

De-Ossifying Internet Routing through Intrinsic Support for End-Network and ISP Selfishness

Aditya Akella, Shuchi Chawla, Holly Esquivel and Chitra Muthukrishnan
Dept. of Computer Sciences, UW-Madison
{akella,shuchi,esquivel,chitra}@cs.wisc.edu

ABSTRACT

We present the S4R supplemental routing system to address the constraints BGP places on ISPs and stub network alike. Technical soundness and economic viability are equal first class design requirements for S4R. In S4R, ISPs announce links connecting different parts of the Internet. ISPs can selfishly price their links to attract maximal amount of traffic. Stub networks can selfishly select paths that best meet their requirements at the lowest cost. We design a variety of practical algorithms for ISP and stub network response that strike a balance between accommodating selfishness of all participants and ensuring efficient and stable operation overall. We employ large scale simulations over realistic scenarios to show that S4R operates at a close-to-optimal state and that it encourages broad participation from stubs and ISPs.

Categories and Subject Descriptors: C.2.m [Computer Communication Networks]: Miscellaneous

General Terms: Algorithms.

Keywords: Inter-domain routing, selfishness.

1. INTRODUCTION

BGP suffers from key inflexibilities that impose constraints on both stub networks and ISPs. BGP offers stub networks exactly one policy-constrained path per destination per ISP connection, with no guarantees on performance or availability. Thus, stub networks cannot flexibly meet the requirements of key network-based applications, such as satisfying the end-to-end performance constraints of real-time video or finance applications, especially during peak traffic periods. One way to overcome this is for stub networks to enter into partial transit or paid peering contracts with multiple ISPs to support sensitive applications. Unfortunately, these contracts are binding and long-term in nature. Other finer-grained approaches, such as overlay routing and multihoming route control, are either undesirable in practice or inadequate: The flexibility offered by overlay routing has undesirable interactions with ISP policies and traffic engineering objectives. Multihoming can offer better performance than single BGP paths, but it still cannot guarantee that the stub networks' end-to-end requirements will be met.

BGP is sub-optimal for ISPs, too. ISPs have little flexibility in controlling their revenues and expanding their services to attract a larger customer base. While BGP import and export policies allow ISPs some control over their revenues, they require ISPs to rely on long-term bilateral contracts with peers and customers. ISPs can offer performance guarantees for traffic within their own domain, but are at the mercy of those they contract with once traffic exits

their own network. Moreover, there are no easy ways for an ISP to expand its customer base to stub networks located in places where the ISP has no "physical presence". Approaches based on tunneling (e.g., MIRO [7]) are inadequate because the tunnels must traverse multiple intermediate ISPs that may not offer the tunneled traffic the same level of high performance.

Some prior efforts [5, 6, 1, 3, 2] have recognized the fundamental shortcomings of routing, namely, that it is neither aligned with important emerging stub network usage scenarios nor with ISP revenue and operational goals. However, these works focus either on (some of the) underlying implementation issues or on economic/theoretical analyses. To date, no work has both described a technical solution and evaluated its viability in practice, especially from an economic standpoint. For example, approaches such as multi-provider MPLS/VPNs [1] consider the technical issues in enabling stub networks to obtain inter-domain paths meeting their requirements, but they do not consider the crucial economic issues for both ISPs and stub networks (e.g., how to price paths to maximize revenue, how to select paths with best cost-performance trade-offs etc.), which impact whether or not such mechanisms are adopted in the first place. At the other extreme, game-theoretic models [3, 2] study selfish interactions among ISPs and stubs, and show that the result can be arbitrarily bad in some network settings; however, it is not clear if these results hold in realistic scenarios.

Our paper brings together both technical as well as economic issues to develop a compelling solution to the above shortcomings. We describe the design and implementation of an economically-grounded routing system, called S4R ("shop-for-routes"). S4R is designed to supplement, not supplant BGP. S4R enables participating stubs and ISPs to behave *selfishly* in order to directly meet their local objectives. Thus, S4R offers its participants a great degree of flexibility, which fosters greater participation from them while not requiring any kind of global oversight. We evaluate S4R in a variety of realistic situations using metrics and models that are similar to those used in prior game-theoretical analyses and show that S4R is desirable for both stub networks and ISPs. We argue the S4R can be implemented using the OpenFlow platform; we present an evaluation of a preliminary OpenFlow-based prototype.

Concretely, ISPs participating in S4R announce (virtual) links connecting different locations of the Internet. ISPs have the flexibility of dynamically altering the link prices so as to control the quality of their links and, more importantly, to attract traffic and maximize their revenue. Stub networks have the flexibility to select (or shift at any time to) paths with optimal cost-performance trade-offs for the specific application at hand. A stub network will always be able to find paths that best meet its application-level requirements as long as it has the willingness to pay for it. S4R's approach to enabling selfishness of its participants directly aligns

with the selfishness models studied in prior worst-case theoretical analyses. However, we find, surprisingly, that S4R leads to robust outcomes in practice contrary to what theory suggests [3, 2].

We conduct an extensive evaluation of S4R using a variety of realistic and synthetic scenarios to answer key technical and economic questions. Our key finding is that, in all scenarios, the net performance derived by S4R’s stub networks (in both the centralized and distributed cases) is roughly 30% away from the best possible social outcome (i.e., where all ISPs are altruistic and provide globally-optimal routes). S4R can support a variety of stub use-cases (which are poorly supported today) equally effectively.

2. S4R OVERVIEW

Stub networks. In S4R, stub networks can obtain end-to-end paths between two network locations with some associated properties. We focus mainly on *performance* guarantees, but S4R can be used for other properties, such as avoiding specific ISPs, routing through intermediaries like DDoS filters and application accelerators, requiring traffic being split over a certain number of non-overlapping paths, requiring backup paths etc. In S4R, stub networks can place requests of four different types that model different likely stub requirements in practice: (1) **Diurnal predicted:** where the stub network has a fixed required bandwidth profile for traffic to a destination. (2) **Peak predicted:** where the stub requests a certain amount of bandwidth for a specific fixed period of time in the day, corresponding to a predicted peak in traffic volume. (3) **Instantaneous:** Based on some initial monitoring, a stub may decide to instantaneously purchase a certain amount of bandwidth for some time period. (4) **Elastic bulk:** This models delay-tolerant bulk transfers (e.g., prefetching VOD content, transfers of large scientific data sets etc.).

S4R stubs provide a *value* associated with the specific traffic to a destination, which is treated as private information. Stub networks are *local utility-maximizing*: a stub network can select routes such that its *utility*—the difference between the value derived by the stub network and the price it pays for the routes—is maximized on a per-destination and per-application basis.

ISPs: Each ISP offers to carry traffic across a “virtual link” between two network locations (e.g., specific PoPs) at some cost per unit bandwidth. ISPs are *revenue-maximizing*, setting prices to maximize the revenue earned from the links owned. An ISP’s revenue per link is the product of the stub network flow routed on the link and the link price per unit traffic.

Equilibrium: There are two possible approaches to accommodating the objectives/requirements of, and the interactions between, ISPs and stubs: centralized and distributed. In the former, a logically central facilitator emulates the selfish interactions between ISPs and stub networks and derives a *correlated equilibrium*. In the latter, ISPs and stub networks interact constantly and organically. Detailed descriptions of these alternatives can be found in [4].

3. S4R EVALUATION

S4R is similar to a real world marketplace where customers are willing to shop around for the best prices for sets of goods and stores try to competitively price goods to attract customers to purchase from them. Since each ISP is interested in maximizing its own revenue, the overall system performance at equilibrium may not be “socially” optimal (compared to a hypothetical third-party computing a globally optimal solution for maximizing the performance). There is a rich body of recent work in algorithmic game theory that studies similar settings, providing bounds on the “price of anarchy” (POA), namely the ratio of the worst-case system performance at equilibrium to the social optimal. The works most relevant to S4R are those of Chawla and Roughgarden [3] and Chawla

and Niu [2]. These show that, for pathological network instances, the price of anarchy can be unbounded, implying that system performance can be significantly far from optimal [3, 2]. The poor efficiency means that few stub networks and ISPs are likely to extract utility from S4R and hence S4R may not be viable. The works also find that when stub values satisfy the *monotone hazard rate (MHR)* condition, the worst-case performance improves significantly: it is worse than optimal by a factor no more than exponential in the number of hops between any source and the sink, and is independent of other parameters such as the values themselves, network size, available capacities, etc. While this is somewhat “positive” for S4R, it still shows that the outcome in practice can be quite far from the optimal, which brings S4R’s viability into question.

To understand if the theoretical worst-case results hold in practice, we conducted a variety of simulation experiments that emulate different realistic scenarios. Our key metric of interest here is the social value derived by the system relative to the optimal social value, that we also refer to as “efficiency”. This measures the ability of stub networks to obtain as much benefit as possible from the system while allowing the ISPs to extract maximal revenue. This metric is different from POA because POA is a measure of efficiency of the worst Nash equilibrium.

Through our evaluation, we found that [4]:

- S4R converges to a stable operating point in all conditions we studied. The overall efficiency is between 65% and 80%, showing that S4R is viable in practice, i.e., it will be of high overall utility, contrary to what the theory predicted. S4R is efficient even when the disparity in stub values is high in practice, contrary to what theory has found.

- The distributed approach converges in all situations as well. Its efficiency is only slightly inferior to the centralized variant.

- ISPs can employ simplistic regret minimizing learning algorithms to set their prices. We find that selfish stub response for rerouting actually leads to better outcomes as it provides more up-to-date information to ISPs about the impact of their price changes.

- S4R effectively supports all the four demands models described earlier. S4R can accommodate a modest amount of churn (up to 10% change of demand due to instantaneous stubs entering and leaving).

- At equilibrium, stubs who have the highest values for their traffic always find paths, and there is no significant skew in ISPs revenues. This shows that both ISPs and stubs will find S4R attractive.

- S4R efficiency suffers when the network has limited path diversity and/or long paths. However, since the barrier to entering S4R is low for ISPs, we expect rich interconnection and path diversity.

4. REFERENCES

- [1] BGP/MPLS IP VPNs. <http://tools.ietf.org/html/draft-ietf-l3vpn-rfc2547bis-03>, 1999.
- [2] S. Chawla and F. Niu. The Price of Anarchy in Bertrand Games. In *ACM Conference on Electronic Commerce*, pages 305–314, 2009.
- [3] S. Chawla and T. Roughgarden. Bertrand competition in networks. In *Symposium on Algorithmic Game Theory*, pages 70–82, 2008.
- [4] H. Esquivel et al. RouteBazaar: An Economic Framework for Flexible Routing. Technical Report 1654, UW-Madison, Mar. 2009.
- [5] C. Estan, A. Akella, and S. Banerjee. Achieving good end-to-end service using Bill-Pay. In *ACM HotNets-V*, Irvine, CA, Dec. 2006.
- [6] V. Valancius et al. Mint: A market for internet transit. In *ReArch*, 2008.
- [7] W. Xu et al. Miro: multi-path interdomain routing. In *SIGCOMM*, 2006.