Improved LP-rounding Approximations for the k-Disjoint Restricted Shortest Paths Problem*

Longkun Guo

College of Mathematics and Computer Science,
Fuzhou University, China

Abstract. Let $G = (V, E)$ be a given (directed) graph in which every edge is with a cost and a delay that are nonnegative. The k-disjoint restricted shortest path (kRSP) problem is to compute k (edge) disjoint minimum cost paths between two distinct vertices $s, t \in V$, such that the total delay of these paths are bounded by a given delay constraint $D \in \mathbb{R}_0^+$. This problem is known to be NP-hard, even when $k = 1$ [4]. Approximation algorithms with bifactor ratio $(1+\frac{1}{r}, r(1+\frac{2(\log r+1)}{r})(1+\epsilon))$ and $(1+\frac{1}{r}, r(1+\frac{2(\log r+1)}{r}))$ have been developed for its special case when $k = 2$ respectively in [11] and [3]. For general k, an approximation algorithm with ratio $(1, O(\ln n))$ has been developed for a weaker version of kRSP, the k bi-constraint path problem of computing k disjoint st-paths to satisfy the given cost constraint and delay constraint simultaneously [7].

In this paper, an approximation algorithm with bifactor ratio (2, 2) is first given for the kRSP problem. Then it is improved such that for any resulted solution, there exists a real number $0 \le \alpha \le 2$ that the delay and the cost of the solution is bounded, respectively, by α times and $2 - \alpha$ times of that of an optimal solution. These two algorithms are both based on rounding a basic optimal solution of a LP formula, which is a relaxation of an integral linear programming (ILP) formula for the kRSP problem. The key observation of the two ratio proofs is to show that, the fractional edges of a basic solution to the LP formula will compose a graph in which the degree of every vertex is exactly 2. To the best of our knowledge, it is the first algorithm with a single factor polylogarithmic ratio for the kRSP problem.

Keywords: LP rounding, flow theory, k-disjoint restricted shortest path problem, bifactor approximation algorithm.

1 Introductions

This paper addresses on the k restricted shortest path problem, whose definition is formally as in the following:

* This research was supported by Natural Science Foundation of China (#61300025), Doctoral Funds of Ministry of Education of China for Young Scholars (#20123514120013), Natural Science Foundation of Fujian Province (#2012J05115), and Fuzhou University Development Fund (2012-XQ-26).

J. Chen, J.E. Hopcroft, and J. Wang (Eds.): FAW 2014, LNCS 8497, pp. 94–104, 2014.

Definition 1. *(The k restricted shortest path problem, kRSP) Let $G = (V, E)$ be a (directed) graph with a pair of distinct vertices $s, t \in V$. Assume that $c : E \to \mathbb{R}_0^+$ and $d : E \to \mathbb{R}_0^+$ are a cost function and a delay function on the edges of E respectively.* The k restricted shortest path problem *is to compute k disjoint st-paths $P_1, \ldots, P_k$, such that $\sum_{i=1,\ldots,k} c(P_i)$ is minimized while $\sum_{i=1,\ldots,k} d(P_i) \leq D$ holds for a given delay bound $D \in \mathbb{R}_0^+$.*

The kRSP problem has broad applications in industry, e.g., end-to-end video transmission with delay constraints, construction of minimum cost time-sensitive survivable networks, design of minimum cost fault tolerance systems subjected to given energy consumption constraint (or other additive constraints) and etc. Before the technique paragraphs, we would like to give the statement of the bifactor approximation algorithms for the kRSP problem first: An algorithm $\mathcal{A}$ is a bifactor (α, β)-approximation for the kRSP problem iff for every instance of kRSP, $\mathcal{A}$ computes k disjoint st-paths whose delay sum and cost sum are bounded by αD and $\beta c(OPT)$ respectively, where OPT is an optimum solution to the kRSP problem and $c(OPT) = \sum_{e \in OPT} c(e)$. We shall use bifactor $(1, \beta)$-approximation and β-approximation interchangeably in the text while no confusion arises.

1.1 Related Work

The kRSP problem have been studied for some fixed positive integral k. When $k = 1$, the problem becomes the restricted shortest path problem (RSP) of finding a single shortest path that satisfies a given QoS constraint. The RSP problem is known as one of Karp's 21 NP-hard problems [4] and admits full polynomial time approximation scheme (FPTAS) [9]. As a generalization of the RSP problem, the single Multiple Constraint Path (MCP) problem of computing a path subjected to multiple given QoS constraints is still attracting interest in the research community. For MCP, the $(1+\epsilon)$-approximation developed by Xue et al [16,10] is the best result in the current state of the art. When $k = 2$, approximation algorithms with bifactor ratio $(1 + \frac{1}{r}, r(1 + \frac{2(\log r+1)}{r})(1+\epsilon))$ and $(1 + \frac{1}{r}, r(1 + \frac{2(\log r+1)}{r}))$ have been developed respectively in [11] and [3]. To the best of our knowledge, there exists no non-trivial approximation that strictly obeys the delay constraint in the literature. However, for general k, the author, together with Shen and Liao, have developed approximation algorithms with bifactor ratio $(1, O(\ln n))$ for a weaker version of the kRSP problem, namely the k bi-constraint path problem, in which the goal is to compute k disjoint paths satisfying a given cost constraint and a given delay constraint simultaneously [7].

There are also some other interesting results that addressed on other special cases of the kRSP problem. When all edges are with delay 0, this problem becomes the min-sum problem of computing k disjoint paths with the total cost minimized, which is known polynomial solvable [13,14]. The min-min and min-max problems are two problems which are close related to the min-sum problem. The former problem is to find two paths with the length of the shorter one minimized, while the latter is to make the length of the longer one minimized. Our previous work, together with Xu et al's and Bhatia et al's [6,15,2], show that the min-min problem is NP-complete and doesn't admit K-approximation for any $K \geq 1$. Moreover, the edge-disjoint min-min problem remains NP-complete and admits no polynomial time approximation scheme in planar digraphs [5]. The min-max problem is also NP-complete. But unlike the min-min problem, it admits a best possible approximation ratio 2 in digraphs [8], which can be achieved immediately by employing Suurballe and Tarjan's algorithm for the min-sum problem [13,14]. In addition, as a variant of the min-max problem, the length bounded disjoint path problem of computing two disjoint paths whose lengths are both bound a given constraint, is also known NP-complete [8].

1.2 Our Technique and Results

In this paper, we first give a LP formula for the kRSP problem. Then by rounding the value of fractional edges of a solution to the LP formula, two approximation algorithms are developed. The first algorithm uses traditional rounding method, and computes solutions with a bifactor ratio of $(2, 2)$. The second algorithm improves the first rounding approach to an approximation with a pseudo ratio of $(\alpha,\, 2 - \alpha)$ for $0 \leq \alpha \leq 2$. That is, for any output solution of the improved approximation algorithm, there always exists $0 \leq \alpha \leq 2$, such that the delay and cost of the solution are bounded by α times and $2-\alpha$ times of that of an optimal solution respectively. By extending the technique in [7], the approximation ratio can be further improved to $(1,\, \ln n)$. To the best of our knowledge, this is the first approximation algorithm with a polylogarithmic ratio for the kRSP problem. Note that the extension of the technique in [7] is non-trivial, since we do not know the bound of the cost sum. Due to the length limitation, this paper will omit the details of the extension.

Like other rounding algorithms, the tricky task is to show that the round-up solutions are feasible for the kRSP problem, as the major results of this paper. The basic idea is to show that the fractional edges of a basic solution to the given LP formula will compose a graph, in which the degree of every vertex is exactly 2. Based on this observation, we show that the round-up edges in the first algorithm could collectively k-connect s and t. The correctness proof of the second algorithm follows a similar line to the first one, but requires a more sophisticated ratio proof because of its more complicated rounding method.

The following paragraphs are organized as follows: Section 2 gives the LP-rounding algorithm and shows that the resulting solution is with a bifactor ratio(2, 2); Section 3 gives the proof of the correctness of the algorithm; Section 4 improves the ratio (2, 2) to a pseudo ratio $(\alpha,\, 2 - \alpha)$ for $0 \leq \alpha \leq 2$ and then to $(1,\, \ln n)$; Section 5 concludes this paper.

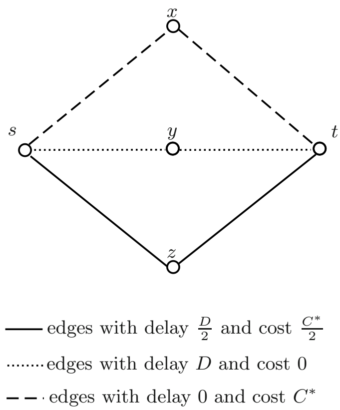

Fig. 1. $\{x_e = \frac{1}{2} | e \in E(G) \setminus \{e(s, z), e(z, t)\}\}$ is a basic optimum solution to LP (1) wrt the given graph G and $k = 1$. It is worth to note that this solution is not integral. $\{x_e = \frac{1}{3} | e \in E(G)\}$ is an optimum solution to LP (1) over this instance, but not a basic optimum solution.

2 An (2, 2)-Approximation Algorithm for the kRSP Problem

The linear programming (LP) formula for the kRSP problem is formally as in the following:

$$\min \sum c(e)x_e \tag{1}$$

subject to

$$\sum_{e\in\delta^+(v)} x_e - \sum_{e\in\delta^-(v)} x_e = \begin{cases} k & \text{for } v = s \\ 0 & \text{for } v \in V \setminus \{s,t\} \end{cases} \tag{2}$$

$$\sum_{e\in E} x_e d(e) \le D \tag{3}$$

$$\forall e \in E(G): \qquad 0 \le x_e \le 1 \tag{4}$$

where $\delta^+(v)$ and $\delta^-(v)$ denotes the set of edges leaving and entering v in G respectively. Before the technique paragraphs, we would like first do some discussion on this LP formula. If $x_e \in \{0, 1\}$, the above will be exactly the integral linear programming (ILP) formula for the kRSP problem. Moreover, with the relaxation over x_e (i.e. x_e is required to satisfy only Inequality (4) instead of $x_e \in \{0, 1\}$), a basic solution to the LP formula remains integral (i.e. , x_e remains integral for each e) for some special cases. When the delay constraint (i.e. Inequality (3)) is removed, any basic optimum solution to LP (1) would be exactly a set of st-paths with minimum cost. The reason is that,

Algorithm 1. A LP-rounding algorithm for the k-RSP problem.

Input: Graph G, distinct vertices s and t, a delay bound $D \in \mathbb{R}_0^+$, a cost function $c(e)$ and a delay function $d(e)$;
Output: k disjoint st-paths.

1. $E_{SOL} \leftarrow \emptyset$;
2. Solve LP (1) and get a basic optimum solution χ by Karmarkar's algorithm [12];
3. **For** each x_e in χ **do**
   (a) **if** $x_e = 1$ **then** $E_{SOL} \leftarrow \{e\} \cup E_{SOL}$
   (b) **if** $\frac{1}{2} \leq x_e < 1$ **then**
      i. Round the value of x_e to 1;
      ii. $E_{SOL} \leftarrow \{e\} \cup E_{SOL}$.

   /* As shown later in Theorem 2, s is k-connected to t by edges of E_{SOL} after the execution of Step 2.*/
4. **For** each e in E_{SOL} **do**
   if s is k connected to t in $E_{SOL} \setminus e$ **then**
   $E_{SOL} \leftarrow E_{SOL} \setminus \{e\}$;
   /* Remove redundant edges of E_{SOL}, such that E_{SOL} exactly compose k-disjoint paths. */
5. Return E_{SOL} as the k-disjoint paths.

in this case, LP (1) is totally unimodular [1,12], so any basic optimum solution to LP (1) must be integral. But with the delay constraint, LP (1) is no longer totally unimodular, and hence a basic optimum solution to LP (1) is no longer integral even when $k = 1$ (as depicted in Figure 1). As the first main result of this paper, we show that a basic optimum solution to LP (1) still acquires an interesting property as stated below:

Lemma 2. *The set of edges with $x_e \geq \frac{1}{2}$ in a basic optimum solution to LP (1) can collectively provide k-connectivity between s and t.*

Since Lemma 2 is one of our major results and its proof is the most tricky part of this paper, its proof is deferred to the next section. Based on this lemma, the key idea of our algorithm is simply as below: compute a basic optimum solution to LP (1), and round x_e to 1 for every edge e with $x_e \geq \frac{1}{2}$ in the solution. Then the algorithm outputs a set of edges with $x_e = 1$ as a solution to the kRSP problem. The detailed algorithm is formally as in Algorithm 1.

Step 2 of Algorithm 1 takes $O(n^{3.5} \log T)$ time to run Karmarkar's algorithm [12], where T is the maximum absolute value of the input numbers. Step 3 and Step 4 take $O(m)$ time to round the edges of E_{SOL} and remove the redundant edges of E_{SOL} respectively. Hence, the time complexity of Algorithm 1 is $O(n^{3.5} \log T)$ in the worst case.

For the approximation ratio of Algorithm 1, obviously E_{SOL} contains only edges with $x_e \geq \frac{1}{2}$ in the basic optimum solution χ, so we have

$$c(E_{SOL}) = \sum_{e \in E(G),\, x_e \geq \frac{1}{2}} c(e) \leq \sum_{e \in E(G),\, x_e \geq \frac{1}{2}} 2x_e c(e) \leq 2 \sum_{e \in E(G)} x_e c(e). \qquad (5)$$

Then because $\sum_{e\in E(G)} x_e c(e)$ is the cost of an optimum solution to LP (1), it is not larger than that of an optimum solution to the kRSP problem. This yields

$$\sum_{e\in E(G)} x_e c(e) \leq \sum_{e\in OPT} c(e), \tag{6}$$

where OPT is an optimum solution to the kRSP problem. Combining Inequality (5) and Inequality (6), we have $c(E_{SOL}) \leq 2c(OPT)$. Similarly, we have $d(E_{SOL}) \leq 2d(OPT)$. Therefore, the time complexity and approximation ratio of Algorithm 1 are as in the following theorem:

Theorem 3. *Algorithm 1 outputs E_{SOL}, a solution to the kRSP problem, in $O(n^{3.5}\log T)$ time. The cost and delay of E_{SOL} are at most two times of that of an optimum solution to the kRSP problem, where T is the maximum absolute value of the input numbers.*

3 Proof of Lemma 2

Before the technique paragraphs, we would like first to give some definitions. Let χ be a basic optimum solution to LP (1), and G_χ be the graph composed by the edges with $x_e > 0$ in χ. We say $e \in G_\chi$ is a full edge if and only if $x_e = 1$. Let Z be the set of full edges of G_χ and $E_{res} = G_\chi \setminus Z$ be the set of edges with $0 < x_e < 1$ in the solution to LP (1). The notation E_{res}or E_{SOL} also denotes the graph composed by the edges of E_{res} or E_{SOL} while no confusion arises.

To prove Lemma 2, the key idea is first to show that each vertex in graph E_{res} is exactly incident with two edges of E_{res} (as Lemma 4). Based on this property, we then show that the edges with $1 > x_e \geq \frac{1}{2}$, together with the edges with $x_e = 1$, collectively k connect s and t. Therefore, we shall first focus on the properties of only the edges in E_{res}, leaving the edges with $x_e = 0$ and $x_e = 1$ for a moment. To make the proof brief, we consider the following residual LP formula instead of LP (1), which is LP (1) except that the edges with $x_0 = 0$ or $x_0 = 1$ of G_χ are removed.

$$\min \sum_{e\in E_{res}} c_e x_e \tag{7}$$

subject to

$$\sum_{e\in\delta^+_{E_{res}}(v)} x_e - \sum_{e\in\delta^-_{E_{res}}(v)} x_e - \sum_{e\in\delta^+_Z(v)} 1 + \sum_{e\in\delta^-_Z(v)} 1 = \begin{cases} k & \text{for } v = s \\ 0 & \text{for } v \in V_{res} \setminus \{s,t\} \end{cases} \tag{8}$$

$$\sum_{e\in E_{res}} x_e d(e) \leq kD - \sum_{e\in Z} d(e) \tag{9}$$

$$\forall e \in E_{res}: \qquad 0 < x_e < 1 \tag{10}$$

It is easy to see that $\{x_e|e \in E_{res}\}$ is a solution to the above LP formula. Let χ_{res} be a basic solution to LP (7). Then χ_{res} is χ except that the x_es of value 0 or 1 are removed. Let $\mathcal{A}$ be the whole constraint matrix of LP (7), $\mathcal{A}_G(v)$ be the row corresponding to v, and $\mathcal{A}_G(D)$ be the row corresponding to the delay constraint, i.e. Inequality (9). Assuming that V_{res} is the set of the vertices of E_{res}, we denote the vector space spanned by the vectors $\mathcal{A}_G(v)$, $v \in V_{res}$, by $Span(V_{res})$, and the space spanned by vectors $\mathcal{A}_G(D)$ together with $\mathcal{A}_G(v)$, $v \in V_{res}$, by $Span(D \cup V_{res})$.

Lemma 4. *Every vertex of V_{res} is incident to exactly two edges in graph E_{res}.*

Proof. Firstly, we shall show that every vertex of $v \in V_{res}$ is incident to at least two edges in E_{res}. Clearly, v must be incident with at least one edge of E_{res}, such that it can belong to V_{res}. Suppose only one edge of E_{res} joins v. Then on one hand, because $0 < x_e < 1$ for every $e \in E_{res}$, the degree of v is not integral. On the other hand, following the LP formula (7), every vertex of V_{res}, including $v \in V_{res}$, must be with an integral degree in the basic optimum solution χ_{res}. Hence a contradiction arises. Therefore, every vertex of V_{res} is incident to at least 2 edges.

Secondly, we shall show that every vertex in V_{res} is incident to at most two edges in E_{res}. Suppose there exists a vertex $v_0 \in V_{res}$ with a degree of at least 3, i.e. $|\delta^-(v_0)| + |\delta^+(v_0)| \geq 3$. Then, on one hand, since every vertex of V_{res} must be incident to at least 2 edges, we have:

$$|E_{res}| = \frac{1}{2} \sum_{v \in V_{res}} |\delta^-(v)| + |\delta^+(v)| \geq \frac{1}{2}\Big(\sum_{v \in V_{res} \setminus \{v_0\}} 2 + \sum_{v=v_0} 3\Big) > |V_{res}|.$$

On the other hand, we can show that $|V_{res}| \geq |E_{res}|$ must hold and hence obtain a contradiction. The proof is as below. Since χ_{res} is a basic optimum solution and there are $|E_{res}|$ edges with $x_e > 0$ in χ_{res}, the dimension of $Span(D \cup V_{res})$ is $|E_{res}|$. Then, since $\sum_{e \in E_{res}} x_e d(e) = D - \sum_{e \in Z} d(e)$ may holds for Inequality (9), the dimension of $Span(V_{res})$ is at least $|E_{res}| - 1$. Then, because the $|V_{res}|$ rows of the constraint matrix for $Span(V_{res})$ contains at most $|V_{res}| - 1$ linear independent vector, $|V_{res}| - 1 \geq |E_{res}| - 1$ holds, and hence $|V_{res}| \geq |E_{res}|$. This completes the proof.

Let V_{full} be the set of vertices of $V(E_{res}) \cap V(E_{full})$. Then, V_{full} are with degree 2 or -2 in E_{res}. Following Lemma 4, the degree of $v \in V_{res} \setminus V_{full}$ is 0. So for each $v \in V_{res} \setminus V_{full}$, E_{res} contains exactly an edge entering and the other edge leaving v. Therefore, E_{res} is actually a set of paths between the vertices of V_{full}. Further, we have the following lemma:

Lemma 5. *The edges of E_{res} compose a set of paths P_{res} between the vertices of V_{full}. For every $v \in V_{full}$, P_{res} contains exactly two paths leaving v or entering v. Moreover, any two paths of P_{res} share no common edge.*

Proof. Following Lemma 4, for every $v \in V_{full}$, E_{res} contains exactly two edges which either enter or leave v, while for $v \in V_{res} \setminus V_{full}$, E_{res} contains exactly one edge entering v and one another edge leaving v. So following flow theory, for each $v \in V_{full}$, E_{res} contains exactly two paths either entering or leaving v. In addition, E_{res} contains no cycles, since E_{res} is a subgraph of G_χ and G_χ contains no cycles because G_χ is a minimum cost fractional flow from s to t. Therefore, the edges of E_{res} compose exactly the paths of P_{res}.

It remains to show that the paths of P_{res} are edge-disjoint. Since every $v \in V_{res} \setminus V_{full}$ is with degree 2, v can appear on only one path of P_{res}. That is, any two distinct paths of P_{res} cannot go through a common vertex of $V_{res} \setminus V_{full}$. So the paths of P_{res} are internal vertex-disjoint, and hence edge-disjoint. This completes the proof.

Now we are to show that the edges with $x_e \geq \frac{1}{2}$ in G_χ could provide k connectivity for s and t. Remind that the output of Algorithm 1, E_{SOL}, is exactly the set of edges with $x_e \geq \frac{1}{2}$ in G_χ. That is, E_{SOL} is equivalently the edges with $x_e \geq \frac{1}{2}$ in E_{res} together with the full edges of Z. Suppose Lemma 2 is not true, then there must exist $k-1$ edges, say $e_1, \ldots, e_{k-1}$, which separate s and t in E_{SOL}. Then according to Lemma 6 as below, the flow between s and t wrt the solution χ to LP (1) is at most of value $k-1$. This contradicts with the fact that the flow wrt χ is of value k, and completes the proof of Lemma 2.

Lemma 6. *Assume $e_1, \ldots, e_{k-1}$ separate s and t in E_{SOL}, then the flow between s and t of the solution to LP (1) is at most of a value $k-1$.*

Proof. Let $G_s \supset \{s\}$ be the component of $E_{SOL} \setminus \{e_1, \ldots, e_{k-1}\}$ which contains s. Then $\{e_1, \ldots, e_{k-1}\}$ would separate G_s and $E_{SOL} \setminus G_S$ in graph E_{SOL}. W.l.o.g., assume $e_1, \ldots, e_h$ are with $x_e = 1$, while $e_{h+1}, \ldots, e_{k-1}$ are with $\frac{1}{2} \leq x_e < 1$. Let $p^1_{h+1}, \ldots, p^1_{k-1}$ be the paths between vertices of V_{full} in $E_{res} \subseteq G_\chi$. W.l.o.g., assume that $e_i \in p^1_i$ for each i. Following Lemma 5, there exists exactly one another path, say p^2_j, which leav*es* the same vertex as p^1_j in E_{res}. Then the set of p^1_js and p^2_js, i.e. the set of paths $\{p^i_j | i \in \{1,2\}, j \in \{h+1, \ldots, k-1\}\}$, can only provide a flow of value at most $k-1$ together with $\{e_1, \ldots, e_h\}$.

It remains to show that there exists neither a full edge outside $\{e_1, \ldots, e_h\}$, nor a path of P_{res} outside $\{p^i_j | i \in \{1,2\}, j \in \{h+1, \ldots, k\}\}$ that leaves G_s in graph G_χ. Suppose otherwise, as the two cases analyzed below, such a full edge or a path must belong to E_{SOL}, and hence it would connect G_s and $E_{SOL} \setminus G_s$ in $E_{SOL} \setminus \{e_1, \ldots, e_{k-1}\}$. This contradicts with the assumption that edges of$\{e_1, \ldots, e_k\}$ separate G_s and $E_{SOL} \setminus G_s$ in E_{SOL}, and completes the proof.

1. Suppose there exists a full edge $e \notin \{e_1, \ldots, e_h\}$ in G_χ that leaves G_s. Since e is a full edge, then $x_e = 1$ and $e \in E_{SOL}$ holds. Hence, $e_1, \ldots, e_k$ can not separate G_s and $E_{SOL} \setminus G_s$, because e connects them.
2. Suppose there exists a path $p \in P_{res} \setminus \{p^i_j | i \in \{1,2\}, j \in \{h+1, \ldots, k\}\}$ that leaves G_s at v. Following Lemma 5, there must be exactly one another path p' that leaves v in G_χ. Then either flow p or p' is with value at least $\frac{1}{2}$. That is, the edges of either p or p' would belong to E_{SOL} and connect

G_s and $E_{SOL} \setminus G_s$. So removal of $\{e_1, \ldots, e_k\}$ can not disconnect G_s and $E_{SOL} \setminus G_s$ in E_{SOL}.

4 An $(\alpha, 2-\alpha)$-Approximation Algorithm for the kRSP Problem

In Section 2, Algorithm 1 adds every edge with $x_e \geq \frac{1}{2}$ to the solution. However, not all the edges with $x_e \in [\frac{1}{2}, 1)$ are good choices for constructing a solution to the kRSP problem. This section will give an improved rounding algorithm which selects the edges of the solution more carefully, such that the algorithm is with a pseudo approximation ratio $(\alpha, 2-\alpha)$. Thus, for any output solution of the algorithm, there always exists $0 \leq \alpha \leq 2$, such that the delay and cost of the solution are bounded by α times and $2-\alpha$ times of that of an optimal solution respectively.

Assume that χ is a basic optimum solution to LP (1). The main idea of the improved rounding algorithm is to select the edges with less cost and delay, rather than to select the edges with $\frac{1}{2} \leq x_e < 1$. To do this, the algorithm combines the cost and delay as one new cost and then selects a set of edges, which are with the new cost sum minimized and provide k connectivity between s and t together with the edges of $x_e = 1$ in χ. Let $c(\chi)$ and $d(\chi)$ be the cost and delay of the basic optimum solution to LP (1). The new mixed cost for every edge is $b(e) = \frac{c(e)}{c(\chi)} + \frac{d(e)}{d(\chi)}$.

According to Lemma 5, the edges of E_{res} compose exactly a set of internal vertex disjoint paths, say $P_{res} = \{p_j^i | i \in \{1, 2\}, j \in \{1, \ldots, h\}\}$, where p_j^1 and p_j^2 leaves a same vertex of V_{full}. Then the task is now to choose h paths $\{p_j^{i_j} | j \in \{1, \ldots, h\}\}$ from P_{res} to provide the k connectivity between s and t. According to Lemma 5, we could divide P_{res} into two path sets $\mathcal{P}_1$, $\mathcal{P}_2$, such that every two paths in $\mathcal{P}_i$ shares no common vertex. Then following the same line of the proof of Lemma2, it can be shown that $Z \cup E(\mathcal{P}_i)$ provides k-connectivity between s and t for either $i = 1$ or $i = 2$. Therefore, the main idea of our algorithm is to divide P_{res} into two path sets $\mathcal{P}_1$, $\mathcal{P}_2$, and then select $\mathcal{P}_i$ with smaller $\sum_{e \in E(\mathcal{P}_i)} b(e)$ for $i = 1, 2$. Formally, the algorithm is as below:

It remains to show the approximation ratio of the algorithm, which is stated as follows:

Theorem 7. *There exists a real number $0 \leq \alpha \leq 2$, such that the delay and cost of E_{SOL} are bounded by α and $2-\alpha$ times of that of the optimum solution of the kRSP problem.*

Proof. Clearly, $b(E_{SOL}) = b(Z) + \beta b(E(\mathcal{P}_i)) + (1-\beta) b(E(\mathcal{P}_i))$ holds. Then because $\sum_{e \in E(\mathcal{P}_i)} b(e) \leq \sum_{e \in E(\mathcal{P}_{3-i})} b(e)$, we have

$$b(E_{SOL}) \leq b(Z) + \beta \sum_{e \in E(\mathcal{P}_i)} b(e) + (1-\beta) \sum_{e \in E(\mathcal{P}_{3-i})} b(e) \leq \sum_{e \in G_\chi} x_e b(e) \leq 2. \quad (11)$$

Algorithm 2. A LP-rounding algorithm for the k-RSP problem.

Input: E_{res} with new cost $b(e)$, a basic optimum solution χ to LP (1);
Output: k disjoint st-paths.

1. $E_{SOL} \leftarrow Z$; /*E_{SOL} is initially the set of edges with $x_e = 1$ in G_χ. */
2. Divide P_{res} into two path sets $\mathcal{P}_1$, $\mathcal{P}_2$, such that every two paths in $\mathcal{P}_i$ shares no common vertex for $i = 1, 2$;
3. Select i, $i \in \{1, 2\}$, such that $\mathcal{P}_i$ is with smaller new cost sum, i.e., $\sum_{e\in E(\mathcal{P}_i)} b(e) \leq \sum_{e\in E(\mathcal{P}_{3-i})} b(e)$, $i \in \{1, 2\}$;
4. Return $E_{SOL} \leftarrow E(\mathcal{P}_i) \cup E_{SOL}$.

Assume that the delay-sum of E_{SOL} is α times of D, where α is a real number and $0 \leq \alpha \leq 2$. Then $b(E_{SOL}) = \sum_{e\in E_{SOL}} b(e) = \alpha + \frac{c(E_{SOL})}{c(\chi)}$ holds. From Inequality (11), $\alpha + \frac{c(E_{SOL})}{c(\chi)} \leq 2$ holds. That is, $c(E_{SOL}) \leq (2-\alpha)c(\chi) \leq (2-\alpha)C^*$. This completes the proof.

By extending the technique of combining cycle cancelation and layer graph as in [7], the approximation ratio can be improved to $(1, \ln n)$. However, we do not know the value of $c(OPT)$, so we can only construct an auxiliary graph based on the delay of edges instead of cost. The auxiliary graph constructed is no longer a layer graph, but the costs of the edges therein are nonnegative. Therefore, we can compute minimum delay-to-cost cycle in the residual graph of G by computing shortest paths in the constructed auxiliary graph, and improve the output of Algorithm 2 by using the cycle cancelation method against the minimum delay-to-cost cycle repeatedly, until the solution satisfies the delay constraint strictly. To the best of our knowledge, this is the first non-trivial approximation algorithm with single factor polylogarithmic ratio for the kRSP problem.

5 Conclusion

This paper investigated approximation algorithms for the k-restricted shortest paths (kRSP) problem. As the main contribution, this paper first developed an improved approximation algorithm with bifactor ratio (2, 2) by rounding a basic optimum solution to the proposed LP formula of the kRSP problem. The algorithm was then improved by choosing the round-up edges more carefully, such that for any output solution there exists $0 \leq \alpha \leq 2$ that the delay and the cost of the solution are bounded, respectively, by α and $2-\alpha$ times of that of the optimum solution. This ratio can be further improved to $(1, \ln n)$ by extending the technique of [7]. To the best of our knowledge, this is the first approximation with single factor polylogarithm ratio for the kRSP problem.

References

1. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network flows: theory, algorithms, and applications (1993)
2. Bhatia, R., Kodialam, M.: TV Lakshman. Finding disjoint paths with related path costs. Journal of Combinatorial Optimization 12(1), 83–96 (2006)
3. Chao, P., Hong, S.: A new approximation algorithm for computing 2-restricted disjoint paths. IEICE Transactions on Information and Systems 90(2), 465–472 (2007)
4. Garey, M.R., Johnson, D.S.: Computers and intractability. Freeman, San Francisco (1979)
5. Guo, L., Shen, H.: On the complexity of the edge-disjoint min-min problem in planar digraphs. Theoretical Computer Science 432, 58–63 (2012)
6. Guo, L., Shen, H.: On finding min-min disjoint paths. Algorithmica 66(3), 641–653 (2013)
7. Guo, L., Shen, H., Liao, K.: Improved approximation algorithms for computing k disjoint paths subject to two constraints. In: Du, D.-Z., Zhang, G. (eds.) COCOON 2013. LNCS, vol. 7936, pp. 325–336. Springer, Heidelberg (2013)
8. Li, C.L., McCormick, T.S., Simich-Levi, D.: The complexity of finding two disjoint paths with min-max objective function. Discrete Applied Mathematics 26(1), 105–115 (1989)
9. Lorenz, D.H., Raz, D.: A simple efficient approximation scheme for the restricted shortest path problem. Operations Research Letters 28(5), 213–219 (2001)
10. Misra, S., Xue, G., Yang, D.: Polynomial time approximations for multi-path routing with bandwidth and delay constraints. In: INFOCOM 2009, pp. 558–566. IEEE (2009)
11. Orda, A., Sprintson, A.: Efficient algorithms for computing disjoint QoS paths. In: INFOCOM 2004, vol. 1, pp. 727–738. IEEE (2004)
12. Schrijver, A.: Theory of linear and integer programming. John Wiley & Sons Inc. (1998)
13. Suurballe, J.W.: Disjoint paths in a network. Networks 4(2) (1974)
14. Suurballe, J.W., Tarjan, R.E.: A quick method for finding shortest pairs of disjoint paths. Networks 14(2) (1984)
15. Xu, D., Chen, Y., Xiong, Y., Qiao, C., He, X.: On the complexity of and algorithms for finding the shortest path with a disjoint counterpart. IEEE/ACM Transactions on Networking 14(1), 147–158 (2006)
16. Xue, G., Zhang, W., Tang, J., Thulasiraman, K.: Polynomial time approximation algorithms for multi-constrained qos routing. IEEE/ACM Transactions on Networking 16(3), 656–669 (2008)