

Rethinking Internet Routing

Jennifer Rexford
Dept. of Computer Science
Princeton University
jrex@cs.princeton.edu

ABSTRACT

Internet routing introduces many interesting challenges, far beyond the basic problem of computing paths on a graph. This talk presents an overview of several open research questions in Internet routing, with the broader goal of placing the design of future routing architectures on a stronger theoretical foundation.

Categories and Subject Descriptors: C.2.1 Network Architecture and Design

General Terms: Algorithms, Design, Theory

Keywords: Routing, Protocols, Internet

1. INTRODUCTION

At a high level, routing is simply the matter of computing paths on a graph. Yet, the Internet has many properties and requirements that make routing immensely challenging:

- **Local policies:** The nodes in the graph are controlled by different Autonomous Systems (ASes) that have their own (possibly conflicting) ideas about which paths are good.
- **Greedy and adversarial nodes:** Some nodes may intentionally try to manipulate the system to their own advantage, or to harm others.
- **Scalability:** The graph is very large, consisting of tens of thousands of ASes, millions of nodes, and billions of end hosts directing traffic through these nodes.
- **Heterogeneity:** The links may vary widely in their capacities and propagation delays, and the nodes in their computation and storage resources.
- **Geographic distribution:** The nodes are distributed all over the planet, and beyond.
- **Topological churn:** The topology changes over time due to equipment failure (and recovery), planned maintenance, and long-term growth. Multiple nodes and links may fail simultaneously, due to shared risks.
- **Fast reaction to change:** Most applications are intolerant of long disruptions in the delivery of data.

- **Diverse traffic:** The network may support a wide range of applications with different performance goals (e.g., low delay vs. high throughput).

These challenges are exacerbated by the incremental way the Internet routing system has evolved, and the practical difficulty of deploying fundamentally new designs. Yet, even ignoring the challenge of incremental deployability, we do *not* know what kind of routing system we would want to have. And, our ability to analyze the existing system is still quite limited. In fact, perhaps the most important goal for a future Internet routing architecture is the ability to define, model, and analyze it precisely, so we *can* make strong statements about its properties. Theoretical computer science can, and should, play a major role in placing the design of routing architectures on a strong intellectual foundation.

To expand on these points, the talk will focus on several broad challenges in Internet routing, including the following:

Reconciling interdomain routing policies: An Internet is, by definition, a “network of networks,” where each independently-administered “Autonomous System” has its own preferences for which routes to use. The responsibility for reconciling conflicting preferences falls to interdomain routing, realized today by the Border Gateway Protocol (BGP). Despite initial progress in understanding when the routing system provably converges to a unique and stable solution, many important questions remain about the trade-offs between routing convergence and flexible policies. These questions become increasingly interesting in the context of multipath routing and greedy/adversarial ASes.

Defining “Autonomous System”: Although most research on interdomain routing treats an AS as a single node, an AS is itself a distributed network consisting of multiple nodes and links. In a multi-router setting, different routers in the same AS may select different interdomain routes; in fact, scalability concerns may constrain whether every router can even *know* about all of the candidate routes. In addition, an AS that connects to a neighboring network in several locations may have some expectation of “consistency” across the routes available at each place. The precise definition of an AS, and the implications for how routing information is propagated within an AS, is not well understood.

Monitoring end-to-end path quality: End hosts and edge networks need to measure the quality of paths through the Internet, to drive routing decisions and hold their service provider’s accountable for poor performance. Computing an accurate and timely estimate of path quality is challenging, not only because the performance may vary over time or be affected by the measurement process itself. In

a competitive environment, nodes along the path may have an incentive to bias the measurement results, to prevent end hosts and edge networks from detecting violations of service-level agreements or moving traffic to paths through other providers. Collecting accurate measurements of path quality, even in the presence of greedy or adversarial nodes, is an interesting and important problem, as is understanding fundamental limits on how well we can localize performance problems to the offending nodes or links.

Distributed traffic management: Routing protocols are just one part of a larger question about how much traffic should flow along each path through the network. In

today's Internet, traffic management has three players: end hosts (that adapt their sending rates in response to congestion), routers (that run routing protocols to compute paths through the network), and operators (who adjust the configuration of the routing protocols). A fresh rethinking of traffic management would revisit the definition and placement of functions necessary to achieve some precisely-stated goal. Broad questions remain about the appropriate "division of labor" between end hosts and routers, whether and how to support multipath routing, and how to adapt dynamically to changing conditions in a stable and efficient fashion.