A Quantitative Evaluation of Traffic-Aware Routing Strategies*

Eric J. Anderson
Thomas E. Anderson, Steven D. Gribble, Anna R. Karlin, Stefan Savage
Department of Computer Science and Engineering
University of Washington, Seattle, WA 98195-2350, USA

{eric,tom,gribble,karlin}@cs.washington.edu, savage@cs.ucsd.edu

Summary

In this research, we address a simple question: how much benefit can be achieved by traffic-aware routing of Internet traffic? Efficient routing on packet-switched networks has attracted considerable research attention from the early days of the Internet to the present day, yet current routing practice still relies on weighted shortest paths to route traffic, using algorithms that do not take the distribution of traffic demand into account. Compared to traffic-aware routing, static routing potentially reduces Internet performance and/or increases operational and infrastructure costs, as networks must overprovision to avoid congestion.

Efficiently mapping traffic demands onto a fixed and limited network topology – a practice called "traffic engineering" – poses several challenges: acquiring a traffic demand matrix, providing a mechanism (such as MPLS or adjustable edge weights) for mapping traffic to links, and developing an algorithm that uses the input traffic demand matrix to compute an overall routing that satisfies some network-wide constraint. Each of these issues is of active research interest today; in our research, we focus on the last problem.

Several algorithms for efficient routing of known point-to-point traffic requirements have previously been developed and extensively analyzed by the theory research community. These research efforts are based on formal models of network routing, but have not yet had much impact on Internet routing practice. One reason, we believe, is the need for design and evaluation of practical versions of these traffic-aware algorithms in the context of Internet-like topologies and workloads. In our work, we consider two well-known theoretically optimal traffic-aware policies, one based on minimizing the maximum link utilization and the other on minimizing queueing delay. We have developed two novel and more practical variants of these algorithms. One algorithm simultaneously "bi-optimizes" for maximum utilization and path length, by treating the maximum link utilization as a constraint in a second optimization step. Our other algorithm produces an equal-cost multipath routing, by rounding off routing table entries derived from optimal (Gallager) routing.

We compare these algorithms to static weighted shortest path routing across a range of topologies and workloads. An example of this comparison is presented in Figure 1. Using a topology based on the ATT-US backbone, and a point-to-point workload matrix derived from the topology (demand proportional to connectivity), we analyze the average network delay across a range of overall traffic volumes, from near-zero utilization to the maximum utilization that can be supported by the network. We have found that the traffic-aware policies significantly outperform the static policies on the topologies and workloads that we studied, and further, that the new, practical traffic-aware routing algorithms offer near-optimal performance on these topologies and workloads.

In this analysis, as a proxy for end-to-end system performance we use a performance metric based on a formal model of "average network delay", which we compute from traffic averages rather than measure through direct packet-level simulation. This approach has enabled us to examine a larger combination of algorithms, topologies, and workloads. However, the formal model does not incorporate the regulation of arrival rates through end-to-end congestion control mechanisms, assuming instead that any necessary buffering of incoming traffic takes place inside the network at router queues. Whether the conclusions we reach based on the formal model hold true under a more detailed and realistic packet-level simulation that includes congestion control is an important question, and is the immediate focus of our current research.

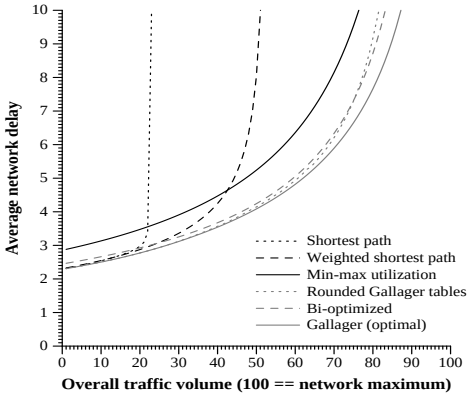

Figure 1: Comparison of average network delay for six routing strategies. Topology based on ATT-US, workload inferred from connectivity. Data points are displayed for volume levels from one percent to one hundred percent of network maximum. The Gallager algorithm optimizes for this metric.

*Presented at Sigcomm 2001, student poster session. More information on this project can be found at `www.cs.washington.edu/homes/eric/Research.html`.