Tussle in Routing: Who Determines Internet Paths?

Z. Morley Mao

As pointed out by Lakshminarayana et al. [3], there is a tussle between end-users and the ISPs in wanting to control the routes on the Internet. We re-examine the tussle among senders, receivers, and ISPs in light of a few new proposals in routing architectures and propose several directions for future work.

End-users or edge networks have limited control over routes.

End-users today have little control over how traffic reaches the intended destination. Most end-users residing in edge networks just rely on their upstream provider to deliver outgoing traffic to the destination networks. Even *Multihoming* provides only limited control in the form of choosing the first hop AS among limited choices for outgoing packets. Such decision is often made based on perceived performance and network cost to select the upstream provider to use for traffic for a given destination.

Similarly, for incoming traffic, end-users today do not have much control in terms of the paths traversed by packets. Multi-homed customers can advertise different prefixes across different provider links, use AS path prepending, or advertise specific MED values across links to the same provider in an attempt to perform inbound traffic engineering.

To overcome these limitations, overlay routing [1] enables end-hosts to select desired application-layer path traversing through intermediate end-hosts. Such overlay paths, however, cannot avoid the first hop AS as determined the immediate upstream providers. Furthermore, the paths between end-hosts serving as overlay nodes are still controlled by the network.

Sender-based control: Given the existing limitations imposed by the networks, end-users have little flexibility in managing both outbound and inbound traffic. There are several proposals to increase route diversity by allowing multi-path BGP [8] and enabling source routing for end-hosts [9]. If end-hosts are exposed multiple paths, they can achieve more flexible control over their paths to arbitrary destination networks for purposes such as load-balancing traffic and gaining better fault-tolerance.

Receiver-based control: To perform receiver-based admission control or traffic management, proposals such as capability-based scheme [10] or "off by default" [2] attempt to enable receivers to decide which sender is permitted to send traffic to them. The recent work of FastPass [7] enhances the capability-based approach by imposing traffic rate-limiting along the path from specific senders.

To summarize, end-hosts today have limited control over routes for reaching external destinations and for external destinations to reach themselves. The control desired by end-hosts can be generalized as follows. (1) Sender: the end-to-end route used to reach a destination. (2) Receiver: the end-to-end route used by a sender host. Note that clearly there can be a conflict between receiver-based and sender-based control; however, some balance can be reached. In terms of detailed control, end-hosts often desire control in the form of permitting certain hosts to send traffic to them, the number of routes the end-host has access to, the ability to bypass a particular network along a path, traffic differentiation among different destinations or senders sharing the same links, the performance characteristics, e.g., fault-tolerance, delay, loss, jitter of the path. Despite potential conflict between sender-based and receiver-based control, receiver should have priority to determine the access control of which senders are allowed to send traffic to it. Ideally, this control should be implemented within the network, close to the sender. Next, we discuss the control over routes from network's perspective.

Individual networks have partial control over routes.

Support from networks is required to achieve the control in managing routes desired by end-hosts. Individual networks, however, have only partial control over routes: each network controls how traffic leaves its network and to a limited extent how traffic enters its network. Therefore, coordination among networks is needed to gain more flexible routing. Pushback [4] is one such form for the purpose of rate-limiting malicious traffic aggregates causing denial of service. We argue a more general form of coordination is required to enable more fine-grained control over traffic flows. A initial starting point is the recent proposal on dissemination of flow specification rules [5].

The current scheme of advertising only the best routes limit the flexibility in routing, as each network can only choose paths based on the best routes advertised by each of its neighbors. The recent work on multi-path BGP

called MIRO [8] enables more flexible routing by exposing multiple AS-level paths. The traditional destination-based forwarding is no longer sufficient and needs to be enhanced to allow more flexible routing, e.g., source prefixes may be used as a criteria for determining forwarding paths, or any other fields of the packet may play a role in the forwarding decisions.

To balance the tussle of the control in routing among senders, receivers, and individual networks. We propose the study of the following issues.

- **New routing paradigms**: e.g., auction-based routing where the immediate upstream providers of the end-hosts use auctions to request its neighbors to satisfy certain QoS requirements. The route with the best cost performance trade-off is adopted. This also provides incentives for networks to explore diverse paths.
- **Renting your network equipment**: Today's routing behavior is limited by existing protocols running on the routers and pre-established commercial relationships governing the routing policies enforced by the routers. Such inflexibility can be overcome by ISPs renting out some of their network equipment to allow third party to freely configure them within prescribed limits (e.g., consumed bandwidth, well-behaved traffic profiles). Users or third-parties can pay for the use of some of such equipments to construct their own desired routing schemes or perform more fine-grained and directly controllable overlay routing.
- **Anycast-based adoption of new routing schemes**: Anycast [6] has been proposed as a mechanism to adopt incrementally deployed new services. We plan to explore how anycast would allow new routes being easily adopted across networks which do not offer the new routing service.
- **Emergency backup routing schemes**: To ensure robust routing in case of attacks or unexpected failures, networks may need to *abandon* restrictive routing policies to ensure connectivity of networks. Although connectivity does not imply reachability on today's Internet, ensuring reachability becomes critical in face of unexpected failures.
- **Content-based or service-based routing**: End-users as consumers of content or users of Internet-based services mostly care about the ability to access the information with good performance guarantees. Application-layer routing schemes should advertise the content or the service along with performance related metrics to offer more flexibility to end-users. Note that both network-based performance and server-based performance are relevant. Existing service discovery protocols can be enhanced by considering actual IP-level routing.
- **Going beyond transit – content filtering as a service**: ISPs today mostly offer transit service or the ability to reach external destinations. With increasing concerns on security, networks can also offer services in content filtering, rate-limiting or blocking malicious traffic, spam or phishing filtering, etc. Traffic desiring such service can choose to go through such networks which must also offer certain performance guarantees.
- **Dynamic virtual overlays or VPNs**: Today's VPNs are mostly statically set up. For any dynamic application community requiring QoS and access control, there is no support to construct a VPN today. Enhanced network signaling to enable a closed group for communication is desirable.

References

[1] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. Resilient Overlay Networks. October 2001.

[2] Hitesh Ballani, Yatin Chawathe, Sylvia Ratnasamy, Timothy Roscoe, and Scott Shenker. Off by Default! In *Proc. Workshop on Hot Topics in Networks*, 2005.

[3] Karthik Lakshminarayana, Ion Stoica, Scott Shenker, and Jennifer Rexford. Routing as a Service. Technical Report UCB/EECS-2006-19, UC Berkeley, 2006.

[4] Ratul Mahajan, Steven M. Bellovin, Sally Floyd, John Ioannidis, Vern Paxson, and Scott Shenker. Controlling High Bandwidth Aggregates in the Network. *ACM Computer Communication Review*, 2002.

[5] P. Marques, N. Sheth, R. Raszuk, B. Greene, J. Mauch, and D. McPherson. Dissemination of flow specification rules. IETF Internet Draft: draft-marques-idr-flow-spec-03, August 2005.

[6] Sylvia Ratnasamy, Scott Shenker, and Steven McCanne. Towards an Evolvable Internet Architecture. In *Proc. ACM SIGCOMM*, 2005.

[7] Dan Wendlandt, David G. Andersen, and Adrian Perrig. FastPass: Providing First-Packet Delivery. Technical Report CMU-CyLab-06-005, CMU, 2006.

[8] Wen Xu and Jennifer Rexford. MIRO: Multi-path Interdomain ROuting. In *Proc. ACM SIGCOMM*, 2006.

[9] Xiaowei Yang and David Wetherall. Source Selectable Path Diversity via Routing Deflections. In *Proc. ACM SIGCOMM*, 2006.

[10] Xiaowei Yang, David Wetherall, and Tom Anderson. A DoS-limiting Network Architecture. In *Proc. ACM SIGCOMM*, 2005.