The Sliding Shortest Path Algorithm

Ramesh Bhandari
Laboratory for Telecommunications Sciences, Department of Defense
College Park, Maryland 20740
rbhandari@ieee.org

1. Introduction

In telecommunication networks, it is often desirable to change the link weights to engender path changes in a certain desired fashion. For example, a bank data transfer in the commercial world between a given pair of offices might need to be occasionally processed at a third intermediate office before arriving at the destination node, in which case one might alter weights of a certain subset of links to cause the route to traverse the intermediate node. In another instance, a high capacity link might suddenly become available over which the traffic between the given source-destination pair of an important customer might need to be rerouted to alleviate congestion on the current route, and maintain the desired quality of service as warranted, perhaps, by a service-level agreement.

The problem of changing routes within a network by altering the weights has received immense attention in the recent past, see, e.g., Ref.. [1-3]. Very often, the techniques involve integer programming formulations because of the requirement of integrality of link weights, leading frequently to approximations such as linear relaxation [2,3].

In this paper, we address the problem of changing the graph weights for a specific demand (a single source-destination pair) to be rerouted through a desired link or vertex not already on the shortest path of the given demand. We first require that the number of links on which the weights are changed be as low as possible in order to reduce the implementation time of link weight changes within the current OSPF-type environment. We then require that the weight changes (increments) be as small as possible in order to minimize the number of other shortest paths (routes of other demands) impacted.

2. Problem Definition

Let $G=(V, E)$ denote an undirected graph, representing a bidirectional network (e.g., a single autonomous system); V is the set of vertices (or nodes), and E is the set of weighted edges (or links); weights are non-negative integers; we assume loops and multiple edges are absent in the graph; routing of traffic demands from one point to another within the network takes place along shortest paths, and the shortest paths are unique.

In what follows, the terms *network* and *graph* will be used interchangeably, and so will the terms *node* and *vertex* , and *link* and *edge*.

Figure 1 shows a network cloud with a spur and a loop-like configuration, and a shortest (least-cost) path connecting vertex s and vertex t within the network; traffic flows from s to t as well as from t to s in the opposite direction along the same path. A desired (or constraint) link pq is also shown. Note that the given constraint link in the problem should not be part of either the spur or

the loop-like configuration or any combination thereof because any path computed to include the constraint link must be a simple path. A simple path refers to a path in which a given node is not visited more than once. Unless otherwise stated, in what follows, the term *path* refers to a simple path
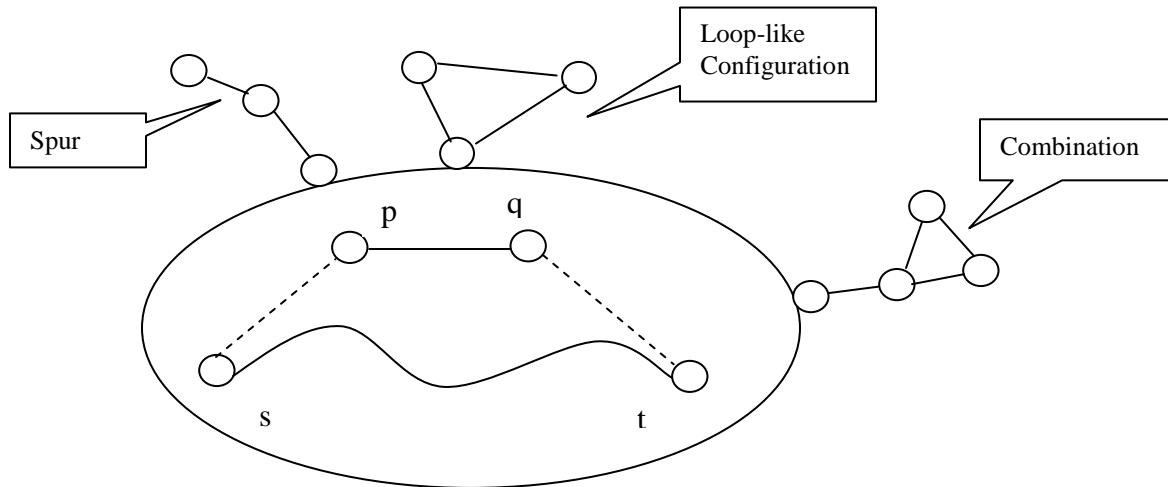

 Figure 1 A network cloud with a path from s to t, a desired (or constraint) link pq, a spur, a loop-like configuration, and a combination of spur and loop-like configurations.

 The problem to solve can be stated as:

Given a constraint link, such as link pq, in Figure 1, what is the smallest set of links whose weights should be changed and what is the minimal increment required on each such link, so that the traffic previously routed between node s and node t now flows over the constraint link pq?

This problem appears to be difficult, and we attempt to solve it using graph-theoretic heuristics. The solution is obtained in two parts via the development of a new algorithm, called the Sliding Shortest Path Algorithm (Using Link Cuts), and the assignment of link weight increments based on a formula derived from simple path analytics (Section 4). The two together constitute what we term as the Sliding Shortest Path Algorithm (Using Minimal Finite Weight Increments). The Sliding Shortest Path Algorithm (Using Link Cuts) is an algorithm (Section 3) which attempts to find the minimal set of links to cut so that the flow of interest (s-t demand) traverses the constraint link pq; this cut-set is then identified as the link set to which the weight increments should apply in the original graph.

 3. The Sliding Shortest Path Algorithm (Using Link Cuts)

This heuristic is an iterative procedure of trimming the network (cutting one link at a time) until the shortest path between s and t "slides" over the given constraint link pq. The link to cut is determined by comparing the s to t shortest path (denoted hereafter by SP(st)) with the "shortest" path[1] from s to p or s to q (denoted by SP1), and finding the first link of SP(st), which does not overlap with SP1; this link is then removed from the network and SP(st) is recomputed in the pruned network; the process of comparing and truncating the first non-overlapping link and re-

---

[1] This will be explained later, but for the present it is to be interpreted to mean the shortest path in the usual sense.

computing the shortest path is repeated until SP(st) passes through the constraint link pq. When this happens the algorithm terminates. Another run is subsequently made, this time comparing the "shortest" path from t to q (or p), denoted hereafter by SP2, with SP(ts), the shortest path from t to s, which is the same as SP(st) in the reverse order.

 Refer to Figure 2, which illustrates the concept of the algorithm. Initially, SP(st) = sabcdt. SP1 = sabep. The first non-overlapping link of the shortest path is link bc, which is removed from the network (Figure 2a). The recomputed shortest path SP(st) in the next iteration is sabef…t (Figure 2b); the first non-overlapping link of SP(st) is link ef, which is subsequently removed from the network. This process is continued until the recomputed shortest path passes through link pq. Note that if the recomputed shortest path after the cut of link bc in Figure 2a had traversed link sg, then link sg, being the first non-overlapping link of the recomputed shortest path, would have been cut in the next iteration. In general, the recomputed SP(st) can take a completely different "direction", depending upon the structure of the graph.


 Figure 2 a) Link bc is cut in the first iteration of the algorithm b) Link eg is cut in the second iteration.

 The algorithm converges without ever disconnecting the source node s from the termination node t; the final shortest path from s to t that passes over link pq is found to be given by

SP(st) = SP1 + arc pq (or qp) + SP2',

where SP2' = SP2 in the reverse order; the two choices for the SP(st) solution are due to the fact that link pq can be traversed in the direction p to q or q to p.

In a similar way, when the algorithm is run via consideration of the shortest path from t to s, the final shortest path is given by

SP(ts) = SP2 + arc qp (or pq) + SP1',

where SP1' = SP1 in the reverse order.

Below we provide a proof, and the meaning of "shortest" path.

Theorem 1: Any shortest path from s to t over the constraint link pq (if it exists) is comprised of the link pq (traversed along the arc pq or arc qp) and the shortest pair of node-disjoint paths, one path connecting s to p (or q) and the other connecting q (or p) to t.

Proof: Let us assume that a path from s to t passing through the constraint link pq exists. The existence of a path from s to t that passes over the constraint link pq necessarily implies that it is made up of either a path from s to p, arc pq, and a path from q to t, or a path from s to q, arc qp, and a path from p to t. Since any computed shortest path between a pair of nodes in a given graph has to be a simple path, the paths s to p and q to t (or alternatively s to q and p to t) must necessarily be node-disjoint, i.e., have no nodes in common.

Furthermore, because the computed path over link pq is a shortest path, the paths s to p (or q) and q (or p) to t must necessarily comprise the shortest pair of node-disjoint paths; shortest means that the sum of the individual lengths of the two paths is the smallest among all possible pairs of node-disjoint paths. ***End of Proof (Theorem 1)***.

The meaning of the term <u>"shortest"</u> is now clear; the "shortest" path from s to p (or q), denoted by SP1, and the "shortest" path from t to q (or p), denoted by SP2, refer to the paths comprising the shortest pair of node-disjoint paths, one connecting s to p (or q) and the other connecting t to q (or p). Algorithms for finding such a pair of paths exist [4,5]. They find the pair of paths simultaneously and automatically determine whether SP1 is a connection from s to p or s to q; if SP1 turns out to be a connection from s to p, SP2 is a connection from t to q, and vice versa.

Theorem 2: The process of cutting one link at a time until the shortest path from s to t slides over link pq does not disconnect t from s, i.e., the algorithm always converges to a feasible solution.

Proof: Refer to Figure 2. During each iteration when a link is cut,

1) *Path SP1 remains intact:* The link that is cut off in each iteration is the first link of SP(st) that does not overlap with SP1. Therefore, it does not belong to the set of links that comprise SP1. Consequently, SP1 remains intact.
2) *Path SP2' remains intact*: The cut link emanates from SP1, and therefore cannot belong to the set of links that comprise SP2', because SP2 is node-disjoint from SP1, i.e., at all points along its path, SP2' is at least one link apart from SP1. Therefore, like SP1, SP2' remains intact at the end of each iteration.
3) *Link pq is not cut*: A cut of link pq would imply that the recomputed shortest path (in the earlier iteration) contained link pq, a contradiction since the algorithm terminates as soon as SP(st) is found to traverse link pq.

Because SP1 and SP2', computed at the outset, remain unaffected during the iterative process, and link pq is never cut, a path always exists from s to t via link pq. Because a path over link pq from s to t is always available during the iterative process, the algorithm will always terminate. ***End of Proof (Theorem 2).***

Theorem 3: Path SP(st) after termination of the algorithm comprises SP1 , SP2', and link pq (arc pq or qp)

Proof: At termination, SP(st) passes through link pq. By Theorem 1, this final SP(st) must comprise the shortest pair of node-disjoint paths in the final truncated graph and link pq (arc pq or qp). Because paths SP1 and SP2' exist in this (final) truncated graph (see proof, Theorem 2) and SP1 ad SP2' comprise the shortest pair of node-disjoint paths in the original graph, SP1 and SP2' must necessarily be the shortest pair of node-disjoint paths in this final truncated graph (which is a sub-graph of the original graph). *End of Proof (Theorem 3).*

 In a similar way, proof is constructed for the case when SP(ts) is compared with SP2. The cut-set obtained for this case can be different.

The Sliding Shortest Path Algorithm (Using Link Cuts) can now be stated as follows:

1) Compute the initial flow path, which is the shortest path from s to t (call it SP(st)). If this path contains arc pq or arc qp terminate; otherwise, go to Step 2.
2) Compute the shortest pair of node-disjoint paths, one path connecting s to p (or q) (call it SP1), and the other connecting node t to q (or p) (call it SP2); see Ref. [4,5] for computation of node-disjoint paths.
3) Assign $\Gamma(1) = \varnothing$, where $\Gamma(1)$ denotes the set of cut links.
4) Find the first link of SP(st), which does not overlap with SP1; remove this link from the network (or graph); denote this link by *L*.
5) Set $\Gamma(1) = \Gamma(1) \cup L$.
6) Compute new SP(st) in the trimmed network.
7) If the new path contains arc pq or qp, terminate; otherwise go to Step 4.

Repeat the above algorithm in the original graph by considering the shortest path from t to s (i.e., replacing SP(st) with SP(ts) above) and comparing it with SP2, instead of SP1. Denote the cut-set obtained in this new run by $\Gamma(2)$, i.e., replace $\Gamma(1)$ with $\Gamma(2)$ above. Note that SP(ts) = SP(st) in reverse order, with Steps 1 and 2 not required in the second run.

The two runs of the algorithm yield two cut-sets, $\Gamma(1)$ and $\Gamma(2)$, which can be different. Choose the one, which has less number of links. If there is a tie, either one may be chosen.


Figure 3 A graph of 8 nodes and 12 links with their weights indicated.

Refer to Figure 3. Assume s = A and t = H, and p=B and q=C. SP(st) = ADFGH of length = 1 + 1 + 1 + 2 = 5. SP1 = ADFGB and SP2 = HC comprise the shortest pair of node-disjoint paths, one

connecting A to B and the other connecting H to C. SP(st) deviates from SP1 at node G; the first non-overlapping link of SP(st) is link GH. Upon deleting this link from the network, the recomputed SP(st) is ADFH. This path deviates from SP1 at node F again, and the first non-overlapping link is FH. Upon deleting this link, we find that the new SP(st) path is ADFGBCH, which traverses the desired constraint link BC. The algorithm terminates at this point for the shortest path run from s to t; $\Gamma(1) = \{GH, FH\}$, i.e., two links are cut to force the traffic over link BC.

Repeating the algorithm with s=H and t=A and using SP2 = HC for comparison purposes, we find that SP(ts) = HGFDA. Comparing with SP2, the first non-overlapping link to cut is HG. This leads to SP(ts) = HFDA in the truncated graph. It traverses the constraint link BC and the algorithm terminates. $\Gamma(2) = \Gamma(1) = \{GH, FH\}$; they need not be identical in general as seen from the example, where p = B and q = G; here $\Gamma(1) = \{AD, AE\}$ and $\Gamma(2) = \{GF\}$. Similarly, for p = E and q = F, $\Gamma(1) = \{AD\}$ and $\Gamma(2) = \{DF\}$.

4. Sliding Shortest Path Algorithm (Using Finite Weight Increments)

Instead of cutting the first non-overlapping link encountered as in Section 3, increment its length:

$$l_i = l_i + l(P_f) - l(P_i) + e,$$

where l_i is the length of the first non-overlapping link encountered in the ith iteration of the algorithm; $l(P_i)$ is the length of the corresponding shortest path (SP(st)); $l(P_f)$ is the length of the final desired path; $l(P_f) = l(SP1) + l(SP2) + l_{pq}$; e is an infinitesimally small positive number. For the integral weights, e = 1. The change defined above is the minimal amount needed to make path P_i greater (in length) than the desired path P_f; as a result, the latter becomes the shortest path in the final (modified) graph.

In Figure 3, for the case of p = B and q = C, changing l_{GH} from 2 to 8 and l_{FH} from 7 to 9 (in accordance with the above formula) changes SP(st) from ADFGH to ADFGBCH.

For the case, $|\Gamma(2)| = |\Gamma(1)|$ (such as p =E and q = F in Section 3), the tie may be broken by the additional criterion of selecting the set for which the number of other shortest paths impacted is the lower of the two, and so on.

Efficiency: In the worst case scenario, a shortest path algorithm such as the Dijkstra algorithm will have to be run as many times as the sum of the degrees of the vertices on SP1 (or SP2 as the case may be). Fortunately, telecommunication networks are sparse. The number of times the Dijkstra algorithm may have to be run tends to be of the order of the maximum number of edge-disjoint paths possible between vertices s and t (the minimum number of link cuts that disconnect s from t). This number may vary from 1 to 5. The developed computer code has been successfully tested on graphs as large as 200,000 nodes and 500,000 links.

5. Discussion

In this paper, we have presented a heuristic algorithm called the Sliding Shortest Path Algorithm to solve the problem of determining link weight changes within a given network to alter the route of a given flow to include a specified link (the case of the route including a specified vertex is obtained simply by collapsing the link pq into a single vertex p); the algorithm identifies a minimal set of links in accordance with a certain cutting criterion and calculates the minimal

increments needed on this set of links. The algorithm always converges, finding a feasible solution in a finite number of steps. For sparse networks, it is expected to be fast, basically mimicking the performance of a shortest path algorithm like the Dijkstra. It has been successfully tested on graphs as large as 200,000 nodes with 50000 links. Note that, although both the versions of the Sliding Shortest Path Algorithm (Sections 3 and 4) accomplish the goal of rerouting the flow of interest over a specific edge or vertex, the latter is to be preferred because of the relative less turmoil in the network due to the finite, minimal, weight increments as opposed to link cuts (equivalent to "infinite" weight increments); turmoil here is measured by the number of other shortest paths affected, i.e., shortest paths for other pairs of nodes in the network. The algorithm is also easily extended to the case where the shortest paths in the given graph are not unique.

References:

1. B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional IP routing protocols*", IEEE Communications Magazine*, 40(10), pp118-124, 2002.
2. W. Ben-Ameur and E. Gourdin, "Internet Routing and Related Topology Issues", *SIAM Journal of Discrete Mathematics*, 17(1), pp18-49, 2003.
3. Pal Nilsson, "On the Inverse Shortest Path Algorithm", *17th Nordic Teletraffic Seminar (NTS 17),* Fornebu, Norway, August 2004.
4. J.W. Suurballe, "Disjoint Paths in a Network", *Networks* 4, 125-145 (1974).
5. Ramesh Bhandari, "Survivable Networks: Algorithms for Diverse Routing", Kluwer Academic Publishers (1998).