ISPs as Nodes or Sets of Links?

Praveen K. Muthuswamy, Koushik Kar, Aparna Gupta
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
Email: {muthup,kark,guptaa}@rpi.edu

Hasan T. Karaoglu, Murat Yuksel
University of Nevada, Reno
Reno, NV 89557, US
Email:{karaoglu,yuksem}@cse.unr.edu

***Abstract*—We consider the contract-switching paradigm for studying the inter-domain traffic engineering problem. In the contract-switching paradigm, each ISP in the Internet is abstracted as a set of edge-to-edge contract links. We formulate the optimal routing problem for the contract-switching paradigm by considering three objectives, namely: 1) maximizing throughput, 2) minimizing delay, and 3) minimizing bandwidth usage. We solve the optimization problems on realistic network topologies and show that the routing solutions developed using the contract-switching paradigm provides significant improvement in performance compare to the BGP routing framework with respect to the three objectives. Moreover, our simulation study also reveals that the contract-switching paradigm performs close to the best performance that can be achieved in the Internet in the absence of any abstractions.**

I. Introduction

Attaining the goal of a high-efficiency inter-domain traffic engineering (TE) involves many factors ranging from technical capabilities to economics and policy [1]. Due to the multi-provider nature of the Internet, business goals and policies have to be respected while trying to attain this goal. ISPs have to play the game of increasing their connectivity and attracting more demand by establishing more peering and transit points with their fellow ISPs, and yet compete with those same ISPs as a business. We focus on the platform where this ongoing game is being played, i.e., in inter-domain routing. The existing routing architecture of the Internet follows a two-level model. Intra-domain level employs proactive approaches focusing on reliability and quality. Intra-domain TE has become a common practice with numerous methods such as MPLS [2], DiffServ [3], RSVP [4] and link weight optimization for Intra-Gateway protocols (e.g., OSPF and IS-IS) [5]. Using these methods, many researchers have proposed both online and offline algorithms that consider network throughput optimization, congestion avoidance, interference minimization and increasing network reliability [6], [7]. Given practically-accurate traffic demand matrices and greater control within domain borders, intra-domain traffic engineering can achieve relatively stable performance on even multi-objective combinations of these performance metrics. However, current inter-domain TE methods in general are unable to attain particular performance targets like quality and reliability. Unless ISPs establish a community, which is not easy to do, inter-domain TE techniques are mostly constrained to outbound traffic load balancing; and other TE goals and desired networking practices are only expressed indirectly. Common expression methods include increasing LOCALPREF, ASPATH, or MED at peering points where less inbound traffic is desired [1].

Limitations of inter-domain traffic engineering have attracted many researchers to studying cooperative traffic engineering mechanisms where neighboring ISPs work on feasible traffic management outcomes for the benefit of both parties [8]. Promising benefits of cooperative traffic engineering have driven research community to develop signaling protocols to mediate such negotiation and distributed decision making [9]. Another challenge in this context is the counter-productive effects of intra-domain and inter-domain traffic engineering policies if they are not designed with careful consideration of their interaction with each other [10].

A key characteristic of the Internet's routing architecture is its abstraction of an ISP as a single node - or a set of nodes - in the case of multiple autonomous systems (ASs) per ISP. This abstraction model eases the task of inter-domain routing significantly. However, the cost is the loss of path selection and flexibility beyond shortest-path routing in terms of AS hops. Sufficient flexibility at the routing level is crucial to practicing TE. Thus, recent work explored ways of overcoming this issue by techniques like inter-AS source routing [11], GMPLS [12]. Introducing link-state mechanisms to the inter-domain routing were also considered as a way to enable more controlled and diverse path selections [13] .

Some recent papers have proposed changing the abstraction of an ISP at the inter-domain routing level. Pathlet Routing [14] abstracts an ISP (or an AS) as a *set of virtual nodes* and gives more flexibility to the routing administrator in engineering the inbound or outbound traffic. Our previous work also explored this dimension by abstracting an ISP as a *set of edge-to-edge links* [15] . The core idea in our approach is to allow ISPs to expose themselves to the other ISPs as a set of edge-to-edge contracts, which we call "contract links", overlaid on their internal Interior Gateway Protocol (IGP). An internetworking paradigm of contract-switching [16] arises from these edge-to-edge contract links being used as the basic building blocks of inter-ISP transactions, including SLAs as well as routing. In this paper, we focus one fundamental question: *"Is the abstraction of ISPs as a set of edge-to-edge links sufficient to reach an optimum inter-domain TE practice?"*

We compare our contract-switching paradigm (CSP) with the existing BGP model and an ideal model (OPT), where all ISPs share their full topology information with each other. For all the three scenarios we solve the routing problem with three different goals reaching beyond simple shortest-path at the inter-domain level: maximum throughput, minimum delay, and minimum bandwidth usage. In the BGP model, an ISP is

978-1-4577-2053-6/12/$31.00 ©2012 IEEE

abstracted as a single node and thus all traffic for a particular destination can only be directed to a single next hop neighbor. In the CSP model, incoming traffic for a particular destination can be split into multiple streams depending on the number of edge-to-edge links serving this particular ingress point. Lastly, the OPT model allows selection of all end-to-end paths separately and traffic load balancing across all of them rather than within an ISP. Our evaluation clearly shows that CSP can utilize almost all of the available end-to-end quality which cannot be attained by the BGP model. We also argue that the optimal inter-domain traffic engineering solution within the CSP model can be attained through localized message exchanges among edge routers.

II. Contract-Switching: System Model

Each ISP has a set of edge routers (ingress or egress) at which the ISP peers with some neighboring ISP. The traffic can enter or leave an ISPs domain only through these edge routers. This includes the traffic from upstream ISPs as well as the traffic from hosts within the ISPs domain. In the contract-switching paradigm illustrated in Figure 1, each ISP in the Internet is abstracted as a set of edge-to-edge contract links. The edge-to-edge contract links are advertisable contracts between pairs of edge routers of an ISP.

Let there be N ISPs in the Internet. Let $\mathcal{R}_i$ denote the set of edge routers of the ISP i and let $\mathcal{L}_i^c = \{\{j,k\} : j,k \in \mathcal{R}_i\}$ denote the set of edge-to-edge contract links of the ISP i. Let $\mathcal{L}^p$ denote the set of all inter-domain peering links in the network. Note that the peering links are links between edge routers of two neighbouring ISPs. Let $\mathcal{L} = \mathcal{L}^p \cup \mathcal{L}^c$ denote the set of all contract links which includes all the intra-domain edge-to-edge contract links, $\mathcal{L}^c = \cup_{i=1}^N \mathcal{L}_i$, as well as the inter-domain peering links $\mathcal{L}^p$. Thus, in the contract-switching paradigm, the network can be modeled by the pair $(\mathcal{R}, \mathcal{L})$, where $\mathcal{R} = \cup_{i=1}^N \mathcal{R}_i$ is the set of all edge routers and $\mathcal{L}$ denotes the set of all contract and peering links. Between each pair of edge routers j and k belonging to the same ISP, there exists two contract links $\{j,k\}$ and $\{k,j\}$. If j and k belong to different ISPs, then $\{j,k\}$ and $\{k,j\}$ denotes the peering links between the two ISPs.

Each ISP is a collection of routers (edge routers as well as other internal routers) and router-to-router links. Let R_i and L_i denote the routers and the router-to-router links that belong to the ISP j, respectively. The set of all routers and router-to-router links in the network is $R = \cup_{i=1}^N R_i$, and $L = \cup_{i=1}^N L_i$, respectively. Each router-to-router link, $(x,y) \in L$, has a bandwidth capacity of C_{xy} units. Note that $(x,y) \in L$ denotes the actual router-to-router link, while $\{j,k\} \in \mathcal{L}$ denotes the contract or peering link. For two routers x and y that are physically connected, there exists two router-to-router links, $(x,y),(y,x) \in L$, having bandwidth capacity of C_{xy} and C_{yx}units, respectively. Each contract link $\{j,k\} \in \mathcal{L}^c$ is constructed using a set of physical (router-to-router) links, and therefore represents a path between the edge router j and the edge router k. If $\{j,k\}$ is a inter-domain peering link, then it would typically correspond to just one physical link. Let θ_{jk} denote the set of physical links that are used to construct the edge-to-edge contract link $\{j,k\}$. We assume that θ_{jk} is fixed and given. For example, path θ_{jk} can be the min-hop path between the edge routers j and k, as shown in the Figure 1.

Let D denote the set of all destinations in the network. Each source and destination host is attached to a particular edge router through access links as shown in Figure 1. Let e_d denote the edge router to which the destination d is attached. For simplicity, in our formulation we assume that the source nodes for destination d cannot be attached to the edge router e_d, but it can be attached to any other edge router.

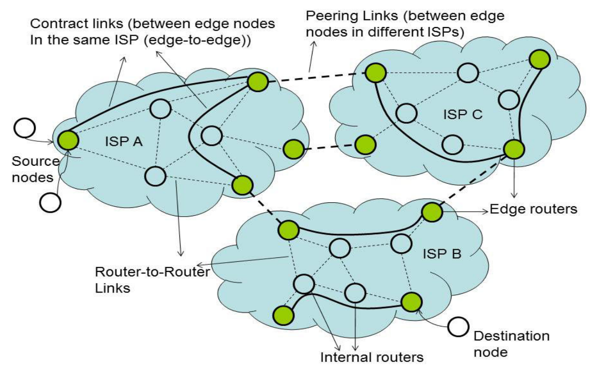

Fig. 1. Contract-switching as a framework: An ISP is abstracted as a collection of (edge-to-edge) contract links. Figure shows three ISPs A, B and C, each advertising two contract links. The dashed lines represent the router-to-router links and peering links, while the solid lines represent contract links.

III. Optimal Inter-domain Traffic Engineering Using Contract-Switching

In the formulations provided below, we assume that each flow on a contract-link follows a single, fixed intra-AS path. While this may be non-optimal, it allows us to focus solely on the performance that inter-domain traffic engineering can attain, even when a simple intra-domain routing policy is used (representative of current intra-domain routing practices).

A. Maximum Throughput Routing

We first formulate the throughput maximization problem for the contract-switched network architecture. We introduce few additional notations before stating the problem. For each edge router z, let μ_z^d denote the total traffic, up to a proportionality constant, originating at the edge router z (from the sources attached to z) for the destination d, where $d \in D$. Note that $\mu_z^d = 0$, if there are no sources attached to the edge router z for the destination d. Given the end-to-end demand μ, we wish to obtain the maximum value of the factor α by which the given end-to-end traffic rates can be scaled while satisfying feasibility constraints. Let h_{jk}^d denote the flow on the contract or peering link $\{j,k\}$ for the destination d ($d \in D$). For a given edge router z, let $\mathcal{Q}_z$ denote the edge routers that are connected to z either through a peering link or through an edge-to-edge contract link. The throughput maximization problem can be stated as follows:

$$\max \quad \alpha \tag{1}$$

$$\text{s.t.} \quad h_{jk}^d \geq 0, \quad \forall \{j,k\} \in \mathcal{L} \text{ and } d \in D, \tag{2}$$

$$\sum_{d \in D} \sum_{\{j,k\} \in \mathcal{L} : (x,y) \in \theta_{jk}} h_{jk}^d \leq \quad C_{xy}, \quad \forall (x,y) \in L, \tag{3}$$

and $\forall d \in D, z \in \mathcal{R}$,

$$\sum_{e \in \mathcal{Q}_z} h^d_{ze} - \sum_{e \in \mathcal{Q}_z} h^d_{ez} = \begin{cases} \alpha \mu^d_z, & \text{if } z \neq e_d \ , \\ -\alpha \sum_{s \in \mathcal{R}} \mu^d_s, & \text{if } z = e_d. \end{cases} \quad (4)$$

Equation (3) states that the total flow on each router-to-router link should be less than or equal its capacity. For a given router-to-router link (x, y), the set $\{\{j, k\} \in \mathcal{L} : (x, y) \in \theta_{jk}\}$, denotes the set of all contract links whose paths use the router-to-router link (x, y). Each edge-to-edge contract link also carries the flows for several destinations. Therefore, the total flow on the link (x, y), is $\sum_{d \in D} \sum_{\{j,k\}:(x,y) \in \theta_{jk}} h^d_{jk}$. The paths used by the edge-to-edge contract links within the ISPs determines this constraint. Equation (4) is the flow conservation constraint to be satisfied by the edge-to-edge flows h^d_{jk} at all the edge routers in the network for all the destinations.

B. Minimum Delay Routing

Next, we formulate the minimum delay routing problem in the contract-switched network. We are given the end-to-end demand μ. We wish to route the traffic along the contract links, in such a way that the total delay in the network is minimized. Towards this end, we associate a convex cost $V_{xy}(f)$ with each router-to-router link (x, y) in the network that represents the delay in the link (x, y). In our simulations, we use $V_{xy}(f) = \frac{1}{C_{xy} - f}$, which represents the $M/M/1$ delay expression. Here C_{xy} denotes the link capacity. Alternate link costs (convex in f) can also be considered. The minimum delay routing problem can be stated as follows:

$$\min \quad \sum_{(x,y) \in L} V_{xy}\Big(\sum_{d \in D} \sum_{\{j,k\}:(x,y) \in \theta_{jk}} h^d_{jk}\Big) \quad (5)$$

The constraints for this problem are given by Equations (2) to (4), with α set to 1. In Equation (5), the term $\sum_{\{j,k\}:(x,y) \in \theta_{jk}} \sum_{d \in D} h^d_{jk}$ represents the total flow along the router-to-router link (x, y).

C. Minimum Bandwidth Routing

Finally, we consider the bandwidth minimization problem. In this case, the objective is to minimize the total bandwidth used for routing the given end-to-end demand μ. Note that this problem can be formulated as a special case of the minimum delay routing problem described earlier, by defining $V_{xy}(f) = f$. The objective can be stated as follows:

$$\min \quad \sum_{(x,y) \in L} \Big[\sum_{d \in D} \sum_{\{j,k\} \in \mathcal{L}:(x,y) \in \theta_{jk}} h^d_{jk}\Big] \quad (6)$$

IV. Alternate Frameworks for Benchmarking

Next, we consider two frameworks that will be used to compare and benchmark the performance of the proposed contract-switching paradigm for inter-domain TE.

A. Global Optimum

This framework considers the complete Internet topology without any ISP or AS-level abstractions. The throughput optimization problem in this framework can be stated as follows:

$$\max \quad \alpha$$

$$\text{s.t.} \quad f^d_{xy} \geq 0, \quad \forall (x, y) \in L \text{ and } d \in D,$$

$$\sum_{d \in D} f^d_{xy} \leq \quad C_{xy}, \quad \forall (x, y) \in L,$$

$$\sum_{y \in Q_z} f^d_{zy} - \sum_{x \in Q_z} f^d_{xz} = \begin{cases} \alpha \mu^d_z, & \text{if } z \neq e_d \\ -\alpha \sum_s \mu^d_s, & \text{if } z = e_d \end{cases}$$

Here f^d_{xy} represents the flow on the router-to-router link (x, y) for the destination d and Q_z denotes the routers that are physically connected to the router z through a router-to-router link. Here μ^d_z denote the total traffic, up to a proportionality constant, originating at the router z (from the sources attached to z) for the destination d, where $d \in D$. $\mu^d_z = 0$, if there are no sources attached to the router z for the destination d. Also, note that the sources will only be attached to the edge routers. This formulation obtains the optimum flows along each router-to-router link in the network. This framework offers maximum flexibility for routing the end-to-end demand and hence the solution represents the maximum throughput that can be achieved in the network. The minimum delay and minimum bandwidth problems for this framework can be stated similarly.

B. BGP framework

The BGP framework abstracts each ISP in the Internet as a single node and finds the path with least number of ISP hops. Within an ISP's domain, routing is done by OSPF or RIP protocols. The BGP framework is simple and widely used, but it provides least flexibility in the choice of paths. We obtain the throughput, delay and the total bandwidth used by the BGP routing approach for comparison.

V. Simulation Results

We solve the inter-domain traffic optimization problems on Internet topologies generated using three techniques. First, we consider random network topologies, in which both the inter-domain and intra-domain topologies are generated randomly. Next, we build more realistic Internet topologies by using the BRITE [17] and the GTITM [18],[19] models for generating inter-domain connectivity. For intra-domain connectivity, we use the topology maps of the ABILENE and the GEANT ISPs that are available online [20],[21]. In Sections V-A, and V-B, we discuss the random and BRITE topology generation models and present the results of the numerical study. The results for GTITM topologies were similar and hence omitted due to space restrictions. For fairness of comparison with the BGP framework (Section IV-B), we use simple min-hop routing at the intra-domain level, when using the contract-switching framework.

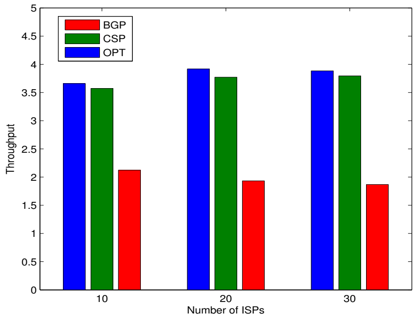

Fig. 2. Average throughput on 50 random Internet topologies. The number of ISPs is increased from 10 to 30, while the size is kept fixed at 10 routers.

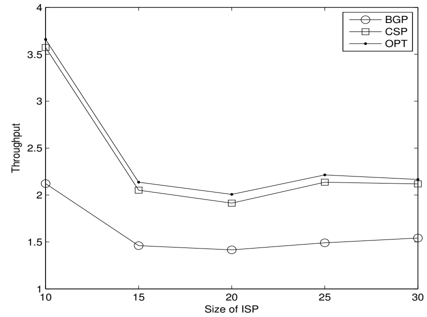

Fig. 3. Average throughput when the size of the ISP increases. 50 random network topology were generated.

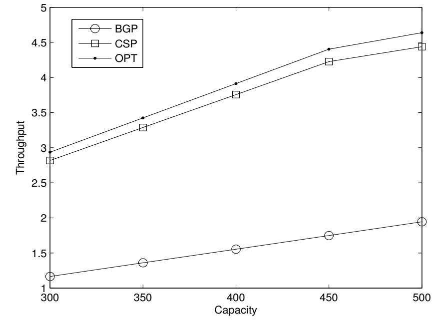

Fig. 4. Average throughput on 50 random Internet topologies for different capacity values. The capacity of all physical links is the same. Each network has 40 ISPs with 10 routers per ISP.

A. Random Topology

We first generate random Internet topologies, where both the inter-domain and intra-domain topologies are generated randomly. To generate a random network, nodes are first distributed uniformly randomly on a square plane. A link is placed between two nodes if they are within a certain distance from each other. We first randomly generate the inter-domain connectivity graph and then replace each node in this graph with another random network that represents the intra-domain topology. For each peering link, the edge routers are chosen uniformly randomly from the routers of corresponding ISPs. The number of routers within each ISP is kept the same for all ISPs. There are 100 sources and 2 destination nodes. Each source node and destination node is attached to a randomly chosen edge router. The destination node for each source is chosen randomly. The traffic demand for each source-destination pair is chosen uniformly between 1 and 10.

Figures 2, and 3, and 4 show the results of the maximum throughput routing on random Internet topologies. In the figures, CSP, OPT and BGP, refers to optimal routing using contract-switching (Section III-A), global optimum (Section IV-A), and the BGP framework (Section IV-B), respectively. For the results shown in these figures, the capacity of all the physical links was set to 500. Figure 2 shows the average throughput when the number of ISPs in the network is increased. The number of routers within each ISP is 10. Each data point is obtained by averaging over 50 networks. From the figure, we observe that the OPT routing approach offers highest throughput compared to BGP and the CSP, as expected. It is important to note that the throughput offered by CSP is significantly higher than that of BGP, and is close to the throughput offered by OPT. This behavior can be observed even when the number of ISPs in the network is increased. Also, similar results are observed in Figure 3, where the number of ISPs is fixed but the size of each ISP is increased.

Next, consider the results shown in Figure 4. We obtain the average throughput for a range of capacity values. The capacity of all the links are kept equal and is increased uniformly from 300 to 500. The throughput increases with capacity as expected. It can be observed that the CSP routing approach offers nearly the same throughput as the OPT for all the capacity values shown. For clarity of figures, we have not shown the error bars for each data point. However, the data from the numerical study showed that performance of CSP was well beyond that of BGP for all network realizations.

In Figure 5, we compare the total delay using the OPT, CSP, and the BGP routing approaches. We consider a sample random topology with 10 ISPs and 10 routers per ISP. Figure 5 shows the total delay in the network, for increasing values of capacity. The end-to-end demand is fixed and given as before. Initially, we find that all the three routing approaches are infeasible, that is, they have infinite delays. OPT becomes feasible when the capacity is around 140, while CSP approach become feasible when the capacity is around 180. However, the BGP approach becomes feasible only when the capacity is around 400. Since BGP uses the same paths for many source nodes, the links tend to be congested. Hence the capacity of the links must be made higher in order to support the BGP routing approach. The total delay for all the routing approaches drops rapidly as soon as they become feasible. OPT has the least delay among the three routing approaches, while CSP is close to OPT.

B. BRITE Topology

Next, we generate more realistic internet topologies using the popular BRITE generation model. The BRITE model can be used to generate the complete Internet (intra and inter domain) topology as well as just the inter-domain topology. We use the BRITE topology generator in accordance with the Albert-Barabasi model to generate the ISP connectivity graph (inter-domain) for a given number of ISPs. For the internal topology of each ISP, we use the topology maps of either the GEANT or the ABILENE ISP (available online) with equal probability. The GEANT ISP consists of 23 routers, while the ABILENE ISP consists of 9 routers. The edge routers are chosen randomly for each peering link between two ISPs.

Figure 6 shows the average throughput for the BRITE topologies. Comparing the Figures 6 and 4, we find that the throughput offered by CSP is closer to OPT than what is observed for random topologies. Next, consider the results shown in Figure 8, where we compare the total bandwidth used for the three routing approaches. We observe that BGP uses maximum network bandwidth. This is due to the fact that even though BGP minimizes the number of ISP hops, the paths used by BGP may not have the least number of router-to-router hops. Therefore, the total bandwidth used tends to

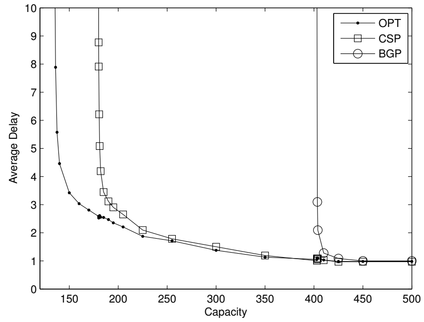

Fig. 5. Average delay on a sample random topology with 10 ISPs and 10 routers per ISP for different capacity values.

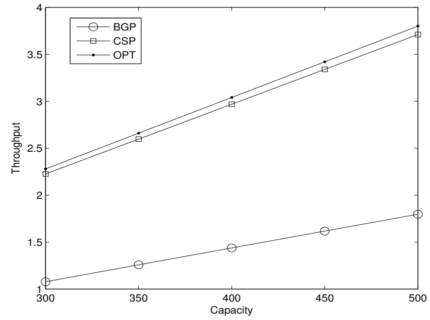

Fig. 6. Average throughput on BRITE topologies for different capacity values. 50 networks with 40 ISPs were generated.

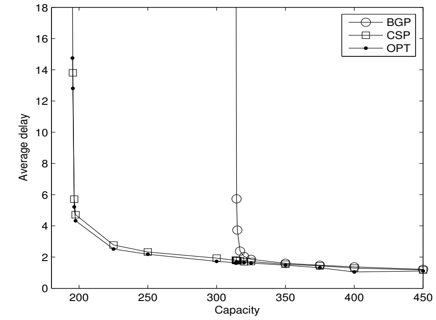

Fig. 7. Figure showing the delay on a sample BRITE topology with 10 ISPs and 10 routers per ISP for different capacity values.

be high. Also, the total bandwidth used increases with the number of ISPs in the network. We also find that both CSP and OPT use the same amount bandwidth. This can be reasoned as follows: the optimal routing approach for minimizing the total bandwidth would to choose the paths with least number of router-to-router links hops. Since both the OPT and CSP have the flexibility to choose this path, the bandwidth turns out to be the same and is also lower than that of BGP. Figure 7 shows that OPT and CSP have similar delay on a sample BRITE topology. This can be expected to hold for other networks as well, since CSP and OPT have roughly the same throughput on average for BRITE topologies.

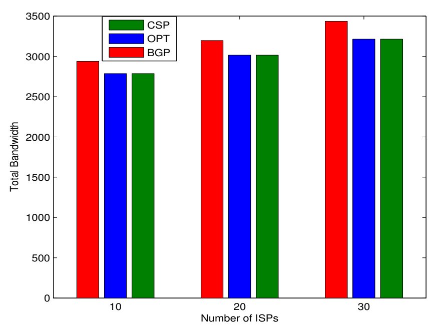

Fig. 8. Average total bandwidth used on 50 BRITE topologies when the number of ISPs in the network is increased.

VI. Conclusion

We studied the optimal routing problem in the Internet by abstracting it as a set of edge-to-edge contract links. Traditionally, the routing problem is studied either by considering the complete topology of Internet without any abstractions or by abstracting the ISPs as nodes. We have shown that the contract-switching abstraction provides significant improvement in throughput, delay and the bandwidth used, over the BGP routing approach. Moreover, our simulation study has revealed that the performance of the contract-switching framework is close to the optimal routing performance that can be achieved in the Internet without any abstractions.

Acknowledgement

This work was supported by the NSF through the NeTS-FIND program (awards CNS-0721609, CNS-0721600, CNS-0831830, and CNS-0831957).

References

[1] Jay Borkenhagen Nick Feamster and Jennifer Rexford. Guidelines for interdomain traffic engineering. *SIGCOMM Computer Communication Review*, 33:19–30, 2003.
[2] A. Viswanathan E. Rosen and R. Callon. Multiprotocol label switching architecture (mpls). RFC 3031, www.ietf.org, 2001.
[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. *IETF Internet RFC 2475*, 1998.
[4] R. Braden and et al. Resource Reservation Protocol (RSVP) - V1 functional Spec. *IETF Internet RFC 2205*, 1997.
[5] B. Fortz and M. Thorup. Internet traffic engineering by optimizing ospf weights. In *Proc. INFOCOM*, 2000.
[6] K. Kar, M. Kodialam, and T.V. Lakshman. Minimum interference routing of bandwidth guaranteed tunnels with mpls traffic engineering applications. *IEEE JSAC*, 18:2566–2579, 2000.
[7] Hao Wang, Haiyong Xie, Lili Qiu, Yang Richard Yang, Yin Zhang, and Albert Greenberg. Cope: traffic engineering in dynamic networks. In *Proc. SIGCOMM*, 2006.
[8] Gireesh Shrimali, Aditya Akella, and Almir Mutapcic. Cooperative interdomain traffic engineering using nash bargaining and decomposition. *IEEE/ACM Transactions on Networking*, 18:341–352, 2010.
[9] N. Bitar JP. Vasseur, R. Zhang and JL. Le Roux. A backward-recursive pce-based computation (brpc) procedure to compute shortest constrained inter-domain traffic engineering label switched paths. RFC 5441, www.ietf.org, 2009.
[10] Simon Balon and Guy Leduc. Combined intra- and inter-domain traffic engineering using hot-potato aware link weights optimization. In *Proc. ACM SIGMETRICS*, 2008.
[11] Xiaowei Yang, David Clark, and Arthur Berger. NIRA: A new inter-domain routing architecture. *IEEE/ACM Transactions on Networking*, 15:775–788, 2007.
[12] E. Mannie. Generalized mpls architecture.
[13] K. Levchenko, G. M. Voelker, R. Paturi, and S. Savage. Xl: an efficient network routing algorithm. In *Proc. SIGCOMM*, 2008.
[14] Scott Shenker P. Brighten Godfrey, Igor Ganichev and Ion Stoica. Pathlet routing. *Proc. SIGCOMM*, 2009.
[15] H. T. Karaoglu and M. Yuksel. Value flows: Inter-domain routing over contract links. In *Proc. IEEE GLOBECOM*, 2010.
[16] M. Yuksel, A. Gupta, and S. Kalyanaraman. Contract-switching paradigm for internet value flows and risk management. In *Proc. IEEE Global Internet Symposium*, 2008.
[17] Alberto Medina, Anukool Lakhina, Ibrahim Matta, and John Byers. Brite: An approach to universal topology generation. In *Proc. MASCOTS*, 2001.
[18] Ken Calvert, Matt Doar, and Ellen W. Zegura. Modeling internet topology. In *IEEE Communications Magazine*, 1997.
[19] http://www.cc.gatech.edu/fac/Ellen.Zegura/graphs.html.
[20] http://itservices.stanford.edu/service/internet2/abilene.
[21] http://informatique.umons.ac.be/networks/igen/.