Geographically Informed Inter-Domain Routing

Ricardo Oliveira* Mohit Lad* Beichuan Zhang† Lixia Zhang*

*{rveloso,mohit,lixia}@cs.ucla.edu †bzhang@cs.arizona.edu
University of California, Los Angeles University of Arizona

***Abstract*— We propose to add *geographic location* information into BGP routing updates to enable Geographically Informed Inter-Domain Routing (GIRO). GIRO departs from previous geographical addressing proposals in that it uses geographical information to *assist* policy-based routing instead of replacing the provider-based IP address allocations. We show that, within routing policy constraints, geographic information can help routers select routing paths with shortest geographic distance and significantly improve the performance of the global Internet routing system. We evaluate GIRO's performance through simulations using a Rocketfuel-measured Internet topology. Our the results show that GIRO can reduce geographic distance for 70% of the existing BGP paths, and the reduction is more than 40% for about 20% of paths.**

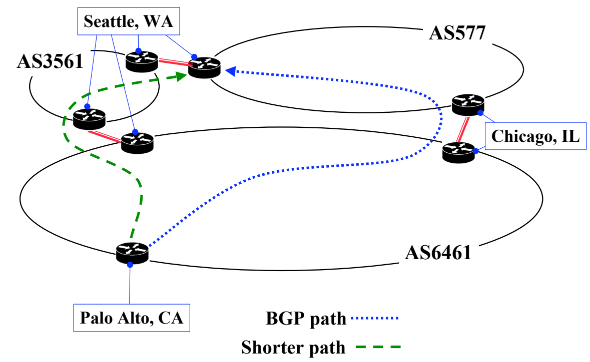

Fig. 1. The route via Chicago travels 3,600 miles, whereas the route via Seattle travels 700 miles.

I. Introduction

Today's Internet consists of a large number of autonomous systems (ASes), which exchange BGP routing messages between each other to propagate reachability information. BGP route announcements and selections are determined by networks' routing policies, which in turn are determined by the business relationship between neighboring ASes. When multiple routes to the same destination have the same policy preference, BGP breaks the tie by picking the route with the smallest AS hop count. However because routes with the smallest AS hops may not have the shortest geographic distance, data may travel longer distances which result in performance reduction.

A. Suboptimal Path Selection

Figure 1 shows an example extracted from BGP log data. AS6461 is a peer of both AS3561 and AS577 and treats the routes from them with equal preference. To reach an AS577's prefix in Seattle, WA, the AS6461's router in Palo Alto, CA, has two candidate routes: [AS6461 AS577] and [AS6461 AS3561 AS577]. BGP picks the first route, which is one AS hop shorter than the second route but five times longer in geographic distance. This may lead to higher end-to-end latency and lower TCP throughput. A previous measurement study [3] reported that about 75% of BGP paths suffer from route length inflation up to more than 15 msec, which were mainly caused by the use of AS-hop count in BGP decision process.

B. Incorporating Geographic Information into Routing Decisions

An inter-domain routing protocol must first choose routes that satisfy given routing policies. Within the policy constraints, the routing protocol should also be able to choose the routes that offer good data delivery performance. Due to the ever increasing density of AS interconnectivity, a router usually has multiple alternative paths to choose within its policy constraints. The performance of the path selection process can be measured either within an ISP (e.g., by the link metric), or end-to-end (e.g., by end-to-end delivery delay). ISPs desire good local performance that can minimize their cost in forwarding data traffic, as well as good end-to-end delivery performance to attract end users. Choosing the path with minimum AS hop count can be viewed as *an attempt to improve end-to-end performance*. Among alternative paths with the same AS hop count, BGP follows a multi-step decision process to nail down the final choice, and one important step is to choose the path with minimum IGP cost, which can be seen as *an attempt to minimize AS internal cost*.

In the current practice, however, BGP lacks the necessary information to make the best routing decisions, as demonstrated by our example mentioned earlier. We propose that BGP path selection can be significantly improved with location information. Instead of minimum AS hop count, we can choose paths with shortest end-to-end distances. Due to the rich connectivity in the Internet topology, we expect geographic distance to have good correlation with end-to-end delay. Short end-to-end delay can provide benefits to interactive, real-time applications, as well as to non-realtime applications. Traversing shorter distance (and fewer routing devices) can also reduce the chance of outage, delay jitter, congestion and packet losses.

C. Previous Efforts in Geographic Addressing

There have been several proposals on geographic location based addressing and routing in the Internet. Although these proposals differ from each other in specifics, they share the fundamental notion of allocating IP addresses solely based on geographic locations. Since location-based addresses do not reflect either the ownership of the addresses or the interconnectivity among network providers, routing based on geo-addresses not only requires that ISPs interconnect at all locations but also is unable to support routing policies. Our solution differs from all the above proposals in that we do not change the nature of the current address allocation or routing practice; we simply add to the BGP updates the location information of each AS hops to assist the routing decisions.

II. Geographically Informed Routing

A. Adding Geographic Information in Routing Announcements

We assume each BGP router is configured with its geolocation information. When an origin AS A announces a route, we define a new BGP attribute to attach geographic location information, (x_{outA}, y_{outA}), to the announcement, where x, y_{outA} represents the geographical latitude and longitude of the exit point of AS A. When this announcement propagates through ASes B and C, the AS path [A B C] goes through three ASes via ingress and egress routers of each AS. Each egress border router attaches its own location information to the BGP update message. Based on the per-hop location information, a router can estimate its distance to the destination prefix and use this information to replace the minimum AS hop count in its route selection.

B. Path Selection

In order to select the shortest geographic distance paths, we replace the default BGP hop count comparison with a geographic distance comparison.However, since distance is measured in miles, we must avoid the pitfall of letting slight difference in distance (which is a rough estimate in any case) make a big influence over routing decisions. Therefore, we introduce a new parameter δ that represents the resolution in geographic distance measurement. The value of δ is an operational parameter configured by each ISP. If δ is small, the decision process is essentially minimizing the end-to-end distance of data delivery paths. On the other hand, if δ is large, the geographical distance tie-break is less selective and the routing decision process essentially optimizes the local cost, by *e.g.* applying early-exit at a later step. Therefore, the parameter δ is a knob that allows each ISP to tune the trade-off between optimizing the global and local costs of each route.

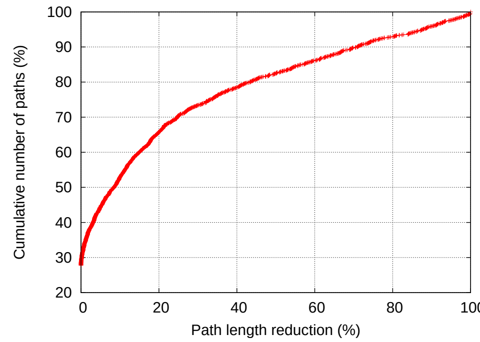

Fig. 2. GIRO path length reduction compared to BGP.

III. Preliminary Evaluation

We simulate GIRO decision process using a PoP level topology measured from Rocketfuel [4]. The topology has 668 AS level links among 67 ISPs, which are mostly tier-1 ISPs and other large ISPs directly connected to tier-1's. Relationships between different ASes are inferred using the PTE algorithm [1], and are classified into either peer-to-peer or customer-to-provider. To simulate BGP paths, we abstract the decision process into the following steps: (1) local preference based on inter-AS relationship, (2) AS hop count, and (3) random tie-breaker. For GIRO path computation, we configured $\delta = \frac{\frac{2}{3}c}{1ms} \simeq 124$ miles, where c is the speed of light in vacuum, and $\frac{2}{3}c$ is the speed of light in fiber optic[2].

Figure 2 shows the reduction of path length (measured in miles) achieved by GIRO. For over 70% of paths, the GIRO paths have shorter distance than BGP paths. For about 20% of the paths, the path length reduction is 40% or higher.

References

[1] L. Gao. On inferring Autonomous System Relationships in the Internet. In *IEEE/ACM Transactions on Networking*, volume 9, pages 733–745, 2001.

[2] R. Percacci and A. Vespignani. Scale-free behavior of the internet global performance. *European Physical Journal B 32*, pages 411–414, 2003.

[3] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *ACM SIGCOMM*, 2003.

[4] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP Topologies with Rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, 2004.