and hence, forwards the packet to aggregation point C.

2. MPLS-LSP: Router C, being an aggregation point for 4.0.0.0/7, has a route for 4.0.0.0/24 with NEXT-HOP set to E. Further, the path to router E involves tunneling the packet with MPLS label l .
3. Map-routed: On receiving the tunneled packet from router C, egress router D looks up its MPLS label map and forwards the packet to external router E after stripping off the MPLS header.

The description above suggests that all of the ISP's traffic would need to be routed through some aggregation point. However, several past studies from as early as 1999 have shown that a large majority of Internet traffic is destined to a very small fraction of prefixes [17–20]. Consequently, routes to these *popular prefixes* will be maintained by all routers so that ViAggre's impact on the ISP's traffic is minimal.

2.3 Design-II

The second design offloads the task of maintaining the full RIB to devices that are off the data path. ISPs commonly use route-reflectors for scalable internal distribution of BGP prefixes and we require only these route-reflectors to maintain the full RIB. For ease of exposition, we assume that the ISP is already using per-PoP route reflectors that are off the data path, a common deployment model.

In the proposed design, the external routers connected to a PoP are made to peer with the PoP's route-reflector.³ This is necessary since the external peer may be advertising the entire DFZ routing table and all these routes obviously cannot reside on any given router. The route-reflector also has

³Note that these will be eBGP multihop peerings since the route-reflector is not directly connected to the external routers.

iBGP peerings with other route-reflectors and with the routers in its PoP. Egress filters are used on the route-reflector’s peerings with the PoP’s routers to ensure that a router only gets routes for the prefixes it is aggregating. This shrinks both the RIB and the FIB on the routers. The data-plane operation and hence, the path of packets through the ISP’s network remains the same as with the previous design.

2.4 Design Comparison

As far as the configuration is concerned, configuring suppression of routes on individual routers in design-I is comparable, at least in terms of complexity, to configuring egress filters on the route-reflectors. In both cases, the configuration can be achieved through a BGP route-map; in design-I, the route-map is applied at individual routers while in design-II, it is applied to the iBGP peerings of the route-reflectors.

Design-II, apart from shrinking the RIB on the routers, does not require the route suppression feature on routers. However, it does require the ISP’s eBGP peerings to be reconfigured which could represent a substantial overhead. It may also seem that the second design impacts the ISP’s robustness since the failure of a route-reflector in a PoP would severely impact the PoP’s routers. However, this is not qualitatively any different from the use of route-reflectors today and is typically accounted for by using redundant route-reflectors.

3 ViAggre Impact

ViAggre causes packets to take paths longer than native paths. Apart from the stretch imposed on traffic, this leads to extra load on the ISP’s routers and links. In the first part of this section, we study how an ISP may choose the aggregation points for its virtual prefixes so as to shrink the FIB on its routers while constraining traffic stretch. We comment on the load increase issue in section 3.3.

3.1 Assigning Aggregation Points

Ideally, an ISP would like to deploy an aggregation point for all virtual prefixes in each of its PoPs such that for every virtual prefix, a router chooses the aggregation point in the same PoP and hence, the stretch imposed on the ISP’s traffic is minimal. However, this is often not possible in practice. This is because ISPs, including tier-1 ISPs, often have some small PoPs with just a few routers and therefore there may not be enough cumulative FIB space in the PoP to hold all the actual prefixes.

Hence, the ISP needs to be smart about the way it designates routers to aggregate virtual prefixes. To this effect, we have implemented a very simple tool that uses an ISP’s topology and information about router memory constraints to determine an




Figure 2: Variation of FIB size and stretch with the constraint on the worst-case stretch.

assignment of aggregation points to the ISP routers. This tool lets us explore the trade-off between traffic stretch and FIB size offered by ViAggre. Specifically, the parameters of interest here include the *Worst-case FIB size* which refers to the largest FIB across the ISP’s routers and the *Worst-case stretch* which refers to the maximum stretch imposed across traffic to all destination prefixes from all PoPs. We also define *Average-case stretch* as the average of the stretch imposed on traffic across all PoPs. The tool uses a greedy algorithm to assign the ISP’s routers to aggregate virtual prefixes so as to minimise the worst-case FIB size while ensuring that the worst-case stretch is within a specified bound. While trivial, such a constraint would probably be critical for a practical deployment so that the ISP can ensure that its existing SLAs with managed Internet customers are not breached due to ViAggre. In the interest of brevity, we don’t discuss the details of our algorithm here; however, below we discuss the application of this tool.

3.2 Tier-1 ISP study

We used the router-level topology and BGP routing tables of a tier-1 ISP to determine the impact of the ISP adopting ViAggre. Instead of using virtual prefixes of the same length, we programmatically selected the virtual prefixes such that the distribution of real prefixes across them is relatively uniform. This led to a total of 1024 virtual prefixes that are in the FIB of every router.

We then used the aforementioned algorithm to determine an assignment of aggregation points that minimizes the worst-case FIB size given a constraint on the worst-case stretch. Figure 2 shows the (average and worst-case) FIB size and stretch for different constraints. As expected, the worst-case FIB size reduces as the stretch constraint is relaxed. For the ISP being studied, ViAggre can yield a more than 20x reduction in FIB size while ensuring that the worst-case stretch is less than 4 msec and the average stretch is less than 0.2 msec. Note that choosing the virtual prefixes such that the distribution of actual prefixes across them is not skewed provides the algorithm with greater flexibility in choosing ag-

gregation points. For instance, simply using /7s as virtual prefixes yields a reduction of $\approx 12x$ with the same 4 msec constraint.

3.3 Router Load

A naïve ViAggre deployment can cause a significant increase in traffic load across the ISP’s routers and links, not to mention the resulting interference with the ISP’s traffic engineering. For instance, for the ISP discussed in section 3.2, calculations using the ISP’s traffic matrix yielded that a deployment with worst-case stretch constrained to 4 msec would reduce the FIB size by more than 20x but would also cause a median increase in router load by 31.3%.

As mentioned earlier, the ISP can alleviate the load concern by taking advantage of the skewed distribution of traffic across Internet prefixes, which also holds for the ISP we studied. For instance, we found that 5% of the most popular prefixes were carrying 96.7% of the ISP’s traffic. Hence, the ISP can maintain routes to these popular prefixes on all its routers to greatly reduce both the load increase and the amount of traffic that gets stretched due to ViAggre. While we don’t present the details of our load analysis, considering 5% of the prefixes to be popular would drop the median and the worst-case load increase across the routers to less than 1% of the router’s native load.

4 Related Work

A number of efforts have tried to directly tackle the routing scalability problem through clean-slate designs. One set of approaches try to reduce routing table size by dividing edge networks and ISPs into separate address spaces [3,7–9,13]. Alternatively, it is possible to encode location information into IP addresses [10–12] and hence, reduce routing table size. Finally, an interesting set of approaches that trade-off stretch for routing table size are *Compact Routing* algorithms; see [21] for a survey of the area.

The use of tunnels has long been proposed as a routing scaling mechanism. VPN technologies such as BGP-MPLS VPNs [22] use tunnels to ensure that only PE routers need to keep the VPN routes. As a matter of fact, ISPs can and probably do use tunneling protocols such as MPLS and RSVP-TE to engineer a BGP-free core [23]. However, edge routers still need to keep the full FIB. With ViAggre, none of the routers on the data-path need to maintain the full FIB. A number of techniques are being used by router vendors to alleviate the impact of routing table growth, including FIB compression [23] and route caching [23]. In recent work, Kim et. al. [24] use relaying, similar to ViAggre’s use of aggregation points, to address the VPN routing scalability problem.

Over the years, several articles have documented the existing state of inter-domain routing and delineated requirements for the future [25–27]; see [26] for other routing related proposals. RCP [28] and 4D [29] argue for logical centralization of routing in ISPs to provide scalable internal route distribution and a simplified control plane respectively. We note that ViAggre fits well into these alternative routing models. As a matter of fact, the use of route-reflectors in design-II is similar in spirit to RCSs in [28] and DEs in [29].

5 Discussion and Future work

Pros. The ViAggre design presented in this paper can be *incrementally deployed* by an ISP since it does not require the cooperation of other ISPs and router vendors. What’s more, an ISP could experiment with ViAggre on a limited scale (a few virtual prefixes or a limited number of routers) to gain experience and comfort before expanding its deployment. Also, the use of ViAggre by the ISP does not restrict its routing policies and route selection. Actually, design-I does not modify the ISP’s routing setup and hence all properties such as convergence times, etc. remain the same. Finally, there is *incentive for deployment* since the ISP improves its own capability to deal with routing table growth.

Management Overhead. ViAggre imposes a significant configuration burden on the ISP. For the first design, this includes configuring route suppression on individual routers and configuring LSP advertisements on the border routers. Further, the ISP needs to make a number of deployment decisions such as choosing the virtual prefixes to use, deciding where to keep aggregation points for each virtual prefix, which prefixes to consider popular, and so on. Apart from such one-time or infrequent decisions, ViAggre may also influence very important aspects of the ISP’s day-to-day operation such as maintenance, debugging, etc.

To study this overhead, we have deployed ViAggre on the WAIL testbed [30] comprising of Cisco 7300 routers. We have already developed a tool that extracts information from existing router configuration files and other ISP databases to generate the configuration files that would be needed for ViAggre deployment. We are also developing a planning tool that would take constraints such as stretch and load constraints and other high-level goals as its input and generate ways that an ISP can deploy ViAggre so as satisfy these. While these tools are specific to the routers, ISP data and other technologies in our deployment, we believe that they can buttress our argument that ViAggre offers a good trade-off between the management overhead and increased routing scalability.

Router changes. Routers can be changed to be ViAggre-aware and hence, make virtual prefixes first-class network objects. This would do away with the configuration complexity that ViAggre entails and hence, make it more palatable for an ISP. We, in cooperation with a router vendor, are exploring this option [31].

Clean-slate ViAggre. Applying the basic concept of virtual networks in an inter-domain setting to induce a routing hierarchy that is more aggregatable can accrue benefits beyond shrinking the router FIB. The idea here is to have virtual networks for individual virtual prefixes span domains such that even the RIB on a router only contains the prefixes it is responsible for. This would reduce both the router FIB and RIB and in general, improve routing scalability.

To summarize, preliminary results show that an ISP can use ViAggre to substantially shrink the FIB on its routers and hence, extend the lifetime of its installed router base. The ISP may have to upgrade the routers for other reasons but at least it is not driven by DFZ growth over which it has no control. While it remains to be seen whether most, if not all, of the configuration and management overhead introduced by ViAggre can be eliminated through automated tools, we believe that the simplicity of the proposal and its possible short-term impact on routing scalability suggest that is an alternative worth considering.

References

- [1] G. Huston, "BGP Reports," Dec 2007, <http://bgp.potaroo.net/>.
- [2] T. Narten, "Routing and Addressing Problem Statement," Internet Draft draft-narten-radir-problem-statement-01.txt, Oct 2007.
- [3] D. Massey, L. Wang, B. Zhang, and L. Zhang, "A Proposal for Scalable Internet Routing & Addressing," Internet Draft draft-wang-ietf-efit-00, Feb 2007.
- [4] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," Internet Draft draft-iab-raws-report-02.txt, Apr 2007.
- [5] T. Li, "Router Scalability and Moore's Law," Oct 2006, http://www.iab.org/about/workshops/routingandaddressing/Router_Scalability.pdf.
- [6] D. Hughes, Dec 2004, PACNOG list posting <http://mailman.apnic.net/mailling-lists/pacnog/archive/2004/12/msg00000.html>.
- [7] S. Deering, "The Map & Encap Scheme for scalable IPv4 routing with portable site prefixes," March 1996, <http://www.cs.ucla.edu/~lixia/map-n-encap.pdf>.
- [8] D. Farinacci, V. Fuller, D. Oran, and D. Meyer, "Locator/ID Separation Protocol (LISP)," Internet Draft draft-farinacci-lisp-02.txt, July 2007.
- [9] M. O'Dell, "GSE—An Alternate Addressing Architecture for IPv6," Internet Draft draft-ietf-ipngwg-gseaddr-00.txt, Feb 1997.
- [10] P. Francis, "Comparison of geographical and provier-rooted Internet addressing," *Computer Networks and ISDN Systems*, vol. 27, no. 3, 1994.
- [11] S. Deering and R. Hinden, "IPv6 Metro Addressing," Internet Draft draft-deering-ipv6-metro-addr-00.txt, Mar 1996.
- [12] T. Hain, "An IPv6 Provider-Independent Global Unicast Address Format," Internet Draft draft-hain-ipv6-PI-addr-02.txt, Sep 2002.
- [13] X. Zhang, P. Francis, J. Wang, and K. Yoshida, "Scaling Global IP Routing with the Core Router-Integrated Overlay," in *Proc. of ICNP*, 2006.
- [14] P. Verkaik, A. Broido, kc claffy, R. Gao, Y. Hyun, and R. van der Pol, "Beyond CIDR Aggregation," CAIDA, Tech. Rep. TR-2004-1, 2004.
- [15] "JunOS Route Preferences," Jul 2008, <http://www.juniper.net/techpubs/software/junos/junos60/swconfig60-routing/html/protocols-overview4.html>.
- [16] "Foundry Router Reference," Jul 2008, http://www.foundrynetworks.co.jp/services/documentation/srcli/BGP_cmds.html.
- [17] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP routing stability of popular destinations," in *Proc. of Internet Measurement Workshop*, 2002.
- [18] W. Fang and L. Peterson, "Inter-As traffic patterns and their implications," in *Proc. of Global Internet*, 1999.
- [19] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: methodology and experience," *IEEE/ACM Trans. Netw.*, vol. 9, no. 3, 2001.
- [20] N. Taft, S. Bhattacharyya, J. Jetcheva, and C. Diot, "Understanding traffic dynamics at a backbone PoP," in *Proc. of Scalability and Traffic Control and IP Networks SPIE ITCOM*, 2001.
- [21] D. Krioukov and kc claffy, "Toward Compact Inter-domain Routing," Aug 2005, <http://arxiv.org/abs/cs/0508021>.
- [22] E. Rosen and Y. Rekhter, "RFC 2547 - BGP/MPLS VPNs," Mar 1999.
- [23] J. Scudder, "Router Scaling Trends," APRICOT Meeting, 2007, http://submission.apricot.net/chatter07/slides/future_of_routing.
- [24] C. Kim, A. Gerber, C. Lund, D. Pei, and S. Sen, "Scalable VPN Routing via Relaying," in *Proc. of ACM SIGMETRICS*, 2008.
- [25] E. Davies and A. Doria, "Analysis of Inter-Domain Routing Requirements and History," Internet Draft draft-irtf-routing-history-07.txt, Jan 2008.
- [26] N. Feamster, H. Balakrishnan, and J. Rexford, "Some Foundational Problems in Interdomain Routing," in *Proc. of Workshop on Hot Topics in Networks (HotNets-III)*, 2004.
- [27] Z. M. Mao, "Routing Research Issues," in *Proc. of WIRED*, 2003.
- [28] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe, "Design and Implementation of a Routing Control Platform," in *Proc. of Symp. on Networked Systems Design and Implementation (NSDI)*, 2005.
- [29] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Meyers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4D approach to network control and management," *ACM SIGCOMM Computer Communications Review*, October 2005.
- [30] P. Barford, "Wisconsin Advanced Internet Laboratory (WAIL)," Dec 2007, <http://wail.cs.wisc.edu/>.
- [31] P. Francis, X. Xu, and H. Ballani, "FIB Suppression with Virtual Aggregation and Default Routes," Internet Draft draft-francis-idr-intra-va-01.txt, Sep 2008.