

Towards cost-aware multipath routing

João Taveira Araújo, Miguel Rio, and George Pavlou

Department of Electronic and Electrical Engineering
University College London
{j.araujo,m.rio,g.pavlou}@ee.ucl.ac.uk

Abstract. Traditional approaches to multipath routing ignore the economic incentives necessary in aligning both networks and users towards a common goal. While theory suggests congestion pricing can be used to maximize social welfare, even within such a framework we have no consistent method for evaluating, constructing and disseminating paths. Our research focuses on building on existing congestion pricing models to provide scalable and differentiated edge selectable routes. By redesigning traffic management to not only provide paths, but also price them accordingly, the resulting cost-aware multipath architecture could alleviate the tension between different stakeholders by evolving the Internet towards a fully functional market.

1 Introduction

The ongoing feud between network operators and bandwidth-heavy applications has shown current resource allocation mechanisms are remarkably inept at mediating between both parties. Users expect to fully utilize the network resource pool, irrespective of the harm they cause others. With no form of protecting sporadic users from this onslaught of traffic, operators are forced to penalize heavy users through the deployment of ad-hoc solutions such as the use of deep packet inspection or degrading performance artificially. Both solutions carry their own problems - they lead to an inefficient use of resources and hinder innovation at the edges - and neither is strictly necessary.

A potential workaround for this resource allocation dilemma was presented by Kelly in his seminal work on network utility maximization [1]. Using an optimization-based model, Kelly proved social welfare could be maximized if sources are charged a shadow price proportionate to the congestion they cause. Initial interest in applying these results however was tepid. On first inspection, the Internet's architecture did not lend itself to such fine grained charging models. Additionally, many feared customer hostility toward price discrimination which has historically affected some transport networks [2]. Recent research however has re-kindled interest in drawing inspiration from economic models, such as Kelly's, to solve some of the Internet's inherent flaws. A recent protocol, re-ECN [3], has not only demonstrated that congestion charging can co-exist with flat priced subscription models, but also that enforcing cost fairness can be implemented scalably at the expense of edge policing [4].

The most significant effect of these results has been to convince the congestion control community to slowly abandon TCP-friendliness and adopt a stronger paradigm [5]. A network with more active participation in congestion control, with the accurate and truthful reflection of costs to customers, opens up the possibility of exploring service models which have thus far remaining largely undeployed due to incentive issues. One such service is multipath routing.

Exploring path diversity

As application needs diversify, the Internet must become increasingly versatile in meeting demands which are at times at odds with each other. Traditional research in Quality of Service has favoured the view that networks should offer specialized traffic classes, each optimized for specific constraints. A radically different approach is for the network to become more agnostic and allow users to explore path diversity and load balance their traffic across multiple paths.

The benefits of multipath routing are well understood, allowing applications to achieve greater end-to-end reliability, higher throughput and optimize traffic according to requirements. Despite various proposals over the past decade (e.g. [6][7]), multipath routing has been notoriously difficult to implement in practice. The simplest path to deployment would rely on the use of source routing. Unfortunately, source routing has always been at odds with traffic engineering - users are empowered to bypass network policy or engineering decisions undermining their effectiveness. With subtle pricing schemes however the balance of power is shifted towards network operators who are able to charge user traffic according to congestion caused. This market will approach equilibrium so long as the cost a network applies is inferior to the utility a user extracts for a particular path.

In the absence of congestion charging, networks have no incentive to provide additional paths. Under such conditions, multipath routing may still be deployed through the use of multihoming, as proposed in [8], requiring no network assistance. While easily deployable, this may further aggravate network operators. Direct competition between providers would be a positive development, but limiting operators to offering a one-size-fits-all path may further accelerate commoditization - effectively a race to the bottom.

As such, we believe congestion pricing should play an integral role in evaluating and constructing multiple paths, with differing characteristics, to be offered by the same provider. A shift towards cost fairness, coupled with the widespread use of multiple paths, could force traffic management to evolve from a tool predominantly responsible for minimizing costs to a valuable asset in generating revenue.

2 Building a cost-aware multipath architecture

The confluence of congestion pricing and edge selectable routes provides an opportunity to rethink the role of traffic management in a network which would behave itself increasingly as a market. Our research focuses on evolving traffic

management to offer multiple paths to users, thereby allowing the same operator to tailor paths to different uses. What distinguishes this multipath solution from related work is the use of a pricing framework to overcome incentive issues and ensure users only use what they are willing to pay for. Under these assumptions, our specific goal is to answer each of the following questions:

- *What pricing strategies must be in place to accurately reflect cost?*

The use of ECN markings as a cost indicator may suggest congestion is the only metric networks should charge users for. However, it is feasible that networks may induce markings to internalize other costs, such as providing low delay. We intend to evaluate whether a coarse accountability framework like re-ECN provides sufficient flexibility for such fine-grained pricing. In particular, we wish to understand whether ECN markings alone are sufficient to accurately convey non-congestion costs, or whether we must extend re-ECN to allow greater flexibility.

An additional concern is how traffic engineering can translate policies through the use of pricing. With the possible advent of user selectable routes, cost would ensure operators maintain control over their network, effectively replacing the role of link weights. How traffic engineering can set costs to ensure network resources are used efficiently and provide an adequate response to market demand is an open issue which is intrinsically related to our next research question.

- *How can we construct path segments locally to satisfy different demands?*

An open question in multipath routing is assessing how much control over path selection can be passed onto the end-user. Ideally, users would have the power to freely specify paths as deemed fit. In reality, this offers little benefit and carries the same scalability issues which plagued strict source routing. Offering users a limited array of paths may therefore strike a more appropriate balance between convenience and diversity. The use of pricing may be valuable in evaluating the relative merit of different paths. What is missing from the current network architecture is a mechanism for creating and diffusing alternative routes.

Within a domain, traffic engineering will be responsible for constructing path segments tailored for different needs. This amounts to evolving traditional offline intradomain traffic engineering, which optimizes routes based on expected traffic matrices, to a tool which calculates multiple path segments based on expected market demand. Our work evaluates the implications pricing has on current traffic engineering practices, and how these must be altered to maximize revenue over multiple paths.

- *How can we scale this solution across economic entities?*

Given a domain can construct multiple paths with differing characteristics, our focus turns to how such paths can be disseminated across economic entities. Pathlet routing [9], where path fragments are propagated throughout the network and assembled by sources to provide an end-to-end route, could provide

a starting point for developing a multipath routing architecture which is able to span the entire Internet. By extending pathlet routing to take into account average congestion, as provided by re-ECN, sources could concatenate path fragments and have an accurate estimate of congestion cost before attempting to balance traffic between links. This would represent a tangible improvement on re-ECN, which only provides a meaningful estimate of path congestion over a long-lived flow.

3 Conclusions

Using the principles of congestion pricing, our research aims to provide operators with means of managing traffic in an increasingly distributed environment as the Internet adopts a multipath paradigm in order to fully utilize the network resource pool. By extending traffic engineering to apply economic pressure on flows and constructing paths to meet the demands of differing application requirements, we believe the work described herein will aid in evolving traffic management to meet the demands of the Future Internet.

References

1. F. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.
2. A. Odlyzko, "The evolution of price discrimination in transportation and its implications for the internet," *Review of Network Economics*, vol. 3, pp. 323–346, September 2004.
3. B. Briscoe, A. Jacquet, T. Moncaster, and A. Smith, "Re-ECN: Adding accountability for causing congestion to TCP/IP," Internet Draft draft-briscoe-tsvwg-re-ecn-tcp-06.txt, Internet Engineering Task Force, July 2008. (Work in progress).
4. A. Jacquet, B. Briscoe, and T. Moncaster, "Policing Freedom to Use the Internet Resource Pool," in *Proc Workshop on Re-Architecting the Internet (ReArch'08)*, 2008.
5. M. Mathis, "Reflections on the TCP macroscopic model," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 1, pp. 47–49, 2009.
6. X. Yang, D. Clark, and A. Berger, "NIRA: A New Inter-Domain Routing Architecture," *IEEE ACM Transactions on Networking*, Jan 2007.
7. W. Xu and J. Rexford, "MIRO: multi-path interdomain routing," in *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, (New York, NY, USA), pp. 171–182, ACM, 2006.
8. D. Wischik, M. Handley, and M. Braun, "The resource pooling principle," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 5, pp. 47–52, 2008.
9. P. Godfrey, S. Shenker, and I. Stoica, "Pathlet routing," in *Proceedings of the Seventh Workshop on Hot Topics in Networks (HotNets-VII)*, ACM SIGCOMM, 2008.