Routing in Multiple Layers: Challenges and Opportunities*

Mostafa Ammar, Srinivasan Seetharaman

Networking and Telecommunications Group, College of Computing
Georgia Institute of Technology, Atlanta, Georgia 30332
{ammar,srini}@cc.gatech.edu

1 Introduction

In the past ten years of the Internet's evolution, we have seen a tremendous increase in complexity and processing power of end-systems, leading to a higher level of collaboration and an increase in the number of application-level services. The increasing popularity of overlay networking is one manifestation of this phenomenon. Overlay networks provide an opportunity for end-systems to break free of their dependency on the stagnant native network and empowers them to overcome native network limitations.

A surge in the number of overlay applications, or in the volume of traffic generated by them can have a significant impact on the operation of the underlying native network. This is particularly true for *performance- aware* overlays that adapt their functioning and routes based on measurements or estimates of the state of the underlying native network. Such adaptation often involves the use of a routing protocol at the overlay layer that competes with the routing protocol in use within the native network.

In this position paper we consider a scenario where overlay networks are deployed on top of a native network where both network layers deploy a routing protocol. The routing protocols may or may not have the same objectives. This can be generalized to more than two network layers in cases where overlay networks can be built on top of other overlay networks[1]. Fig. 1 illustrates this general scenario. In the figure, the traffic between nodes K and Q at the $Overlay_2$ layer is determined by routing at that layer, with traffic between each node being relayed through nodes in $Overlay_1$, with the path traversed by each $Overlay_1$ link determined by native IP routing protocols.

This position statement highlights the challenges in the design and operation of *multi-layer routing* in such networks and summarizes research opportunities derived from these challenges.

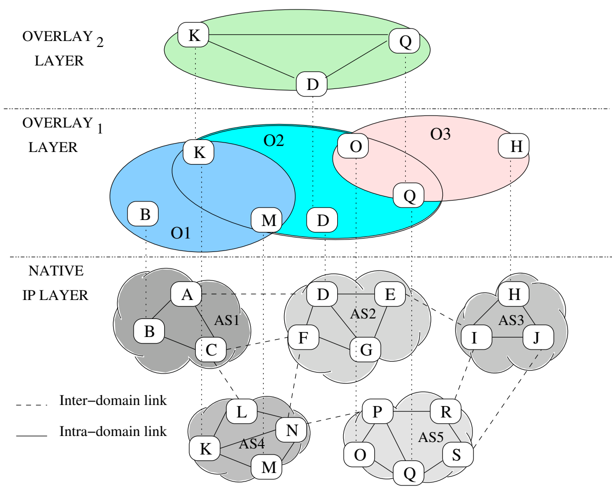

Figure 1. Multiple Overlay Layers

2 Challenges

The multiple layer routing environment presents us with the following challenges:

Functionality Overlap: In a multi-layer routing scenario, each layer often conducts most of the standard routing functions – like link verification, path computation, policy enforcement, route optimization, topology design, load balancing – causing a functionality overlap. Replicating the core functions in multiple layers can lead to route flaps and suboptimal routes (because each layer tends to undo the operations of the other layer), wasteful efforts (caused by an unawareness of the exact layer offering the best service for a particular scenario) and undesired overhead.

It is sometimes argued that overlays are merely applications and native networks have always been tuned to meet application demands. As a result overlay networks do not represent a fundamentally different paradigm. However, we argue here that while overlays can architecturally be considered applications, they are quite different from normal applications, precisely because they act like *networks* and in many respects try to mimic the functionality designed within the native network.

Coexistence and Resource Management: In the multiple-layer environment, overlays compete for native network resources with other overlays[2] (at the same or different level) and with native traffic[3]. While sharing of network resources among native flows has received significant attention in previous research, this multiple-layer environment is considerably more complex. This is especially true because of the fact that overlay layers are generally performance-aware and may deploy routing protocols that attempt to improve performance as network conditions change.

Mismatching Objective: A conflict in objective between the individual routing layers and a misalignment of objective (because of varying parameters and environment) can exacerbate the problems of network instability. For example, it has been shown that the selfish behavior of overlay routing has significant conflict with the traffic engineering efforts of the ISP[4, 5] and with the enforcement of inter-domain policies[6].

*This work was supported in part by NSF grant ANI-0240485.

[1] For instance, we can build an overlay on top of the Planetlab network[1], which is in turn an overlay on top of the native Internet.

Unconstrained Overlay Design: Real networks are constrained by the real world, e.g., population centers, natural obstacle, and national boundaries. Virtual overlay networks, on the other hand, are unconstrained, and tend to admit more flexible designs. Interestingly, the level of "freedom" increases as we go higher up in level of virtualization. While this increased design flexibility might sound like an advantage, it makes the overlay network design problem fundamentally different from the well-studied native network design problem. Additionally, the flexibility is often an illusion, causing resource deficit and performance deterioration for the overlay applications.

3 Research Opportunities

To address the research challenges discussed above, we envision a research agenda revolving around the following research opportunities:

Layer Aware Protocol Design: Much of the problems associated with multi-layer routing can be addressed by designing *layer awareness* in the routing protocols at all layers. It is easy to argue that the overlay layer should be aware of the native layer operations[7]. There is also strong evidence that the native layer itself also needs to be aware at least of the existence, but preferably the operational details, of the overlay network. In our previous work we have called such networks *overlay-friendly native networks*[8].

We believe there is sufficient incentive for native and overlay layers to increase awareness of each other and to even explicitly cooperate in their operations. It has been argued that simple modifications to the native layer can achieve substantial improvement for the overlay application[9]. Similarly, locality-awareness in P2P applications has been shown to reduce the stress on the ISP, thereby making it more willing to support P2P traffic[10].

We believe there still remains many unanswered questions about exactly how the interaction among the various routing layers should be structured. In particular, it is really critical to understand the tradeoffs introduced by such designs.

Virtual Network Assignment: As mentioned earlier, there is potentially considerable flexibility in the design of an overlay network topology. This design can in turn impact the routing performance at the overlay and its interaction with routing in other layers. We call the process of determining where an overlay will be deployed, within the resources of the underlying network, the *overlay network assignment (ONA)* problem.

We envision that certain objectives will need to be met when assigning overlay networks to underlying resources. For example, the goal of the work in [11] is to balance the load imposed on the native network. Other criteria can include the ability to meet traffic demands, to provide a certain quality of service, or to achieve a certain level of isolation among assigned overlay networks. Clearly, overlay routing and its interaction with native routing will play an important role in determining ONA solutions.

In general, ONA methods have to also address the issue of overlay network reconfiguration. We expect this to be necessary as underlying network conditions change due to changes in native traffic, changes in overlay demands, including the arrival or departure of entire overlay networks. The work in [11, 12] provide potential starting points to understanding dynamic reconfiguration at the overlay layer.

As we see it, the assignment process is concerned with providing a mapping between overlay network requirements and underlying network resources. This mapping can then be used to *instantiate* the overlay network, i.e., physically deploy the network within the underlying network resources. Recent efforts[13] have begun to address this important step.

Multi-layer Testbeds: In order for meaningful progress to be made in our understanding of multi-layer routing, we need to have a suitable testbed. The testbed should allow 1) easy creation and deletion of overlay networks, 2) control of routing protocols and their parameters at *multiple layers*, 3) the ability to create interfaces among the routing layers and among the routing protocols within the same layer, 4) the ability to measure some form of user-perceived performance. It is important for the testbed to allow the deployment of alternative scenarios easily, including a benchmark "all-native" scenario.

References

[1] "Planetlab," http://www.planet-lab.org/.

[2] R. Keralapura, N. Taft, C. N. Chuah, and G. Iannaccone, "Can ISPs take the heat from Overlay Networks?," in *Proceedings of HotNets-III*, November 2004.

[3] P. Karbhari, M. Ammar, and E. Zegura, "Optimizing End-to-End Throughput for Data Transfers on an Overlay-TCP Path," in *Proceedings of Networking 2005*, May 2005.

[4] L.Qiu, R.Y.Yang, Y.Zhang, and S. Shenker, "On Selfish Routing in Internet-Like Environments," in *Proceedings of ACM SIGCOMM*, August 2003.

[5] Y. Liu, H. Zhang, W. Gong, and D. Towsley, "On the Interaction Between Overlay Routing and Traffic Engineering," in *Proceedings of IEEE INFOCOM*, 2005.

[6] S. Seetharaman and M. Ammar, "Characterizing and Mitigating Inter-domain Policy Violations in Overlay Routes," in *To appear in Proceedings of IEEE ICNP*, November 2006.

[7] S. Seetharaman and M. Ammar, "On the Interaction between Dynamic Routing in the Overlay and Native Layers," in *Proceedings of IEEE INFOCOM*, April 2006.

[8] S. Seetharaman and M. Ammar, "Overlay-friendly Native Network: A Contradiction in Terms?," in *Proceedings of ACM HotNets-IV*, November 2005.

[9] J. Jannotti, *Network layer support for overlay networks*, Ph.D. thesis, Massachusetts Institute of Technology, 2002, Supervisor-M. Frans Kaashoek.

[10] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?," in *Proceedings of ACM Internet Measurement Conference*, Otober 2005.

[11] Y. Zhu and M. Ammar, "Algorithms for Assigning Substrate Network Resources to Virtual Network Components," in *Proceedings of IEEE INFOCOM*, 2006.

[12] J. Fan and M. Ammar, "Dynamic Reconfiguration of Service Overlay Networks," in *Proceedings of IEEE INFOCOM*, 2006.

[13] A. Bavier, N. Feamster, L. Peterson, and J.Rexford, "In VINI Veritas: Realistic and Controlled Network Experimentation," in *Proceedings of ACM SIGCOMM*, September 2006.