

Improved Routing Metric In The Overlay Routing Delay Nodes

Kalyani Wasnik, Prof. Shrikant v.Sonekar

Department of computer science and engineering

J D College of engineering and Management

Nagpur,India

HOD,Vice Principal

J D College of engineering and Management

Nagpur,India

Abstract— Overlay routing is a very attractive scheme that allows improving certain properties of the routing (such as delay or TCP throughput) without the need to change the standards of the current underlying routing. However, deploying overlay routing requires the placement and maintenance of overlay infrastructure. This gives rise to the following optimization problem: Find a minimal set of overlay nodes such that the required routing properties are satisfied. In this paper, we rigorously study this optimization problem. We show that it is NP-hard and derive a nontrivial approximation algorithm for it, where the approximation ratio depends on specific properties of the problem at hand. We examine the practical aspects of the scheme by evaluating the gain one can get over several real scenarios. The first one is BGP routing, and we show, using up-to-date data reflecting the current BGP routing policy in the Internet, that a relative small number of less than 100 relay servers is sufficient to enable routing over shortest paths from a single source to all autonomous systems (ASs), reducing the average path length of inflated paths by 40%. We also demonstrate that the scheme is very useful for TCP performance improvement (results in an almost optimal placement of overlay nodes) and for Voice-over-IP (VoIP) applications where a small number of overlay nodes can significantly reduce the maximal peer-to-peer delay.

Keywords— Overlay network, resource allocation.

I. INTRODUCTION

Overlay routing has been proposed in recent years as an effective way to achieve certain routing properties, without going into the long and tedious process of standardization and global deployment of a new routing protocol. For example, in overlay routing was used to improve TCP performance over the Internet, where the main idea is to break the end-to-end feedback loop into smaller loops. This requires that nodes capable of performing TCP Piping would be present along the route at relatively small distances. Other examples for the use of overlay routing are projects like RON and Detour where overlay routing is used to improve reliability. Yet another example is the concept of the “Global-ISP” paradigm introduced in [4], where an overlay node is used to reduce latency in BGP routing. In order to deploy overlay routing over the actual physical infrastructure, one needs to deploy and manage overlay nodes that will have the new extra functionality.

This comes with a non negligible cost both in terms of capital and operating costs. Thus, it is important to study the benefit one gets from improving the routing metric against this cost. In this paper, we concentrate on this point and study the minimum number of infrastructure nodes that need to be added in order to maintain a specific property in the overlay routing. In the shortest-path routing over the Internet BGP-based routing example, this question is mapped to: What is the minimum number of relay nodes that are needed in order to make the routing between a groups of autonomous systems (ASs) use the underlying shortest path between them? In the TCP performance example, this may translate to: What is the minimal number of relay nodes needed in order to make sure that for each TCP connection, there is a path between the

connection endpoints for which every pre-defined round-trip time (RTT), there is an overlay node capable of TCP Piping?

Regardless of the specific implication in mind, we define a general optimization problem called the Overlay Routing Resource Allocation (ORRA) problem and study its complexity. It turns out that the problem is NP-hard, and we present a non-trivial approximation algorithm for it. Note that if we are only interested in improving routing properties between a single source node and a single destination, then the problem is not complicated and finding the optimal number of nodes becomes trivial since the potential candidate for overlay placement is small, and in general any assignment would be good. However, when we consider one-to-many or many-to-many scenarios, then a single overlay node may affect the path property of many paths, and thus choosing the best locations becomes much less trivial. We test our general algorithm in three specific such cases, where we have a large set of source–destination pairs, and the goal is to find a minimal set of locations, such that using overlay nodes in these locations allows to create route (routes are either underlay routes or routes that use these new relay nodes) such that a certain routing property is satisfied.

The first scenario we consider is AS-level BGP routing, where the goal is to find a minimal number of relay node locations that can allow shortest-path routing between the source–destination pairs. Recall that routing in BGP is policy-based and depends on the business relationship between peering ASs, and as a result, a considerable fraction of the paths in the Internet do not go along a shortest path. This phenomenon, called path inflation, is the motivation for this scenario. We consider a one-to-many setting where we want to improve routing between a single source and many destinations. This is the case where the algorithm power is most significant since, in the many-to-many setting, there is very little overlap between shortest paths, and thus not much improvement can be made

over a basic greedy approach. We demonstrate, using real up-to-date Internet data, that the algorithm can suggest a relatively small set of relay nodes that can significantly reduce latency in current BGP routing.

The second scenario we consider is the TCP improvement example discussed above. In this case, we test the algorithm on a synthetic random graph, and we show that the general framework can be applied also to this case, resulting in very close-to-optimal results.

The third scenario addresses overlay Voice-over-IP (VoIP) applications such as Skype (<http://www.skype.com>). Such applications are becoming more and more popular offering IP telephone services for free, but they need a bounded end-to-end delay (or latency) between any pair of users to maintain a reasonable service quality. We show that our scheme can be very useful also in this case, allowing applications to choose a smaller number of hubs, yet improving performance for many users.

Note that the algorithmic model we use assumes a full knowledge of the underlying topology, the desired routing scheme, and the locations of the required end points. In general, the algorithm is used by the entity that needs the routing improvement and carries the cost of establishing and maintaining overlay nodes, using the best available topology information. For example, in the VoIP case, the VoIP application is establishing the overlay nodes, and thus the application can gain by using our approach. The main contributions of this paper are as follows.

- We develop a general algorithmic framework that can be used in order to deal with efficient resource allocation in overlay routing.
- We develop a nontrivial approximation algorithm and prove its properties.
- We demonstrate the actual benefit one can gain from using our scheme in three practical scenarios, namely BPG routing, TCP improvement, and VoIP applications.

II. LITERATURE SURVEY

Using overlay routing to improve network performance is motivated by many works that studied the inefficiency of varieties of networking architectures and applications. Analyzing a large set of data, Savage et al. [6] explore the question: How “good” is Internet routing from a user’s perspective considering round-trip time, packet loss rate, and bandwidth? They showed that in 30%–80% of the cases, there is an alternate routing path with better quality compared to the default routing path. In [7] and later in [1], the authors show that TCP performance is strictly affected by the RTT. Thus, breaking a TCP connection into low-latency sub connections improves the overall connection performance. In [5], [8], and [9], the authors show that in many cases, routing paths in the Internet are inflated, and the actual length (in hops) of routing paths between clients is longer than the minimum hop distance between them.

Using overlay routing to improve routing and network performance has been studied before in several works. In [3], the authors studied the routing inefficiency in the Internet and used an overlay routing in order to evaluate and study experimental techniques improving the network over the real

environment. While the concept of using overlay routing to improve routing scheme was presented in this work, it did not deal with the deployment aspects and the optimization aspect of such infrastructure. A resilient overlay network (RON), which is architecture for application-layer overlay routing to be used on top of the existing Internet routing infrastructure, has been presented in [2]. Similar to our work, the main goal of this architecture is to replace the existing routing scheme, if necessary, using the overlay infrastructure. This work mainly focuses on the overlay infrastructure (monitoring and detecting routing problems, and maintaining the overlay system), and it does not consider the cost associated with the deployment of such system.

A. Problem Definition

Given a graph $G=(V,E)$ describing a network, let P_u be the set of routing paths that is derived from the underlying routing scheme, and let P_o be the set of routing paths that is derived from the overlaying routing scheme (we refer to each path in P_u and in P_o as the underlying and overlaying path sets, respectively). To Note that both P_u and P_o can be defined explicitly as a set of paths, or implicitly, e.g., as the set of shortest paths with respect to a weight function $W : E \rightarrow \mathbb{R}$ over the edges. Given a pair of vertices $s, t \in V$, denote by $P_{s,t}$ the set of overlay paths between s and t , namely $P_{s,t} \subset P_o$ and $\forall p \in P_{s,t}$, the endpoints of p are s and t .

Definition 3.1: Given a graph $G = (V,E)$, a pair of vertices (s,t) a set of underlay paths P_u , a set of overlay paths $P_{s,t}$, and a set of vertices $U \subset V$. We say that U covers (s,t) if such that is a concatenation of one or more underlying paths, and the endpoints of each one of these underlay paths are in U . Intuitively speaking, the set of vertices U , also called relay nodes, is used to perform overlay routing, from sources to destinations such that packets can be routed from one relay node to another using underlay paths. The Overlay Routing Resource Allocation (ORRA) problem is defined as follows:

Definition 3.2: Given a graph $G = (V,E)$, a set of source–destination pairs $Q = \{(s_1,t_1),(s_2,t_2),\dots,(s_n, t_n)\}$ (where $Q \subset V \times V$), a set of underlay paths P_u , and a set of overlay paths P_o find a subset of vertices such that $\forall 1 \leq i \leq n$, U covers (s_i,t_i) . Note that, in this work, we do not deal with the mechanism and the implementation aspects required by the clients and the relay nodes in order to utilize the system. This may be subject for further research. Using the assumption that single-hop paths are always in P_u , the set $U = V$ is a trivial feasible solution to the ORRA problem. Our objective, as discussed in Section I, is to minimize the deployment cost of relay nodes, thus we define the MIN-ORRA problem.

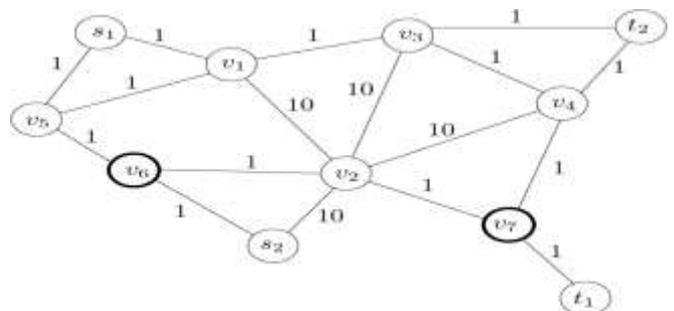


Fig. 1. Overlay routing example: Deploying relay server on v_6 and v_7 enables overlay routing

Definition 3.3: Given an instance of the ORRA problem, and a nonnegative weight function over the vertices, one needs to find a set such that: 1) is feasible; and 2) the cost of is minimal among all feasible sets.

For instance, consider the graph depicted in Fig. 1, in which the underlying routing scheme is minimum hop count, and the overlay in grouting scheme is the shortest path with respect to the edge length. In this case, the underlay path between s_1 and t_1 is $(s_1, v_1, v_2, v_7, t_1)$, while the overlay path between them should be $(s_1, v_1, v_3, v_4, v_7, t_1)$ or $(s_1, v_5, v_6, v_2, v_7, t_1)$. Similarly, the underlay path between s_2 and t_2 is (s_2, v_2, v_4, t_2) , while the overlay path between them should be $(s_2, v_6, v_2, v_4, v_7, t_2)$ or $(s_2, v_6, v_5, v_1, v_3, t_2)$. Deploying relay nodes on v_6 and v_7 implies that packets from s_1 to t_1 can be routed through the concatenation of the following underlay paths (s_1, s_5, s_6) , (v_6, v_2, v_7) , and (v_7, t_1) , while packets from s_2 to t_2 can be routed through the concatenation of the following underlay paths (s_2, v_6) , (v_6, v_2, v_7) , and (v_7, v_4, v_2) . Thus, is a feasible solution to the corresponding ORRA problem. If all the nodes have an equal weight then one may observe that this is also an optimal solution.

III. CASE STUDY AND EXPERIMENTAL RESULTS

BGP Routing Scheme BGP is a policy-based interdomain routing protocol that is used to determine the routing paths between autonomous systems in the Internet [18]. In practice, each AS is an independent business entity, and the BGP routing policy reflects the commercial relationships between connected ASs. A customer-provider relationship between ASs means that one AS (the customer) pays another AS (the provider) for Internet connectivity, a peer-peer relationship between ASs means that they have mutual agreement to serve their customers, while a sibling-sibling relationship means that they have mutual-transit agreement (i.e., serving both their customers and providers). These business relationships between ASs induce a BGP export policy in which an AS usually does not export its providers and peers routes to other providers and peers. In [10] and [11], the authors showed that this route export policy indicates that routing paths do not contain so-called valleys nor steps. In other words, after traversing a provider-customer or a peer-peer link, a path cannot traverse a customer-provider or a peer-peer link. This routing policy may cause, among other things, that data packets will not be routed along the shortest path. For instance, consider the AS topology graph depicted in Fig. 4. In this example, a vertex represents an AS, and an edge represents a peering relationship between ASs. While the length of the physical shortest path between AS6 and AS4 is two (using the path AS6, AS7, AS4), this is not a valid routing path since it traverses a valley. In this case, the length of the shortest valid routing path is five (using the path AS6, AS5, AS1, AS2, AS3, AS4). In practice, using real data gathered from 41 BGP routing tables, Gao and Wand [5] showed that about 20% of AS routing paths are longer than the shortest AS physical paths.

While routing policy is a fundamental and important feature of BGP, some application may require to route data using the shortest physical paths. In this case, using overlay routing, one can perform routing via shortest paths despite the policy. In this case, relay nodes should be deployed on servers located in certain carefully chosen ASs.

TPC level improvement:

We consider is the TPC level improvement in the wireless networks as explained in the above module. In this case, we test our proposed algorithm on a synthetic random graph, and we show that the general framework can be applied also to this case, resulting in very close-to-optimal results.

Voice-over-IP:

Voice-Over-IP type of applications are becoming more and more popular offering IP telephone services for free, but they need a bounded end-to-end delay (or latency) between any pair of users to maintain a reasonable service quality. We show that our scheme can be very useful also in this case, allowing applications to choose a smaller number of hubs, yet improving performance for many users.

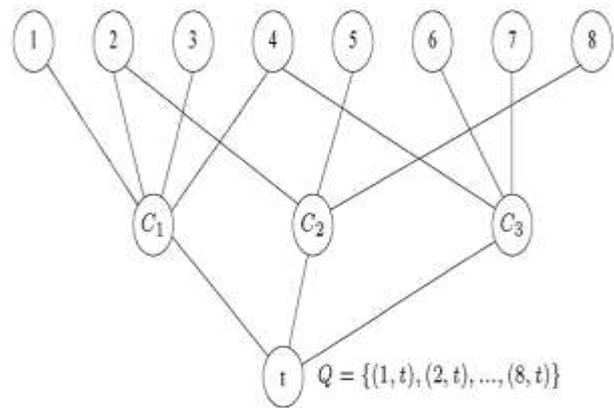


Fig. 2. Set Cover Diagram for proposed System

IV. CONCLUSION

While using overlay routing to improve network performance was studied in the past by many works both practical and theoretical, very few of them consider the cost associated with the deployment of overlay infrastructure. In this paper, we addressed this fundamental problem developing an approximation algorithm to the problem. Rather than considering a customized algorithm for a specific application or scenario, we suggested a general framework that fits a large set of overlay applications. Considering three different practical scenarios, we evaluated the performance of the algorithm, showing that in practice the algorithm provides close-to-optimal results. Many issues are left for further research. One interesting direction is an analytical study of the vertex cut used in the algorithm. It would be interesting to find properties of the underlay and overlay routing that assure a bound on the size of the cut. It would be also interesting to study the performance of our framework for other routing scenarios and to study issues related to actual implementation of the scheme. In particular, the connection between the cost in terms of establishing overlay nodes and the benefit in terms of performance gain achieved due to the improved routing is not trivial, and it is interesting to investigate it. The business relationship between the different players in the various use cases is complex, and thus it is important to study the economical aspects of the scheme as well. For example, the one-to-many BGP routing scheme can be used by a large content provider in order to improve the user experience of its customers. The VoIP scheme can be used by VoIP services (such as Skype) to improve call quality of their customers. In both these cases, the exact

translation of the service performance gain into actual revenue is not clear and can benefit from further research.

V. REFERENCES

- [1] H. Pucha and Y. C. Hu, "Overlay TCP: Multi-hop overlay transport for high throughput transfers in the Internet," Purdue University, West Lafayette, IN, USA, Tech. Rep., 2005.
- [2] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," in Proc. 18th ACM SOSP, 2001, pp. 131–145.
- [3] S. Savage, T. A. A. Aggarawl, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan, "Detour: A case for informed internet routing and transport," IEEE Micro, vol. 19, no. 1, pp. 50–59, Jan.–Feb. 1999.
- [4] R. Cohen and A. Shochot, "The "global-ISP" paradigm," Comput. Netw., vol. 51, no. 8, pp. 1908–1921, 2007.
- [5] L. Gao and F. Wang, "The extent of as path inflation by routing poli- cies," in Proc. IEEE GLOBECOM, 2002, vol. 3, pp. 2180–2184.
- [6] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The end-to-end effects of Internet path selection," in Proc. ACM SIG- COMM, 1999, pp. 289–299.
- [7] R. Cohen and S. Ramanathan, "Using proxies to enhance TCP performance over hybrid fiber/coaxial networks," Comput. Commun., vol. 20, no. 16, pp. 1502–1518, Jan. 1998.
- [8] N. Spring, R. Mahajan, and T. Anderson, "The causes of path infla- tion," in Proc. ACM SIGCOMM, 2003, pp. 113–124.
- 646 IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 22, NO. 2, APRIL 2014
- [9] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The impact of routing policy on Internet paths," in Proc. IEEE INFOCOM, 2001, pp. 736–742.
- [10] L. Gao and J. Rexford, "Stable Internet routing without global coordination," IEEE/ACM Trans. Netw., vol. 9, no. 6, pp. 681–692, Dec. 2001.
- [11] L. Gao, T. Griffin, and J. Rexford, "Inherently safe backup routing with BGP," in Proc. IEEE INFOCOM, 2001, pp. 547–556.