

Can ISPs and Overlay Networks form a Synergistic Co-existence?

Ram Keralapura^{1,2}, Nina Taft², Gianluca Iannaccone² and Chen-Nee Chuah^{1*}

¹Univ. of California, Davis ²Intel Research

1 Introduction

Overlay networks are becoming increasingly popular for their ability to provide effective and reliable service catered to specific applications [1][2][3]. For instance, this concept has been used in peer-to-peer services like SplitStream, content delivery networks like Akamai, resilient networks like RON and so on.

Overlay networks consist of a number of nodes (spanning one or many domains) that collaborate in a distributed application. One of the underlying paradigms of overlay networks is to give applications more control over routing decisions, that would otherwise be carried out solely at the IP layer. Overlay networks typically monitor multiple paths between pairs of nodes, and select the one based on its own requirements (e.g., delay, bandwidth, etc.). Allowing routing control to take place at both the application layer and the IP layer, could have profound implications on how ISPs design, maintain and run their networks. The architecture and protocols that carriers use to run their networks are based on assumptions about how their customers and traffic behave. It is possible that the approach used by overlay networks could call into question some of carriers' assumptions thus rendering it more difficult for them to achieve their goals.

In this paper, we identify some potentially problematic interactions between overlay and layer-3 networks. We raise a key question: *Can overlay networks and underlying IP networks form a synergistic co-existence?* Given the recent rise in popularity of overlay networks, we believe that now is the time to address these issues. We hypothesize that it could be problematic to have routing control in two layers, when each of the layers is unaware of key things happening in the other layer. ISPs may be unaware of which nodes are participating in an overlay and their routing strategy. Overlay networks are unaware of an ISP's topology, load balancing schemes, routing protocol timer values, etc. We believe that ISPs need to clearly understand the implications of overlay network behavior.

2 Sample Interaction Issues

Traffic Matrix (TM) Estimation: A traffic matrix specifies the traffic demand from origin nodes to destination nodes for a single domain. A TM is a critical input for many traffic engineering tasks (e.g, capacity planning, failure provisioning, etc.) and hence ISPs undergo considerable effort to estimate TMs. Many flows whose ultimate destination lies outside an ISP's domain still appear inside the TM which specifies an exit router in the domain. If this traffic belongs to an overlay network that uses its own path selection mechanism and spans multiple domains, the overlay can alter the egress router for that flow from a

* This work was partly supported by the NSF CAREER Grant No. 0238348.

particular domain. For example, consider two domains with two peering links between them. Suppose the layer-3 path from a node in the first domain to a destination node in the second domain uses the first peering link. An overlay network can decide to route to the destination node via an intermediate overlay node that causes the path taken to traverse the second peering link, thereby changing the exit node from the first domain and subsequently the TM of that domain. If this were to happen for large flows, it could affect a significant portion of the TM. If this were to happen often, it would increase the dynamic nature of the TM which might then require more frequent updates to remain accurate. **Failure Reaction:** It has been shown recently [4] that there are a large range of failure types in the Internet, some intermittent and some long-lasting, and that overall failures happen surprisingly often. As overlay networks use frequent active probing to assess the quality of their paths, they could react to failure events at a time scale either faster or similar to ISPs. If multiple overlay networks that have their nodes in a domain experiencing a failure, react to the failure closely in time, then it is easy to see that all of them might choose the same next best path, causing congestion on that path. If two overlays have similar values for their probe timeouts, they could become synchronized in their search for a new path, thereby leading to load thrashing (or traffic oscillations). All of this is due to the relationship between the failure reaction timers of a carrier's routing protocol, and the path probing timeouts of each of the multiple overlays. We believe that careful guidelines should be developed for the selection of such timeouts. Traffic oscillations in the network due to such race conditions are undesirable for ISPs that are held accountable for performance degradation.

Load Balancing: ISPs usually have a target for the distribution of load across their network; for example, they want a low average and variance of link loads network-wide. Failures trigger overlay networks to redistribute their traffic by probing a limited set of alternate paths (constrained by where other overlay nodes reside). As overlay networks lack the global knowledge of the ISP's domain, the resulting distribution of load across all links could differ significantly from what would happen if an ISP handles the load shift all by itself. In this way overlays can undermine an ISP's load balancing strategy.

3 Summary

We have identified a few critical problems that ISPs may face due to overlay networks. Using simple examples, we show that it is important to address these issues before the impact of overlay networks is detrimental to the Internet. The co-existence of multiple overlays is only like to exacerbate these problems.

References

1. D. Anderson, H. Balakrishna, M. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *SOSP*, Oct. 2001.
2. "Akamai." [Online]. Available: <http://www.akamai.com>
3. B. Zhao, L. Huang, J. Stribling, A. Joseph, and J. Kubiawicz, "Exploiting Routing Redundancy via Structured Peer-to-Peer Overlays," in *ICNP*, 2003.
4. A. Markopoulou, G. Iannaccone, S. Bhattacharya, C. N. Chuah, and C. Diot, "Characterization of Failures in an IP Backbone Network," in *INFOCOM*, 2004.