A Pluralist Approach to Interdomain Communication Security

Ioannis Avramopoulos and Jennifer Rexford
Princeton University

1 Introduction

The best way to support secure communication in the Internet is the subject of much debate. The role of secure routing, in particular, has received considerable attention. The debate has been dominated by a "purist" philosophy that advocates the ubiquitous deployment of a secure version of BGP. The purist approach seems natural, if not mandatory, since BGP is the glue that holds the disparate parts of the Internet together. Purist solutions are advocated in public forums, such as the RPSEC working group of the IETF [2] and the North American Network Operators Group [3]. In fact, the debate focuses primarily on *which* secure routing protocol should be adopted (e.g., S-BGP or soBGP) [4], rather than *whether* a single solution should prevail. In fact, the Internet policy community has also discussed the possibility that the U.S. government might mandate S-BGP deployment [1].

Although ensuring that routing-protocol messages are authorized is clearly useful, we find the purist approach discomforting, for both economic and technical reasons:

Ubiquitous deployment would require the cooperation of more than 20,000 Autonomous Systems (ASes). The large size of the group prevents market forces from driving deployment, implying the need for government regulation—an outcome that may be both hard to realize (due to the global nature of the Internet) and undesirable (since it may stifle innovation).

Smaller groups of like-minded ASes are much more likely to deploy a security solution. Market forces can drive smaller-scale deployments, either because one (presumably large) AS is willing to bear a large part of the cost, or because adoption by some ASes has a noticeable effect on other members of the group.

Groups benefit from deploying customized security solutions. No one interdomain security solution satisfies all of the security objectives, and the choice of a secure routing protocol is just one part of any solution. Different groups may want to strike different trade-offs, based on their customer requirements and deployment costs.

Instead, we argue for a "pluralist" approach that enables graceful coexistence of multiple customized solutions, deployed by smaller groups of various sizes. We envision that each group forms an *archipelago*—an overlay of *islands*, where each island is a contiguous collection of ASes.[1] Security derives from the mechanisms the group voluntarily deploys within the archipelago (e.g., a secure routing protocol), as well as mechanisms (that we collectively call the SBone) that provide a secure virtual topology for interconnecting the islands. Unlike the archipelago, which can deploy any security solution it wishes, the SBone is constrained to provide security *on top of* uncooperative, sometimes hostile, non-member ASes. We argue that this is, in fact, possible by leveraging IP-compatible mechanisms, without requiring any changes in the non-member ASes.

Overlays have been a popular research topic recently, since they enable clean-slate design without the cooperation of the underlying network. Our approach differs from this past research in two important respects:

An archipelago is an overlay of *networks*, rather than individual end hosts or servers. In traditional overlays, the participating hosts have little or no control over the ASes they connect to. In contrast, an archipelago is created by the administrators of the participating ASes. As such, the SBone nodes at island boundaries have access to the routers (and may even run directly on the routers). For example, an SBone node could switch a virtual link from one underlying path to another, in response to a failure in a non-member network. Data-plane support in the routers can substantially improve the performance and robustness of the SBone mechanisms.

The SBone connects islands through virtual links with built-in security capabilities. In traditional overlays, virtual links have limited security capabilities, if any. For example, active probes used to detect performance problems are not robust to adversaries that treat probe packets preferentially. In contrast, the SBone has mechanisms for secure availability monitoring, as well as access control, confidentiality, and integrity.

In the next section, we present our economic arguments for the pluralist approach. Next, we present a brief overview of the SBone, followed by several examples of archipelagos that provide secure interdomain communication. Then, we discuss related work before concluding the paper with a discussion of future research directions.

[1] I.e., each AS in an island is able to reach every other AS in the island through a path consisting only of ASes in the island.

2 The Economic Case for Pluralism

In this section, we analyze the formation of groups that provide communication-security goods and the role that network architecture plays in incentivizing group formation. After a brief overview of groups and goods, we argue that a *purist* solution of a ubiquitously deployed secure routing protocol prevents market forces from driving adoption. Then, we argue for a *pluralist* approach that supports customized security solutions for groups of various sizes and is consistent with market forces.

2.1 Economics of Groups and Goods

Secure interdomain communication requires collective action. Although there are techniques that an AS acting alone can use to reduce the likelihood of attacks (such as applying protective filters to routing protocol messages and data packets), these techniques are not sufficient to ensure confidentiality, integrity, and availability for interdomain communication. Symmetric encryption instead, the simplest technique to ensure confidentiality, requires bilateral cooperation to establish a security association and encrypt/decrypt the data. The formation of *groups* of ASes is, therefore, essential for interdomain communication security. The ASes in the commercial Internet are independent, rational, and payoff-maximizing entities, making it important to consider their incentives for collaborating to provide security services. This is in sharp contrast to (say) military networks where security services may be deployed by fiat.

The group's goal is to provide secure communication as a good to its members by deploying common security mechanisms. Goods are, in general, classified as (1) purely public, (2) purely private, and (3) impurely public, with different economic implications. Pure public goods are *non-rival*, i.e., consumption of the good by one member does not diminish the availability of the good to other members, and *non-excludable*, i.e., the privilege of consumption of the good is unrestricted. An example of a pure public good is public television broadcasting. In contrast, pure private goods are rival and excludable, for example, recorded music sold in music stores. Impure public goods are partially rival or partially excludable, such as cable television broadcasting.

The appropriate classification of a good depends on the group's incentive structure for production and consumption. In fact, technological innovations can transform a good from one class to another. For example, encryption changed television broadcasting from a pure public good to an impure public good. As another example, peer-to-peer file-sharing applications are rapidly transforming recorded music from a pure private good to a pure public good. The rest of this section argues that the purist view treats secure interdomain communication as a pure public good, which prevents market forces from driving adoption, whereas a pluralist approach would better match the economic incentives of smaller groups.

2.2 Purism is not Economically Viable

Ubiquitous deployment of a secure routing protocol is unappealing because it implies *non-excludability*. Consider, for example, an exclusion mechanism based on fees. The option of charging a fee to prospective customer networks for connecting them to your secure routing protocol, implies the possibility of networks that decline to pay the fee. In the absence of other sources of revenue (e.g., advertising), non-excludability leads to *market failure*, i.e., no supply of the good, or a level of provision that is grossly inefficient. This is the situation today, where no secure interdomain routing protocol is deployed, despite a pressing need for better security.

Avoiding market failure under non-excludability typically requires government intervention, such as regulation [10]. However, regulation would be a significant departure from the way ASes interconnect today—through bilateral relationships in an unregulated fashion. In fact, the Federal Communications Commission has considered whether regulation of Internet backbones is necessary, but has declined to intervene thus far [15]. Regulation of the Internet infrastructure has been debated extensively in the academic community [25] without a definitive answer emerging. Sustaining competition and fostering innovation are at the heart of the debate, but advocates and opponents of regulation disagree on the incentive structure that will best nurture them.

Providing secure interdomain communication requires innovation because ASes are reluctant to deploy existing solutions, such as secure routing protocols, for both technical reasons[2] and the economic reasons explained above. The question about the best way to foster innovation naturally arises. Government intervention through regulation is certainly one possibility. However, we argue that regulatory action to mandate the ubiquitous deployment of a secure routing protocol is unnecessary, and in fact may stifle the creation and deployment of superior alternatives. Instead, we believe it is possible for market forces to drive the deployment of security mechanisms, including the existing and novel secure routing protocols, just not based on the purist view. In the rest of this section, we advocate *pluralism* and discuss market-based incentive structures for secure communication.

[2]For example, secure routing protocols like S-BGP do not protect against colluding adversaries (where two ASes falsely claim to have a link between them) or data-plane attacks (where an adversarial router drops packets or deflects them from the advertised path).

2.3 Smaller Groups are More Effective

Whether a group will form to counteract a threat will ultimately depend on the economic incentives of individual ASes to join. Counteracting a threat incurs costs to deploy security mechanisms, including start-up costs to upgrade the network and ongoing costs to maintain the new functionality. The value of the investment will depend on a network's individual needs, which are also likely to change with time, and the degree to which other networks support the same or similar functionality.

Common interest in counteracting a threat is a necessary condition for the formation of a group, but it is by no means sufficient. Although the members of a very large group, like the complete collection of ASes in the Internet, may have no incentive to provide a good, smaller groups can be more successful. The theory of collective action [20] argues that small and medium-sized groups are more effective in providing public goods than large ones. In a small group, one large member may have sufficient incentive to provide the good by himself, essentially financing the participation of the other members. For example, a large corporation may finance the deployment of encryption devices at smaller business partners for business-to-business transactions. In a medium-sized group, a good can be provided by strategic interaction and bargaining. For example, large backbone providers may form a coalition to deploy a secure routing protocol to protect their customers.

Accommodating independent variable-sized groups would enable market forces alone to drive the provision of communication-security goods, in two main ways. First, as noted above, a small or medium-sized group may provide a pure public good based solely on alignment of incentives or bargaining. Second, *exclusion mechanisms* can be leveraged to provide goods as impurely public or purely private. For example, a coalition of networks that had deployed a secure routing protocol may charge non-member networks to use its routes. Without such exclusion mechanisms, the possibility of free-riding would be a disincentive for the deployment of security mechanisms. In fact, fees on non-members may provide an incentive for them to join the coalition, ultimately leading to wider deployment.

2.4 Custom Security Solutions Per Group

The primary benefit of the ubiquitous deployment of a secure routing protocol is that it makes it harder for an adversary to intercept remote traffic, limiting the adversary's ability to breach the confidentiality, integrity, or availability of remote communications. However, as noted earlier, secure routing protocols alone cannot protect against colluding adversaries and data-plane attacks. In fact, the advantages of secure routing protocols with respect to one set of threats might even be disadvantages with respect to another set. Though secure routing protocols protect reachability, a victim of a denial-of-service attack may, in fact, prefer (selective) *unreachability*.

An AS, or group of ASes, have many options for improving communication security. There is encryption to ensure confidentiality, authentication to ensure integrity, availability monitoring to detect communication failures and multipath routing to circumvent them, filtering and capability-based systems to block DoS-attacks, etc. Each course of action has merits and deficiencies. Because no single mechanism addresses the full gamut of threats, individual networks are likely to prefer different combinations based on their individual goals. As such, we argue that the network architecture should support the graceful coexistence of different mechanisms rather than impose a bias toward one particular solution.

As a simple example, consider two ways to protect confidentiality: a secure routing protocol (that prevents interception of remote traffic) and encryption ciphers. Because encryption ciphers can be effective based solely on bilateral negotiations, they can serve as a building block to build groups of arbitrary sizes. In contrast, secure routing protocols offer little to no support for partial deployment. Although encryption ciphers are widely deployed in protocols such as the secure sockets layer (SSL) and in virtual private networks (VPNs), the deployment of secure routing protocols has received minimal traction. In the next section, we present a new framework that can support groups of ASes in deploying security solutions that are traditionally hard to deploy.

3 Security Backbone (SBone) Framework

A group of ASes cannot successfully deploy an interdomain security solution without an effective way to handle *deployment gaps*—non-member networks that may be uncooperative or even hostile. Our architectural framework solves this problem by connecting *islands* (contiguous collections of ASes) via *a secure mesh of virtual links* (called an SBone). Security is derived from cryptographic mechanism incorporated into the virtual links and the selection of the underlying routes traversing non-member networks. After discussing the threats imposed by non-members, we describe how to construct secure virtual links using existing network mechanisms. Then, we discuss how a group can form an *archipelago* that deploys security solutions within and between islands.

3.1 Threats From Non-Member Networks

Adversaries can launch denial-of-service attacks on destinations inside the archipelago, as well as the physical links along the paths between islands. In addition, an adversarial router may lie in the data path between two islands, or launch routing-protocol attacks to intercept the

cross-island traffic; either way, the adversary can drop, snoop, modify, discard, or misdirect the packets. The goals of the attacker include the breach of confidentiality and integrity of cross-island communication, as well as the disruption of connectivity among islands. The SBone is designed to detect these attacks and limit their effectiveness, substantially increasing the resources the adversary must expend to attack the archipelago.

3.2 Secure Cross-Island Virtual Links

A secure virtual link (or "surelink") connects a *relay point* in one island to a relay point in another, over one or more non-member networks. The sending relay point encapsulates a data packet and directs it to the receiving relay point, which decapsulates the packet and forwards it to the next hop in its journey. Encryption ensures the confidentiality of the cross-island traffic. Authentication ensures integrity, and prevents denial-of-service attacks by allowing the relay point to drop packets that fail the integrity check. In essence, surelinks enhance the service model of a vanilla IP tunnel with the cryptographic capabilities of IPsec [16], which offers point-to-point cryptographic protection at the IP layer. The relay points can capitalize on existing hardware support for IPsec in commercial routers, which are typically capable of supporting up to 10 Gbps links.

Although IPsec prevents attacks on confidentiality and integrity, the adversary can still affect the *availability* of the surelink. As such, the relay points must monitor the quality of the underlying path through the non-member networks. However, conventional monitoring techniques are not sufficient, since an adversary may bias the accuracy of the measurements by treating the probe traffic preferentially. For example, an adversary in the data plane could identify active probes (based on the packet header, size, or timing) and successfully deliver these packets while dropping or delaying the other traffic. Instead, the relay points should apply passive sampling, where the data packets serve as implicit probes, using a hash function to sample the same packets at both relay points without revealing the identity of the sampled packets to the adversary [8]. Such secure availability monitoring provides the relay points with an accurate estimate of the packet loss and delay on the virtual link.

3.3 Secure Topology per Archipelago

An archipelago constructs a virtual topology, composed of multiple surelinks, to provide secure interdomain communication among a group of islands. Having a rich collection of surelinks gives the archipelago significant influence over how the traffic flows between islands, in two main ways. First, the archipelago may have multiple underlying paths from one relay point to another, and can direct the surelink traffic over paths with higher availability. Second, the archipelago may have multiple surelinks (or virtual paths that are a sequence of multiple surelinks) that can direct traffic between two islands.

Control over the underlying paths can be leveraged in several ways. The first is to *proactively prevent routing attacks* from compromising availability. This can be achieved, for example, by routing cross-island traffic over shorter underlay routes, which are less likely to be victimized than longer ones. The second is to *proactively bypass untrusted nonmember ASes*, preventing them from receiving archipelago traffic. The third is to *proactively spread traffic* over multiple paths, reducing the overall amount of traffic carried over any single path. This substantially increases the resources an adversary would have to invest to perform a targeted attack, such as a denial-of-service attack on cross-island communication. The fourth is to *reactively reroute traffic* to an alternate path upon detecting an availability problem.

The collection of surelinks in an archipelago is controlled by software processes called *control points*. The control points may run on the same routers that implement the data-plane relay points, or on servers [12] that coordinate routing within and between islands. Communication between control points in different islands is carried over surelinks to ensure confidentiality, integrity, and availability. For example, since the relay points discard packets that fail an integrity check, the control-point software itself is not vulnerable to denial-of-service attacks. In addition to selecting the underlying paths for surelinks, the control points can implement new functionality such as secure routing protocols (e.g., S-BGP).

In addition to deploying customized routing solutions, the archipelago defines its own policies for which traffic can enter, to enforce exclusion mechanisms and prevent attacks. More generally, since an AS may participate in multiple archipelagos simultaneously, each data packet must be mapped to the appropriate archipelago (or none at all). The AS may employ a variety of mechanisms, including packet classifiers that match packets on header fields (e.g., source and destination IP addresses, TCP/UDP port numbers, etc.) and incoming link. Based on the classification, the packet may be marked (e.g., the Type-of-Service bits), tagged (e.g., with a VLAN tag or MPLS label), or encapsulated (e.g., within an IP packet destined to the nearest relay point) to ensure proper treatment by the intermediate routers. The routers inside the archipelago may have separate resources, such as forwarding tables or packet queues, for each archipelago.

4 Examples of Archipelagos

In this section, we present two examples of archipelagos. In the first example, edge networks form a group driven by one member who receives the bulk of the benefit. In the second example, backbone providers form a

coalition to offer a security service to multinational customers. Both examples illustrate how a group can collaborate to deploy a Virtual Private Network (VPN) service,[3] traditionally only available within a single institution. We also show how the groups can counteract threats against availability that traditional VPNs do not.

4.1 Edge-Network Secure VPN

Consider a large corporation planning to deploy a secure network for business-to-business transactions with its suppliers. The corporation may be willing to bear the bulk of the deployment cost of a security solution to prevent costly disruptions in its supply chain. In addition, the large business can amortize the financial investment over its large base of suppliers, whereas a smaller-in-size supplier might be reluctant or unable to invest.

Deploying a dedicated network (say, using leased lines) is an unattractive option because of the significant deployment cost. Leveraging the common IP infrastructure is preferable, except that the Internet is both insecure and unreliable. Instead, the business could deploy a *secure VPN* with each supplier to protect the confidentiality and integrity of communication. However, this solution does not detect or circumvent availability problems, and also places a large administrative burden on the business to maintain many independent VPNs.

In contrast, an SBone would provide secure availability monitoring to detect failures in the primary path between the business and a supplier, and routing through alternate paths when the primary path fails. In rerouting the traffic, the business can leverage alternate paths through other suppliers by essentially reflecting the packets off their relay points. In fact, since the corporation is bearing the cost of the security service, it can require the suppliers to support these packet deflections.

4.2 Backbone-Provider Trusted VPN

Consider a company that has offices in the US and Australia, where AT&T is the Internet Service Provide (ISP) of the US branch and Telstra is the ISP of the Australian branch. Suppose the customer would like to have a *trusted VPN* for both security and predictable communication performance between the sites. For the sake of the example, assume that AT&T and Telstra do not connect directly to each other (i.e., they do not peer). Upgrading their networks to provide a multi-provider VPN service would allow the two ISPs to tap into a market of customers with a large geographic footprint.

[3]The two dominant deployment cases of a virtual private network (VPN) is a *secure* VPN deployed by edge networks and a *trusted* VPN deployed by backbone providers. In first case, traffic protection is derived from cryptography to ensure confidentiality and integrity and, in the second case, traffic protection is derived from the security of the backbone provider's network without resorting to cryptography.

However, AT&T and Telstra cannot rely solely on the security and performance of their own networks, because the traffic must cross ASes outside their control. One possibility is that the two ASes peer directly in several locations, but this may require a significant financial investment. Instead, AT&T and Telstra can connect their networks using multiple surelinks. Encryption and authentication in the surelinks can protect the confidentiality and integrity of the traffic as it traverses nonparticipant ASes. Furthermore, availability monitoring and rerouting would allow the ISPs to improve availability, even in the presence of adversaries in the intermediate ASes. Therefore, the ISPs would be able to offer stronger service-level agreements (SLA) for the VPN traffic, improving on the poor (or non-existent) SLAs typically offered today for interdomain traffic.

The revenue from the multi-provider VPN service can serve as a bootstrapping mechanism for other backbone providers to join the group, expanding the service to customers that have sites connected to other ISPs. Today, collaboration between ISPs is largely limited to providing a basic IP reachability service. However, there are promising signs that providers are willing to collaborate to support this underserved base of customers. IPsphere [19], for example, is a consortium of providers and vendors trying to facilitate business transactions through collaborative services. These frameworks, coupled with the SBone mechanisms, can be helpful in overcoming the barriers for ISPs to collaborate in providing interdomain security services.

5 Related Work

The SBone relates to previous work that uses overlays to deploy services that the underlying Internet infrastructure lacks, such as multicast [11], reliability [5, 6], and quality of service [22]. Our work is unique in focusing on *security* and on overlays that connect entire *networks* rather than end hosts or servers.

The architecture we propose in this paper complements other proposals to evolve the Internet architecture [7, 13, 21, 23] but looks at evolvability from a security standpoint. We suspect though that our economic argument in favor of pluralism can be extended in areas beyond communication security.

Hu et al. [14] present a model for the incremental deployment of a public key infrastructure for securing BGP and Chan et al. [9] study the *adoptability* of secure routing protocols, defined as how attractive a protocol is with respect to full adoption. Both incremental deployability and adoptability imply a target objective of ubiquitous deployment. In contrast, our economic argument for pluralism advocates *partial* deployment of security mechanism, including secure routing protocols, and the SBone framework is designed to support this objective.

An argument for pluralism in secure interdomain routing is also presented in [18]. The proposed framework is based on financial *dis*incentives for initiating and propagating harmful routing-protocol messages, different from our emphasis on incentives for deploying a variety of interdomain security solutions. In addition, the work in [18] does not propose a technical framework for supporting multiple customized security solutions.

6 Conclusion

In this paper, we showed that *purism*, i.e., the ubiquitous deployment of a secure version of BGP, is not incentive compatible. Instead, deployment of communication security mechanisms should be based on *pluralism*, i.e., the formation of variable-sized groups deploying custom solutions tailored to their individual needs. We also presented the SBone, a framework for mitigating the security vulnerabilities of deployment gaps, which are inevitable in a pluralist architecture.

In the future, we plan to quantify the benefits of *partial deployment* of security mechanisms such as secure routing protocols [17, 24] and denial-of-service protection systems [26]. We also plan to explore the incentives that drive small and medium-sized groups to collaborate in deploying security solutions like the ones outlined in Section 4, to identify strategies for forming coalitions.

Acknowledgments

This work was supported by HSARPA grant 1756303. We would like to thank Nick Feamster, Joan Feigenbaum, Sharon Goldberg, Wenjie Jiang, Nikitas Konstantinidis, Carmela Lutmar, and Martin Suchara for invaluable feedback on earlier versions of this paper.

References

[1] *ARIN IX Public Policy Meeting*, `http://www.arin.net/meetings/minutes/ARIN_IX/ppm_minutes.html`, April 2002.

[2] *http://www.ietf.org/html.charters/rpsec-charter.html*.

[3] *http://www.nanog.org/*.

[4] *S-BGP/soBGP Panel: What Do We Really Need and How Do We Architect a Compromise to Get It?, http://www.nanog.org/mtg-0306/sbgp.html*, June 2003.

[5] *Akamai Sureroute.* `http://www.akamai.com/dl/feature_sheets/fs_edgesuite_sureroute.pdf`.

[6] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilent overlay networks. In *Proc. of ACM Symposium on Operating System Principles*, Oct. 2001.

[7] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. *IEEE Computer*, 38(4):34–41, Apr. 2005.

[8] I. Avramopoulos, D. Syrivelis, J. Rexford, and Spyros Lalis. Secure availability monitoring using stealth probes. Technical Report TR-769-06, Princeton University Computer Science Department, October 2006.

[9] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocols. In *Proc. ACM SIGCOMM*, Sept. 2006.

[10] R. Cornes and T. Sandler. *The Theory of Externalities, Public Goods and Club Goods*. Cambridge University Press, second edition, 1996.

[11] H. Eriksson. MBONE: The multicast backbone. *Communications of the ACM*, 37(8), 1994.

[12] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe. The case for separating routing from routers. In *Proc. Future Directions in Network Architecture*, pages 5–12, August 2003.

[13] N. Feamster, L. Gao, and J. Rexford. How to lease the Internet in your spare time. *ACM SIGCOMM Computer Communication Review*, pages 61–64, January 2007.

[14] Y.-C. Hu, D. McGrew, A. Perrig, B. Weis, and D. Wendlandt. (R)Evoluationary bootstrapping of a global PKI for securing BGP. In *Proc. ACM SIGCOMM HotNets Workshop*, Nov. 2006.

[15] M. Kende. The digital handshake: Connecting Internet backbones. OPP Working Paper 32, Federal Communications Commission, Sept. 2000.

[16] S. Kent and R. Atkinson. Security architecture for the Internet protocol. RFC 2401, IETF, Nov. 1998.

[17] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, Apr. 2000.

[18] R. Mahajan. A polytheistic approach to secure interdomain routing. position statement, *Workshop on Internet Routing Evolution and Design*, http://wired2006.org/position/mahajan.pdf, Oct. 2006.

[19] T. Nolle. A new business layer for IP networks. *Business Communications Review*, Jul. 2005.

[20] M. Olson. *The Logic of Collective Action*. Harvard University Press, 1971.

[21] S. Ratnasamy, S. Shenker, and S. McCanne. Towards an evolvable Internet architecture. In *Proc. ACM SIGCOMM*, Aug. 2005.

[22] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz. OverQoS: An overlay based architecture for enhancing Internet QoS. In *Proc. Networked System Design and Implementation*, Mar. 2004.

[23] J. Touch. Dynamic Internet overlay deployment and management using the X-Bone. *Computer Networks*, 36(2-3):117–135, Jul. 2001.

[24] R. White. Securing BGP through secure origin BGP. *The Internet Protocol Journal*, 6(3), 2003.

[25] T. Wu. The broadband debate: A user's guide. *Journal of Telecommunications and High Technology Law*, 3(69), 2004.

[26] X. Yang, D. Wetherall, and T. Anderson. A DoS-limiting network architecture. In *Proc. ACM SIGCOMM*, Aug. 2005.