



Performance model and evaluation on geographic-based routing

Taoyu Li^{a,*}, Yuncheng Zhu^a, Ke Xu^b, Maoke Chen^c

^a Department of Electronic Engineering, Tsinghua University, Beijing 100084, China

^b Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

^c Network Research Center, Tsinghua University, Beijing 100084, China

ARTICLE INFO

Article history:

Received 20 January 2008

Received in revised form 31 October 2008

Accepted 2 November 2008

Available online 17 November 2008

Keywords:

Inter-domain routing

Geographic-based routing

Modeling

Performance

Simulation

ABSTRACT

In order to improve the inter-domain routing performance as well as its route aggregation, people have proposed to use geographical information to assist the routing system in both addressing architecture and routing mechanisms. Researches on such mechanism focus on reducing the geographical length of the path selected by routing mechanism. However, to the best knowledge of authors, there is no evaluation about how such reduction influences the actual end-to-end data delivery performance. In this paper, a simple model considering both processing and propagation delays is developed and a method to evaluate the performance of a routing mechanism is presented. Also, a heuristic algorithm to optimize routing performance is presented base on the model. After that, the performance of Geographic-Based Routing (GBR), compared to current practice using Border Gateway Protocol (BGP), is evaluated through a simulation using concrete statistical data accordingly. Result shows that GBR did not show performance improvement in most of the time, and current practice can be improved by more than 50% in performance theoretically. It is believed that this is the first work to model and examine the performance of inter-domain routing in the context of actual packet delivery behavior.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Although Border Gateway Protocol (BGP) [1] is almost the only operating instance among a number of protocols used for global routing between Autonomous Systems (AS), or inter-domain routing protocols, proposals have been raised to improve various aspects of inter-domain routing practice. Among them the performance properties of inter-domain routing has been attracting more and more attention by both researchers and network operators.

In order to improve the routing performance as well as the route aggregation, people have proposed to use geographical information to assist the routing system in both addressing architecture [2] and routing decisions [3]. People deem that Geographic-Based Routing (GBR) can result in path length reduction compared to BGP, since many paths inflate because of the use of AS-hop count as a tie-breaker metric in BGP [4].

However, to the best knowledge of authors, there is no evaluation about how those geographical path length reduction influence the actual data delivery performance. Moreover, the experience of network traffic analysis, tells people the rule of thumb that network delay is caused mainly by the network congestion rather than the propagation delay. As a result, the possible congestion resulted from the route change of GBR may counteract the amelioration in performance by reducing the geographical length.

In this paper, we intend to validate the performance of GBR through simulation methods. Rather than simply considering end-to-end propagation delay, we combine both queuing and propagation ones within total end-to-end transmission latency to formulate a model to represent the actual data delivery performance.

Since GBR is a routing mechanism still in discussion, to perform an evaluation by running GBR on actual network is impossible. To deal with this problem, we build a simple model and then take topology and traffic data from a specific network. This method makes it possible to evaluate a routing mechanism which is not yet deployable in actual networks while keeping the result closely related to the actual ones.

The remaining part of this paper is organized as follows. Section 2 introduces the geographic-based addressing architectures and geographic-based routing mechanisms. Section 3 illustrates the modeling of end-to-end transmission performance, which is in turn used in Section 4, a experiment with its result demonstrated and analyzed. The improvement of inter-domain routing aspired by the simulation result is discussed in Section 5. And finally, the paper is concluded in Section 6.

2. Background

The idea of addressing and routing based on geographical information was first proposed by Finn [5]. More proposals bearing the same basic notion have been brought forward in the passing years,

* Corresponding author. Tel.: +86 10 62603203.

E-mail address: taoyu@ns.6test.edu.cn (T. Li).

including Geographically Informed Inter-domain Routing (GIRO) by Oliveira et al. recently [3]. Most of these work involves two aspects, the addressing architecture and the inter-domain routing mechanism.

2.1. Geographic-based addressing

One important aim of geographic-based addressing is to improve the aggregation of IP addresses. As early as in 1990s, it was indicated that the current address space aggregates badly because of the provider-based allocation and the confusing intra-domain address distribution.¹ Several drafts have been proposed to bring forth a scheme of provider-independent address allocation [6].

This scheme is also reminded as geographic-based addressing, for it involving geographical information – the longitude and the latitude of the network – in the first bits of IP address. It is believed that the involvement of geographical information in IP address can help in the aggregation of IP prefixes [3].

The GIRO architecture evolved on this subject still further. It is presented that by including AS number and a traffic slice ID, an identifier to facilitate traffic engineering in multi-homed networks, as well as by including geographical location into IP address, route aggregation can be managed even better than simple geographic-based addressing [3]. Research has shown that a more than 70% reduce of numbers of BGP prefixes is expected by introducing GIRO.

2.2. Geographic-based routing

Besides addressing architecture, inter-domain routing is also believed to have close relationship with geography. Several researches have tried to analysis the geographic properties of Internet routes [4,2], among which GIRO presents a geographically informed inter-domain routing mechanism based on previous work. The fundamental notion of GIRO is that when all other policies are compliant, a router should choose a path with the shortest geographical distance, not the path with the shortest AS-hop count, as current BGP does [1].

A typical AS topology, where GBR is believed to be able to improve routing, is illustrated in Fig. 1. Supposing that AS64514 is in peering relationship with both AS64512 and AS64513 and treats routes from its two peers with same policy preference. In order to reach the Point-of-Presence (POP) of AS64512 in Beijing from the POP of AS64514 in Shanghai, two routes with same policy preference is given. BGP break the tie by picking the route with the lowest AS-hop count, so the path through Xi'an, with an AS-hop 64514–64512 is chosen. However, the path AS 64514–64513–64512, though has more AS-hops, has a less geographical length. In GIRO, geographical distance is used as a tie-breaker prior to AS-hop count, in order that those shorter paths will be chosen.

By using geographical information engaged in IP address as an important routing metrics, it is believed that GBR can result in path length reduction compared to BGP, as that its measurement study reported that in current inter-domain routing, many paths inflate because of the use of AS-hop count as a tie-breaking metric in BGP [4]. Simulation has shown that GBR design can reduce the whole routing path lengths for a mass of the paths.

However, to the best knowledge of the authors, there is no evaluation about how those geographical path length reduction influence the actual data delivery performances. Moreover, the rule of thumb come from the experience of network traffic analysis that network delay is caused mainly by the network congestion rather than the propagation delay, will results in that the congestion re-




Fig. 1. Use shorter geographic distance or less AS-hop count as the tie-breaker results in the selection of different paths.

sulted by the route changing of GBR may counteract the economize of the geographical length.

3. Modeling

The methodology, also the parameters of the performance model are presented here in the order that how these elements build a simulation model. This model can be a better reflect actual performance since it is developed in the context of actual packet delivery behavior.

3.1. Basic parameters

In order that the models well represent the performance of protocol on specific network, we established two groups of parameters, the first group decides the capacity of the network infrastructure while the second represents the behaviors of users in the network.

In the model, a POP is denoted as P_i , and a link sending data from P_i to P_j is denoted as L_{ij} . For the first group of parameters, we combine the data of actual average delay $\bar{d}(P_i)$ of each POP P_i , bandwidth capacity $B(L_{ij})$ and actual traffic load $t_0(L_{ij})$ of each link L_{ij} in every hour to generate the capacity parameter of each POP on each of its links. Also, we take $g(L_{ij})$, the geographical length of link L_{ij} , as an important parameter.

For the second group of parameters, we adopt the inbound and outbound traffic load $t_i(P_i)$ and $t_o(P_i)$ of each POP to calculate the traffic flow between every POP-to-POP pair. The inbound traffic of a POP is divided to all other POP according to the ratio of their outbound traffic. That is, $t(P_i, P_j)$, the traffic from a certain P_i to P_j is calculated as

$$t(P_i, P_j) = t_i(P_i) \frac{t_o(P_j)}{\sum_{k \neq i} t_o(P_k)} \quad (1)$$

3.2. Route selection

To evaluate the performance of BGP and GBR, the routes, between every two POP on the topology established above, are calculated, respectively. If the two POP are in the same AS, only intra-domain routing applies, one of the Interior Gateway Protocols (IGPs) is adopted and $g(L_{ij})$ is chosen as the routing weight. If the two POP are in different AS, different inter-domain routing protocol applies. Noticing that both BGP and GBR carry out a same series of routing policies except the tie-breaker, the relationship between

¹ See S. Deering, Metro-Based Addressing, <ftp://parcftp.xerox.com/pub/net-research/metro-addr-slides-jul95.ps>.

AS can be simplified to be all peering with transit allowed. In such condition, GBR will choose a egress link with minimum geographical distance $g(L_{ij})$, while BGP will pick up routes with minimum AS-hop count and then early-exit policy, which is the default policy of BGP, will apply to break the ties.

After we have computed a route $R_{ij}^{(\mathcal{P})} = \{L_{ab} | L_{ab} \text{ is a link traversed by route from } i \text{ to } j \text{ by protocol } \mathcal{P}\}$, the aggregated traffic load $t^{(\mathcal{P})}(L_{ij})$ on a single link L_{ij} under protocol \mathcal{P} can be calculated as

$$t^{(\mathcal{P})}(L_{ij}) = \sum_{a,b:L_{ij} \in R_{ab}^{(\mathcal{P})}} t(P_a, P_b) \quad (2)$$

3.3. Performance criteria

Finally, a performance criterion is needed to represent the user experience under different routing protocol. We chose average end-to-end transmission delay as performance criteria because it can reflect the inflation of geographical path as well as the performance degradation when network is congested and a number of packets are dropped.

The end-to-end transmission delay can be divided onto links, on each link, the total transmission delay is actually composed by two parts: queuing delay and propagation delay, as indicated in Fig. 2. Thus, denoting queuing delay as $d_{\text{queue}}^{(\mathcal{P})}(L_{ij})$ and propagation delay as $d_{\text{prop}}^{(\mathcal{P})}(L_{ij})$, we have

$$d^{(\mathcal{P})}(L_{ij}) = d_{\text{queue}}^{(\mathcal{P})}(L_{ij}) + d_{\text{prop}}^{(\mathcal{P})}(L_{ij}) \quad (3)$$

where $d_{\text{prop}}^{(\mathcal{P})}(L_{ij}) = g(L_{ij})/v$, in which v is the average speed of light in optic fiber, which is about $\frac{2}{3}c$ [7].

To derive $d_{\text{queue}}^{(\mathcal{P})}(L_{ij})$, since the queuing delay is brought on the POP but not on the link, there is

$$d_{\text{queue}}^{(\mathcal{P})}(L_{ij}) = d_{\text{queue}}^{(\mathcal{P})}(P_i) \quad (4)$$

and we divide this delay happening on POP into two parts: one is the processing and sending time, which we assumed to be a constant for a certain POP, denoted as $d_0(P_i)$. The other part is the queuing time monotonic increasing according to the average intension of the traffic load. Consulting queuing models like M/G/1 [8], it is known that this queuing delay would have a linearly proportion with $1/(\mu - \lambda)$, where λ is the current traffic load, and μ is the processing capacity [9]. Thus we have queuing delay on POP i being

$$d_{\text{queue}}^{(\mathcal{P})}(P_i) = d_0(P_i) + \frac{k(P_i)}{\mu(P_i) - \lambda(P_i)} \quad (5)$$

in which $\lambda(P_i)$ is the current traffic load, and $d_0(P_i)$, $k(P_i)$, and $\mu(P_i)$ are undetermined coefficients for POP i , which can be decided by a least squares fitting with a group of data of $(d(P_i), \lambda(P_i))$ pairs.

Now that $d_{\text{prop}}^{(\mathcal{P})}(L_{ij})$ and $d_{\text{queue}}^{(\mathcal{P})}(L_{ij})$ have been determined, the total transmission delay of a link, $d^{(\mathcal{P})}(L_{ij})$, can be calculated. And there-

fore, we can calculate the total end-to-end transmission delay of a route by summing delay of each link traversed by the route together, as

$$d^{(\mathcal{P})}(R_{ab}^{(\mathcal{P})}) = \sum_{L_{ij} \in R_{ab}^{(\mathcal{P})}} d^{(\mathcal{P})}(L_{ij}) \quad (6)$$

Thus the average delay of the traffic applied routing protocol P can be calculated as

$$d^{(\mathcal{P})} = \frac{\sum_{i \neq j} t(P_i, P_j) d^{(\mathcal{P})}(R_{ij}^{(\mathcal{P})})}{\sum_{i \neq j} t(P_i, P_j)} \quad (7)$$

$d^{(\text{BGP})}$ and $d^{(\text{GBR})}$ are the major performance index in our experiment. The simulations of BGP and GBR shares the same algorithm for performance index introduces above, the difference of their performance thus results from their traffic goes in different routes, due to their adoption of different tie-breakers.

3.4. Optimal route estimation

Now that we have developed a model for routing performance, it is intuitive to estimate the optimal performance. Instead of looking for a clear, theoretical upper bound of performance, we develop a heuristic algorithm to approach it. Inspired by traffic-aware routing proposed by Basu et al. [10], we set up a path-vector algorithm with the performance criteria, the average transmission delay, used as the routing weight for feedback. This mechanism, since it avoids heavy congestion on minority links, can have an optimized performance. The detailed description of the algorithm is as Algorithm 1.

Algorithm 1. Optimizing algorithm running on router P_i

```

While true do
  if received_update(Route(prefix, weight, next_hop)) then
    local_update(Route);
    if Dedicated[Route.prefix] = Route then
      send_update(Route);
    end
  end
  for egress link  $L_{ij}$  do
     $n = \text{new\_delay}(L_{ij})$ ;
    if  $|n - \text{delay}[L_{ij}]| > \delta$  then
      For Route.next_hop =  $L_{ij}$  do
        local_update(Route);
        if Dedicated[Route.prefix] = Route then
          send_update(Route);
        end
      end
    end
  end
end

```




Fig. 2. The composition of total transmission delay on a link.

For a more detailed description, each router P_i get the delay of $d^{(\mathcal{P})}(L_{ij})$ and $d^{(\mathcal{P})}(L_{ij})$ on each of outbound links regularly and thus calculate a new weight based on delay index introduced above for each prefix passed through this link. Routers then compare the new weight to previous announced one. If the difference is larger than some previously configured δ value, they will send weight information update for corresponding prefixes to neighbors.

We believe that this abstract routing algorithm can result in a sub-optimal routing on each certain topology and can be used to used as a contrast for current routing mechanisms. Since the weight is strictly positive one, the convergence property of this mechanism can be assured by path-vector algorithm itself.

Although its stability and related properties are not theoretically discussed in this paper, our simulation does show that such an algorithm could converge in short period of time.

4. Experiment

After establishing the performance model developed in Section 3, we can run simulation of different routing mechanisms, to compare the difference among the performance of GBR, BGP, and the theoretical sub-optimal value. In this section, we first present the topology and data used in our experiment, and then demonstrate and analyze the experiment result.

4.1. Topology and data

AS-level topology of Internet could be inferred from several monitoring sites dumping BGP routing tables from routers in the tier-1 AS.² However, POP-level or router-level topology of a network is confidential due to security and business reasons.

On the other hand, although there have been several random algorithm to generate a router-level topology mathematically, it was found recently that the ISP router-level topologies are highly unlikely formed by these random mechanisms [11].

Therefore, we are forced to use the topology data from a specific network instead. In fact, we adopted the network topology of China Education and Research Network (CERNET)³ and virtually divide its network into several AS to suit for our purpose to simulate an inter-domain environment. AS-level and POP-level topologies used in our experiment is shown in Figs. 3 and 4, respectively. Note that since BGP and GBR adopts the same policy preference priority, all AS in our experiment are set to peering relationship with transit allowed, in order to make full use of our all 39 POP as well as their possible data.

The parameters used in the experiment, i.e. $\bar{d}(P_i)$, $c(L_{ij})$, $t_0(L_{ij})$ and so on, also come from the data gathered from CERNET's backbone monitoring system. Note that $g(L_{ij})$, the geographical length of each link L_{ij} , is inferred from the railway length between the cities where the two POP are located, since most optic fiber are constructed along with the railway,⁴ except for L_{bj-isa} , the only satellite link in CERNET, the length is computed to adhere to common propagation delay on satellite link.

4.2. Result and analysis

According to the model, we first calculate the parameter k for each egress port of each pop using actual statistical data. And then we run a simulation based on a set of 24-h data to get routing result, using Open Shortest Path First (OSPF) [12] as intra-domain routing, as well as BGP and GBR as inter-domain routing protocol, respectively, thus the queuing delay of each link $d_{queue}^{(\varphi)}(L_{ij})$ in both routing mechanisms. Together with propagation delay $d_{prop}^{(\varphi)}(L_{ij})$ calculated with $g(L_{ij})$, the average performance index

$$d^{(BGP)} = d_{prop}^{(BGP)} + d_{queue}^{(BGP)} \quad (8)$$

and

$$d^{(GBR)} = d_{prop}^{(GBR)} + d_{queue}^{(GBR)} \quad (9)$$

is worked out, respectively. The result of the simulation is illustrated in Figs. 5 and 6.

Fig. 5 presents two proportions made up the performance index, $d_{queue}^{(\varphi)}$ and $d_{prop}^{(\varphi)}$, which are the former and latter item in (3), respectively.




Fig. 3. AS-level topology used in the experiment.

From the histograms we can infer that in most of the time, $d_{prop}^{(GBR)}$ is less than $d_{prop}^{(BGP)}$, while $d_{queue}^{(GBR)}$ is greater than $d_{queue}^{(BGP)}$. This result fit with intuitive assumption that GBR can reduction propagation delay since it optimize the route geographically, where BGP, with its AS-hop count policy, can lighten the congestion on the already crowded inter-domain links, and result in the reduce of queuing delay caused by congestion.

Fig. 6 indicates the comparison between $d^{(GBR)}$ and $d^{(BGP)}$, from which we can infer that though GBR results in a lower propagation delay, it still behave worse than BGP since queuing delay plays the main role in end-to-end transmission delay.

However, Fig. 7 shows the difference between BGP and GBR delays, on queuing and propagation, respectively. In can be inferred that GBR result in an almost invariable reduce in propagation delay, while the advantage of BGP on queuing delay depends on the traffic load. When the network bandwidth is much higher than required, the propagation delay may play the main role and thus GBR, which is optimized for propagation delay, would have better performance in the sense that it has less average end-to-end transmission delay.

4.3. Optimal route estimation

After discovering that GBR does not improve overall performance of BGP well, it is intuitively that estimation on the optimal route performance in current topology by our performance index model is wanted. The theoretical algorithm presented in Section 3 would be a proper contrast.

As the algorithm is executed in our simulation, we get the performance indicator of average transmission delay in optimal situation and compare it with the better one between GBR and BGP, two current mechanisms discussed in this paper, as illustrated in Fig. 6. It can be inferred from the curves that great disparity exists between the performance current routing mechanisms and the optimal ones, which is greater than 50% in some cases.

From this result, we can infer that there is great different between the performance of current BGP and sub-optimal value. And GBR did not contribute significant performance improvement from BGP. There is still much to invest on the performance of routing mechanisms.

5. Discussion

In this section, we discuss some open issues and some possible ways to improve current inter-domain routing practice that deserve further attention.

5.1. Problems of geographic-based routing

Previously, our result shows that GBR methods such as GIRO can result in a reduction in propagation delay on the link; however do not result in improvement of actual data delivery performance in common case. Moreover, more concerns relative to GBR methods need to be discussed.

² For example, see <http://www.routeview.org/>.

³ See <http://www.edu.cn/>.

⁴ See <http://www.shike.org.cn/>




Fig. 4. POP-level topology used in the experiment.





Fig. 5. Detailed simulation result of GBR and BGP.

Intuitively, in any GBR such as GIRO, relevant geographical information should be propagated within routing information. In




Fig. 6. Result comparison of GBR, BGP, and estimated optimal performance.

fact, GIRO attempts to involve that information into IP address. However, AS are basic visible entities in the current inter-domain routing system [13] and this attempt to propagate router-level or POP-level information among AS is in fact impossible due to security as well as scalability concerns. As a result, geographical information can only be propagated in an aggregated form. For example, total geographical distance of a route, may be propagated globally as the weight of a route. This compromise restricts the variety of the usage of geographical information.

Moreover, cheating may also be an issue under such a situation. A network operator may decrease the fore-mentioned aggregated geographical information, the distance value to attract, or increase the distance value to avoid traffic to a specific destination through it. What is worse, if AS-hop metric is absent in the routing decision process, such manipulation may result in routing loop.

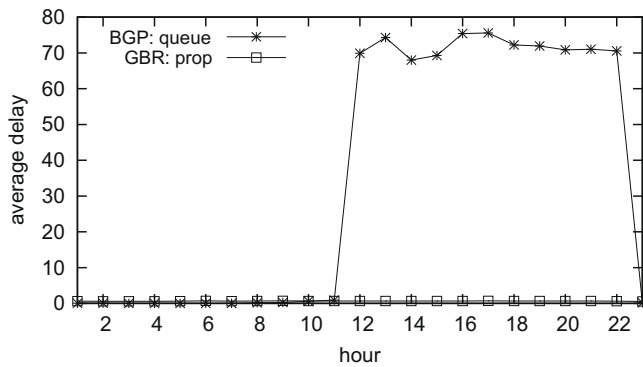


Fig. 7. Difference between BGP and GBR advantages, on queuing and propagation delay, respectively.

5.2. Inter-domain routing metric

In previous research on inter-domain routing, the most focused topics are about variety of policies and routing convergence. The performance property of a routing protocol, however, is not concerned broadly by researchers. Whereas, recently, as GIRO and some other routing mechanisms indicated [14], performance also becomes an important factor in routing protocol design [13]. Thus a performance factor of routing metric should be presented. As indicated by our evaluation, choosing shortest geographical distance as the tie-breaker for inter-domain routing does not necessarily yield the best performance in all circumstances than traditional AS-hop count metrics. It is also posted in our experiment that a performance optimal routing mechanism may result in a 50% of performance improvement according to our modeled index of total transmission delay.

It is clear that a good performance factor of routing metric should also consider capability of network infrastructure as well as the aggregated behavior of its users. We also argue that for a commercial network operator, the economical factor should also be embraced into the routing metric.

Besides, the idea showed in GIRO that sharing router-level information between different AS implicates a drawback of current AS-based routing architecture. In current routing mechanisms, AS are only entities visible globally in the inter-domain routing system [1], and router-level information sharing between AS is impossible. In order to enable more permissive and effective routing mechanisms such as [15], it should not be assumed that AS are necessarily needed to be only visible entities exposed by the inter-domain routing protocol.

6. Conclusion

The emerging idea of geographic-based routing shows that the performance properties of an inter-domain routing protocol have

been attracted more and more attention by both researchers and network operators.

In this paper, a performance model based on transmission delay is carefully developed. The model not only considers the propagation delay, but also embraces the queuing delay within the route.

Based on the model, a performance evaluation on geographic-based routing is carried out. The result shows that geographical distance cannot serve as a better metric for inter-domain routing in current network circumstances.

The performance model also leads to the design of an optimizing algorithm indicating the sub-optimal performance of routing mechanisms on a certain topology. Experiment also shows that there is still a large distance between the performance of current routing mechanisms and this optimal value. Our model of performance criteria presented in our paper may become one of the guide lines in the design of novel routing mechanisms where performance is a concern.

According to our knowledge, this article is the first to model and evaluate inter-domain routing performance in the context of actual packet delivering behavior.

References

- [1] Y. Rekhter, T. Li, S. Hares, A Border Gateway Protocol 4 (BGP-4), IETF RFC4271, 2006.
- [2] M. Freedman, M. Vutukuru, N. Feamster, H. Balakrishnan, Geographic locality of ip prefixes, in: Proceedings of the Internet Measurement Conference (IMC2005), Berkeley, USA, 2005.
- [3] R. Oliveira, M. Lad, B. Zhang, L. Zhang, Geographically informed inter-domain routing, in: Proceedings of 15th IEEE International Conference on Network Protocols (ICNP2007), Beijing, China, 2007.
- [4] N. Spring, R. Mahajan, T. Anderson, Quantifying the causes of path inflation, in: Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany, 2003.
- [5] G. Finn, Routing and Addressing Problems in Large Metropolitan-Scale Internetworks, Tech. Rep. ISI-TR-2001-552, Information Science Institution, University of South California, 1987.
- [6] T. Hain, Application and Use of the IPv6 Provider Independent Global Unicast Address Format, I-D draft-hain-ipv6-pi-addr-use-10, 2006.
- [7] R. Percassi, A. Vespignani, Scale-free behavior of the internet global performance, European Physical Journal B 32 (2003) 411–414.
- [8] D. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi, C. Diot, Analysis of measured single-hop delay from an operational backbone network, in: Proceedings of 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM2002), New York, USA, 2002.
- [9] M. Schwartz, BroadBand Integrated Networks, Prentice Hall, Inc., 1996.
- [10] A. Basu, A. Lin, S. Ramanathan, Routing using potentials: a dynamic traffic-aware routing algorithm, in: Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany, 2003.
- [11] W. Thong, G. Chen, K. Ko, Searching ISP router networks for footprints of engineering design consideration, in: Proceedings of 15th IEEE International Conference on Networks (ICON2007), Adelaide, Australia, 2007.
- [12] J. Moy, OSPF Version 2, IETF RFC2328, 1998.
- [13] O. Bonaventure, Reconsidering the Internet Routing Architecture, I-D draft-bonaventure-irtf-rrg-rira-00, 2007.
- [14] H. Xiong, M. Chen, M. Zhang, PIRP: a performance-aware inter-domain routing protocol, in: Proceedings of International Conference on Communication Technology (ICCT2006), Guilin, China, 2006.
- [15] L. Subramanian, M. Caesar, T. Cheng, M. Handley, M. Mao, S. Shenker, I. Stoica, HLP: a next generation interdomain routing protocol, in: Proceedings of ACM SIGCOMM 2005, Philadelphia, USA, 2005.