# Research Statement

Amogh Dhamdhere

February 2009

My research focuses on the measurement, analysis and modeling of network protocols, applications and services. I am interested in how networks perform "in the wild", after they have been designed, tested and implemented. We can obtain tremendous insight by carefully measuring the various interacting components of a network system, and the external factors that affect the performance of that system (*e.g.*, how users actually use the system). Measurements can better inform analysis and modeling, and these models in turn serve as inputs for better design. My thesis work has followed this theme, starting from measurements of the evolution of the Internet ecosystem – a complex system of interacting networks. My thesis also proposed models for the evolution of the Internet ecosystem, drawing on insights obtained from the measurements. I intend to continue this theme of measurement, modeling and analysis of networks in my future work.

Part of my research has followed an *evolutionary* approach to improve network performance. I believe that evolutionary and clean-slate research approaches are complementary to each other. Evolutionary approaches are useful when large-scale changes in the existing infrastructure are not possible. Clean-slate approaches are important because they aim to redesign network architectures and protocols from scratch, and thus focus on the fundamental issues affecting network performance. Along with these approaches, it is important to study *the evolution of the Internet*. Such a study can provide valuable inputs for both clean-slate and evolutionary design approaches. Network protocols and architectures can be designed either in a clean-slate or an evolutionary manner. Once deployed, however, they need to run on an Internet that evolves continuously due to changing traffic patterns, application popularity, and economic conditions. It is thus important to understand how the Internet evolves, to reason about the future performance of these new protocols and architectures. This is a direction that I plan to pursue in my future research.

## 1 Thesis research and current work

**Provider and peer selection in the evolving Internet**
My thesis work focuses on measurement and modeling of the evolution of the Internet ecosystem. In the first part of this work, we measure the evolution of the Internet ecosystem over the last decade. Our findings highlight some important trends in the evolution of the Internet, and hint at what the Internet is heading towards, in terms of topological and economic organization. Next, we propose a first-principles model for the interactions between ASes in the Internet ecosystem. We use the model to predict the effects of provider and peer selection strategies used by different types of networks on the topology, economics and performance of the resulting Internet. The model can also be used to reason about the provider and peer selection strategies that are likely to be profitable for different types of ISPs. This work is of practical importance for ISPs who are concerned with using the best possible provider and peer selection strategies. Further, understanding the evolution of the Internet ecosystem is important for both evolutionary and clean-slate approaches to designing protocols and architectures for the future Internet.

**The impact of router buffer size on network performance**
Packet buffers in router/switch interfaces constitute a central element of packet networks. It is well-known that the size of router buffers can significantly impact network performance in terms of link utilization, loss rate and application-level throughput. This work follows an evolutionary approach to improving network performance by appropriately sizing the buffers of Internet routers, without requiring any changes in the underlying transport protocols. Our results provide guidelines for dimensioning the router buffers to satisfy constraints on the

minimum utilization, maximum loss rate and maximum queueing delay at the bottleneck link. We show that to limit the loss rate at the link, the buffer size should be proportional to the number of long-lived flows that are bottlenecked at that link. We show that the buffer requirement for optimal application-level throughput depends on the nature of the flows (short vs. long-lived) and the nature of the flow arrival process (open-loop vs. closed-loop).

**Network troubleshooting**
The distributed nature of the Internet makes it difficult for a single service provider to troubleshoot disruptions experienced by its customers. In this work, we develop NetDiagnoser, a troubleshooting algorithm to identify the location of failures in an internetwork environment. The NetDiagnoser approach is novel in that it uses information both from end-to-end measurements as well as routing data such as routing update messages from a provider's network and Looking Glass servers. NetDiagnoser can identify various types of failures that are common in an internetwork environment, such as multiple link failures and logical failures (*e.g.*, misconfigurations of route export filters). Our results show that NetDiagnoser can successfully identify a small set of links, which almost always includes the actually failed/misconfigured links.

**Network monitoring using passive flow-level measurements**
Understanding the behavior of an intradomain routing protocol in a large network is a challenging problem. Mechanisms that monitor the control plane rely on routing messages to infer intradomain routing changes. Control plane information alone, however, cannot be used to infer when routers actually converged to a new routing state after a link or router failure, or to quantify the amount of traffic that shifted due to a routing change. In this work, we develop novel techniques to detect routing changes using passive flow-level measurements from the data plane. Our approach has several applications, such as the ability to study the effect of routing changes on the data plane, the convergence time of routers, and the network-wide changes in traffic volume and link utilization due to OSPF events.

# 2 Proposed research

**Internet evolution and its implications on the economics and performance of the future Internet**
The Internet at the interdomain level consists of multiple autonomous networks that continually rewire their connectivity to optimize a certain utility function (possibly different for each network). We can thus think of the Internet as a complex system of interacting, selfish agents. There are several questions about Internet evolution that I plan to investigate in future work. First, which provider and peer selection strategies should different types of networks (enterprise customers, content providers, transit providers) use to maximize their utility? What happens to global Internet properties such as end-to-end path lengths, when each network uses its preferred strategy? Is there a centralized optimization problem that the Internet solves through distributed optimizations? For instance, do the distributed optimizations minimize end-to-end path lengths or maximize economic efficiency? Does the Internet evolve towards a situation that is "favorable" or "unfavorable", in terms of the overall topology, economics, or performance? The answer to this question has significant implications in light of recent efforts to redesign the Internet using clean-slate approaches. This study can identify mechanisms that force the Internet to evolve towards a state that is favorable in terms of topology, economics, or performance. I plan to use the modeling framework developed in my thesis work to approach the aforementioned questions. This approach will use a combination of modeling, agent-based computation, and game-theoretic methods.

**Interdomain traffic matrix estimation**
Networks in the Internet exchange traffic according to bilateral contracts, where the payments depend on the amount of traffic exchanged. The amount of traffic exchanged by two networks depends on the interdomain traffic matrix, *i.e.*, the traffic that flows from each source network to each destination network on the Internet. The interdomain traffic matrix is thus a crucial input for Internet economics. We do not, however, have a good understanding of the properties of this traffic matrix and how it evolves over time. In future work, I plan to develop techniques to estimate the interdomain traffic matrix, using data collected from the core of the network, *e.g.*, from a large tier-1 ISP. I propose to study the properties of this traffic matrix, such as the presence of "communities" (meaning networks that exchange large traffic volumes among themselves), heavy hitter sources

or sinks of traffic, and locality effects with respect to geography or network distance. I also propose to study how the interdomain traffic matrix evolves over time. For example, has the traffic matrix changed from a predominantly client-server model to a peer-to-peer model? I plan to follow a measurement-based approach to answer these questions. The approach will rely on passive measurements from a large ISP, and measurements of the Internet's interdomain topology.

**Network monitoring and troubleshooting for ISP networks**

Large ISP networks run instances of multiple protocols, each of which provides a different functionality. For example, large ISPs typically have to deal with intradomain (OSPF) and interdomain (BGP) routing protocols, MPLS, and customer VPNs. These protocols often interact in subtle and unexpected ways. This makes troubleshooting network problems challenging, as the problem could be in any of the various protocols and system components. For example, an unreachability issue could be due to a link/router failure or protocol misconfiguration in the local network, routing loops caused by late-converging routers, or due to failures outside the local network. This behavior cannot be observed by studying a single protocol in isolation. In future work, I plan to develop a scalable measurement and monitoring system that focuses on understanding the *interactions between various protocols and system components*. This system will involve a combination of active and passive measurement techniques. For instance, active measurements could be collected using probes from end systems (thus utilizing the diversity of vantage points), or by instrumenting dedicated probing nodes. Passive measurements are already collected in most networks in the form of Netflow records or logs of protocol messages. An important component of this system is a module that will monitor and correlate alarms from these different measurement sources. I plan to use the insights obtained from these measurements to design better algorithms for troubleshooting network disruptions.

**Social networks as evolving graphs**

In my thesis work, I studied the Internet at the interdomain level as an evolving graph. Networks are seen in various other forms such as social networking websites, video and photo sharing applications, blogs and forums. The interactions among participants of a social network can be viewed as links in a graph. These graphs evolve over time, as people join social networks, interact with their friends and form new relations. There are several questions about the evolution of social network graphs. How do these graphs evolve as participants join and add links? Do we see the presence of "attractors", meaning persons that attract a large number of connections? What are the semantics of these links? All links between participants of a social network are not the same. People add links to their close friends, colleagues, and mere acquaintances on social networks. Can we infer the semantics associated with these links based on the interactions between participants on the social network? How do the evolution of the underlying graph and the semantics of the links affect the propagation of information on these networks? In this project, I will follow an approach that starts from measurements and datasets collected from social networking applications. I will use these measurements to answer questions about the evolution and semantics of the links in these graphs. I propose to use the results obtained from these measurements to develop models and algorithms for routing and information dissemination in social networks.