Network Utility Maximization with Path Cardinality Constraints

Yingjie Bi*, Chee Wei Tan† and Ao Tang*
*School of Electrical and Computer Engineering, Cornell University
†Department of Computer Science, City University of Hong Kong
yb236@cornell.edu, cheewtan@cityu.edu.hk, atang@ece.cornell.edu

Abstract—**In this paper, we study network utility maximization over both routing choice and path rate assignment for any given path cardinality constraint. We provide a novel convex relaxation, which leads to a randomized algorithm with performance guarantees. The new relaxation also enables distributed algorithm design and allows us to obtain performance estimation for nonconvex routing optimization problems that is significantly better than previous work based on the multipath routing relaxation. Convergence and performance of the proposed randomized algorithm are characterized theoretically and further illustrated numerically through examples to demonstrate its superiority over existing work.**

I. Introduction

Many current Internet routing protocols are single-path based. For example, the Open Shortest Path First (OSPF) protocol allows a user (source-destination pair) to use only one path from the source to the destination, with the exception that traffic may split evenly among equal-cost paths. Recently, the IETF has published the *multipath specification* as an experimental standard in RFC 6824 to promote the feasibility of multipath TCP. It is expected that the use of multipath routing will increase the resource usage efficiency and provide better load balancing [1].

More generally, the number of paths (W) allowed, i.e., the path cardinality constraint, greatly affects the attainable performance of a given routing scheme, theoretical tractability of routing optimization, as well as actual implementation complexity. On one hand, using all the available paths (multipath routing, $W = \infty$) can potentially achieve the best possible performance and also make routing optimization problems convex and tractable. But it is often too expensive for implementation in terms of protocol overhead. On the other hand, allowing only one path between each source-destination pair (single-path routing, $W = 1$) is much easier to realize and has been the dominating practice. Its performance nevertheless is usually suboptimal and the corresponding optimization formulations and algorithms are typically very hard to analyze due to the intrinsic nonconvexity in the problem structure.

Take the network utility maximization over joint congestion control and routing as an example, which is also the main topic of this paper. If multipath is allowed, the corresponding routing optimization problems are convex and tractable [2], [3]. However, due to various difficulties such as out-of-order arrival of packets, the actual use of multipath TCP only starts to gain attention recently [4], [5]. If we only look at congestion control plus single-path routing, the authors in [6] showed that this problem is in general NP-hard. They proposed a relaxation based on Lagrange duality that is akin to a multipath TCP utility maximization framework, and a dual-based algorithm to maximize the system utility when the Lagrange duality gap is zero. The duality gap thus measures the *cost of not splitting*. The authors in [7] further studied this duality gap characterization using ideas from sparse recovery, and obtained sufficient conditions under which the gap is zero. In the event that the duality gap is non-zero, the authors proposed algorithms to project the solution obtained by solving the multipath TCP problem to a feasible single-path TCP solution. Besides the network utility maximization framework, the benefit of using multipath has also been studied within the traffic engineering framework [8].

In all these studies, an interesting finding [7], [8] has been that in terms of performance, single-path routing may not be too far away from multipath routing, especially when the network size and the number of users are both large. This motivates a practical implementation to consider using only a few number of paths. Typically, the number of allowed paths W is small, and hence we call this the *sparse routing* in the paper. In fact, a natural question that generalizes the above single-path and multipath routing studies arises: assuming W ($1 \le W \le \infty$) paths are allowed for each user, what is the optimal routing performance and how to achieve it.

This paper investigates the above general problem and makes two contributions:

- We study network utility maximization over joint congestion control and routing with path cardinality constraints, which is generally a nonconvex problem. To characterize its performance, we propose a convex relaxation of this problem that is significantly tighter than the standard multipath relaxation.
- Based on the new convex relaxation, we develop a novel dual-based distributed algorithm that can be interpreted as *sparse multipath* TCP/IP joint congestion control and routing algorithm. Even with path cardinality constraints, it can load balance a large portion of the traffic with a small number of paths. We demonstrate the effectiveness

This work was supported by the Research Grants Council of Hong Kong under RGC CityU 11212114, by the Huawei Technologies Co. Ltd. under the CityU-Huawei joint Project 9231159 and by the Office of Naval Research, Arlington under N00014-11-1-0131.

978-1-4673-9953-1/16/$31.00 ©2016 IEEE

of this solution both theoretically and numerically.

This paper is organized as follows. We introduce the system model and formulate the network utility maximization problem in Section II. In Section III, we propose a convex relaxation to this NP-hard problem. In Section IV, we focus on deriving stronger results in the linear utility special case for throughput maximization. In Section V, we propose a dynamic dual algorithm to solve the convex relaxation in a distributed manner. In Section VI, we illustrate the numerical performance of our algorithm. We conclude the paper in Section VII.

II. Model and Notation

We summarize the key notations, especially those related to the vector norms, used in the paper below. For a vector $u = (u_1, \dots, u_n) \in \mathbb{R}^n$:

- The l^1 norm $\|u\|_1$ denotes $\sum_{k=1}^n |u_k|$.
- The infinity norm $\|u\|_\infty$ denotes $\max_{k=1,\dots,n} |u_k|$.
- $\|u\|_0$ denotes the number of nonzero entries in u.[1]

Let $(u_{[1]}, \dots, u_{[n]})$ be a rearrangement of $(u_1, \dots, u_n)$ sorted in nonincreasing order, i.e., $u_{[1]} \geq \cdots \geq u_{[n]}$, then the sum of W largest components in u

$$\sum_{k=1}^{W} u_{[k]}$$

is a convex function over u [9, Sec. 3.2.3]. When $u \geq 0$ and $W = 1$, this is just the infinity norm.

A network consists of L uni-directional links with positive and finite capacities $c = (c_1, \dots, c_L)^T$, which support N source-destination pairs or users indexed by i. There are K^i acyclic paths available for user i. Each path is a sequence of links over which the data of user i can flow to its destination. The paths of user i are represented by an $L \times K^i$ matrix R^i, where $R^i_{lk} = 1$ if path k of user i passes through link l, and $R^i_{lk} = 0$ otherwise. The overall routing matrix R is the $L \times K$ matrix defined by

$$R = \begin{pmatrix} R^1 & R^2 & \cdots & R^N \end{pmatrix}.$$

Here $K = \sum_{i=1}^N K^i$.

For each user i, x^i is a $K^i \times 1$ vector whose kth entry x^i_k denotes the sending rate of user i through its path k. Let

$$x = \begin{pmatrix} x^1 \\ x^2 \\ \vdots \\ x^N \end{pmatrix}$$

be the complete rate allocation. Each user i has a utility $U^i(\cdot)$ as a function of its total transmission rate $\|x^i\|_1$. We assume $U^i(\cdot)$ to be concave, increasing and continuous, which is the case in most TCP algorithms [10].

[1] Note that $\|u\|_0$ is not a vector norm. Our notation for the number of nonzero entries in u follows from the compressed sensing literature, and is commonly known as "l^0 norm".

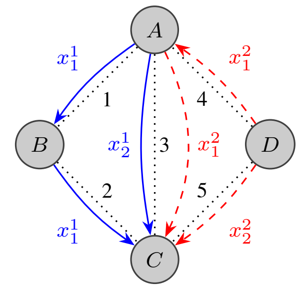

Fig. 1. A sample network with five links (dotted line) and two users. The first user, whose source is A and destination is C, has rate x^1_1 and x^1_2 on its two paths (solid line). The second user, whose source is D and destination is C, has rate x^2_1 and x^2_2 on its two paths (dashed line).

For example, as shown in Fig. 1, a five-link network supports two users, each of which has two available paths. The corresponding routing matrices are given by

$$R^1 = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad R^2 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

We now state the network utility maximization problem over joint congestion control and routing with path cardinality constraints. The goal is to maximize the aggregate utility of all users with the restriction that each user can only split its data into at most W flows. Here, W is relatively small enough such that the rate allocation x^i is a *sparse* vector. This *sparse routing* problem can be formulated as the following optimization problem:

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} U^i\left(\|x^i\|_1\right) \\ \text{s.t.} \quad & Rx \leq c, \\ & x \geq 0, \\ & \|x^i\|_0 \leq W, \quad \forall i = 1, \dots, N. \end{aligned} \tag{1}$$

Let opt_S denote the optimal value of (1). Note that when $W = 1$, we have the *single-path routing* special case studied in [6], [7].

By dropping the last N nonconvex constraints in (1), interestingly, this yields the *multipath routing* problem:

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} U^i\left(\|x^i\|_1\right) \\ \text{s.t.} \quad & Rx \leq c, \\ & x \geq 0. \end{aligned} \tag{2}$$

Indeed, (2) can also be viewed as a relaxation of (1), and its relaxation gap has been studied in [6], [7]. If x is a feasible solution to (2), define a new rate allocation y by

$$y^i_k = \begin{cases} x^i_k & \text{if } x^i_k \text{ is among the } W \text{ largest rates of user } i, \\ 0 & \text{otherwise.} \end{cases}$$

Then y is feasible for the original sparse routing problem (1). In the following, we say y is the *projection* of x to a sparse routing configuration.

In fact, the multipath relaxation (2) is hitherto the only known relaxation to (1) in the literature, but, in general, it can be very loose and algorithms based on the multipath relaxation cannot solve (1) distributedly or with guarantees. In Section III, we will provide stronger convex relaxation techniques to (1) that can even lead to distributed algorithms with guarantees.

III. Convex Relaxation

We start by considering a reformulation of (1) that essentially moves the nonconvexity from the path cardinality constraints to the objective function.

Lemma 1: The sparse routing problem (1) and the following problem (3) have the same optimal value:

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} U^i \left(\sum_{k=1}^{W} x^i_{[k]} \right) \\ \text{s.t.} \quad & Rx \le c, \\ & x \ge 0. \end{aligned} \tag{3}$$

Proof: Let γ be the optimal value of (3). The optimal solution to (1) is obviously feasible for (3), and the corresponding objective values are the same in both problems, which implies that $opt_S \le \gamma$. Conversely, if we have an optimal solution to (3), projecting it to a sparse routing configuration will give a feasible solution to (1) without changing the objective value. Thus, $\gamma \le opt_S$. ■

Next, we want to relax problem (3) in order to make its objective function concave. The objective function of (3) is separable for each user and the only constraints coupling the different users are from $Rx \le c$, so (3) is a special case of the separable problems studied in [11]. For each user i, define

$$f^i(x^i) = \begin{cases} U^i \left(\displaystyle\sum_{k=1}^{W} x^i_{[k]} \right) & \text{if } 0 \le x^i \le \|c\|_\infty, \\ -\infty & \text{otherwise.} \end{cases} \tag{4}$$

Then problem (3) can be rewritten as

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} f^i(x^i) \\ \text{s.t.} \quad & Rx \le c, \end{aligned} \tag{5}$$

because $x^i \le \|c\|_\infty$ in (4) is automatically implied by the constraints in (3).

For each function f^i in (5), its *concave envelope* $\hat{f}^i$ is a concave function defined by

$$\hat{f}^i = \inf \left\{ g | g \text{ is concave and } g(x^i) \ge f^i(x^i), \forall x^i \in \mathbb{R}^{K^i} \right\}.$$

Define

$$\rho^i = \sup_{x^i} \left\{ \hat{f}^i(x^i) - f^i(x^i) \right\},$$

where we consider $(-\infty) - (-\infty) = 0$. ρ^i measures the nonconcavity of the function f^i [12].

Replacing each function f^i by its concave envelope $\hat{f}^i$, we get the *convex relaxation* for sparse routing problem (1):

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} \hat{f}^i(x^i) \\ \text{s.t.} \quad & Rx \le c. \end{aligned} \tag{6}$$

The Lagrange dual problem of (6) is

$$\begin{aligned} \min \quad & \sup_x \left\{ \sum_{i=1}^{N} \hat{f}^i(x^i) + p^T(c - Rx) \right\} \\ \text{s.t.} \quad & p \ge 0. \end{aligned}$$

Remark 1: The multipath relaxation (2) can also be understood as replacing f^i by

$$g^i(x^i) = \begin{cases} U^i\left(\|x^i\|_1\right) & \text{if } x^i \ge 0, \\ -\infty & \text{otherwise.} \end{cases}$$

Since $U^i(\cdot)$ is increasing and concave, $g^i(x^i) \ge f^i(x^i)$ and g^i is also concave. Thus $g^i(x^i) \ge \hat{f}^i(x^i)$ and our convex relaxation (6) is always tighter than or at least the same as the multipath relaxation (2).

Solving the convex relaxation (6) can provide a suboptimal solution to the original sparse routing problem (1) within the bound given below.

Theorem 2 (modified from [11, Theorem 2]): Let $w \in \mathbb{R}^{K \times 1}$ be a random vector drawn from the uniform distribution on a unit sphere. Assume (x^*, p^*) is an optimal primal-dual pair for (6). Then with probability 1, the following problem

$$\begin{aligned} \min \quad & w^T x \\ \text{s.t.} \quad & Rx = Rx^*, \\ & \sum_{i=1}^{N} \left(\hat{f}^i(x^i) - \hat{f}^i(x^{*i}) \right) \ge p^{*T}(Rx - Rx^*) \end{aligned}$$

has a unique optimal solution $\hat{x}$, which is also an optimal solution to (6) and satisfies

$$opt_S - \sum_{i=1}^{N} f^i(\hat{x}^i) \le \sum_{i=1}^{\min\{N,L\}} \rho^i. \tag{7}$$

Here we assume users are sorted such that $\rho^1 \ge \cdots \ge \rho^N$.

Remark 2: It is important to note that not every optimal solution to (6) satisfies the bound (7). In fact, the auxiliary optimization problem in Theorem 2 can be viewed as using randomization to find such a desired optimal solution.

IV. Throughput Maximization

In general, calculating the concave envelope $\hat{f}^i$ can be hard. In this section, we restrict ourselves to the special case when $U^i(s) = s$, in which we can say something stronger and improve the convex relaxation (6). Maximizing a linear utility is equivalent to maximizing the total network throughput, and [13] analyzed the relationship between throughput maximization and fair rate allocation.

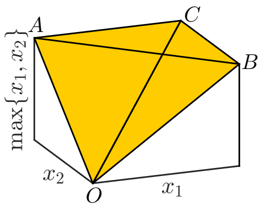

Fig. 2. The function $\max\{x_1, x_2\}$ (plane OAC and OBC) on the unit square, which is a special case of the f^i defined in (4) for linear utility and $W = 1$. Its concave envelope (plane OAB and ABC) is the smallest concave function bounded by $\max\{x_1, x_2\}$ from below.

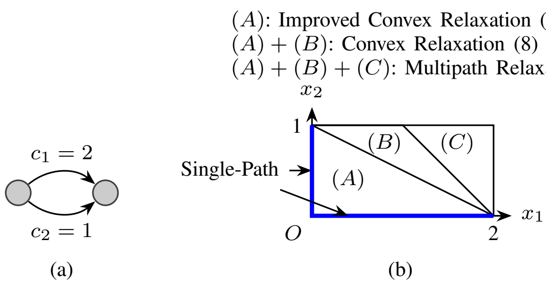

Fig. 3. (a) A simple network with one user and two links. The user can choose one link to send data from the node on the left to the node on the right. (b) Comparison among the feasible region of sparse (single-path) routing problem (1), convex relaxation (8), improved convex relaxation (9) and multipath relaxation (2) for the network in (a).

Lemma 3: For linear utility function $U^i(s) = s$, the concave envelope $\hat{f}^i$ for f^i defined in (4) is

$$\hat{f}^i(x^i) = \begin{cases} \|x^i\|_1 & \text{if } 0 \le x^i \le \|c\|_\infty, \|x^i\|_1 \le W\|c\|_\infty, \\ W\|c\|_\infty & \text{if } 0 \le x^i \le \|c\|_\infty, \|x^i\|_1 > W\|c\|_\infty, \\ -\infty & \text{otherwise.} \end{cases}$$

Proof: Due to the space constraint, we omit the proof but simply provide a graphical illustration (Fig. 2) here. ■

The convex relaxation (6) then becomes

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} \|x^i\|_1 \\ \text{s.t.} \quad & Rx \le c, \\ & x \ge 0, \\ & \|x^i\|_1 \le W\|c\|_\infty, \quad \forall i = 1, \dots, N. \end{aligned} \tag{8}$$

Let us illustrate how the convex relaxation (8) can be stronger than the multipath relaxation (2) and, in fact, can be further strengthened by considering a simple illustrative example. In Fig. 3(a), there are two links with capacity $c_1 = 2, c_2 = 1$ and a single user who can choose only one from them ($W = 1$). The feasible region of the original sparse (single-path) routing problem (1), convex relaxation (8) and multipath relaxation (2) are drawn in Fig. 3(b). In this example, the feasible region of (8) is not the smallest convex set containing the feasible region of the original problem (1), which demonstrates that (8) can be further improved by restricting the feasible region to the convex hull (A) in Fig. 3(b).

Define

$$\hat{c}_k^i = \min_{l=1,\dots,L} \{c_l | R_{lk}^i = 1\}$$

to be capacity of the bottleneck link on path k of user i. Tightening the last N constraints in (8), we obtain the following improved convex relaxation by using weighted l^1 norm instead of l^1 norm:

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} \|x^i\|_1 \\ \text{s.t.} \quad & Rx \le c, \\ & x \ge 0, \\ & \sum_{k=1}^{K^i} \frac{x_k^i}{\hat{c}_k^i} \le W, \quad \forall i = 1, \dots, N. \end{aligned} \tag{9}$$

Let opt_C denote the optimal value of (9). For any feasible solution x of sparse routing problem (1), $x_k^i \le \hat{c}_k^i$, so x is also feasible for (9), thus $opt_S \le opt_C$. On the other hand, for a feasible solution to (9), we can project it to a suboptimal solution to (1) by transmitting packets only on the paths that have W largest rates for each user.

Next, we will investigate the quality of the solution obtained by the above method. Among all the optimal solutions of (9), we shall focus on *vertex optimal solutions*, because the following result shows that the maximal violation of path cardinality constraints can be bounded for a vertex.

Lemma 4: Assume x is a vertex of the feasible region of (9). Let N' be the number of users with at least $W+1$ positive flows in rate allocation x, then

$$N' \le \frac{L}{W}.$$

At the same time, x has at most $N' + L$ positive flows.

Proof: Recall that a vertex must have $K = \sum_{i=1}^{N} K^i$ independent active constraints (i.e., constraints that hold with equality). If in vertex x, a user i has less than W positive flows, its corresponding constraint

$$\sum_{k=1}^{K^i} \frac{x_k^i}{\hat{c}_k^i} \le W \tag{10}$$

must be inactive.

Now assume user i has exact W positive flows and its corresponding constraint (10) holds with equality. Without loss of generality, we can assume that its first W paths are used. Then the only possible case is

$$\begin{aligned} x_k^i &= \hat{c}_k^i, & k &= 1, \dots, W, \\ x_k^i &= 0, & k &= W+1, \dots, K^i. \end{aligned}$$

For $k = 1, \dots, W$, if l is the bottleneck link for path k, then $\hat{c}_k^i = c_l$. The capacity constraint for link l must have the form $x_k^i = c_l$, i.e., the link l is fully occupied by user i. Therefore, the constraint (10) of user i is not an independent constraint because it can be written as a linear combination of the active constraints in $Rx \le c$ and $x \ge 0$.

Hence, there are at most N' independent active constraints from (10). Since at most L active constraints can be obtained from $Rx \le c$, there are at least $K - N' - L$ active constraints among $x \ge 0$, i.e., x must have at most $N' + L$ positive flows.

However, N' is the number of users who have at least $W+1$ positive flows. By the previous result, we have

$$(W+1)N' \le N' + L,$$

which implies $N' \le L/W$. ■

Based on the above lemma, we can bound the loss of the total rates after the projection of a vertex, which leads to the following main result.

Theorem 5: Assume x is a vertex of the feasible region of (9). Let y be the projection of x. Then

$$\sum_{i=1}^{N} \|x^i\|_1 - \sum_{i=1}^{N} \|y^i\|_1 \le \Psi(L, W)\|c\|_\infty,$$

where

$$\Psi(L, W) = \max_{n=1,\ldots,\lfloor L/W \rfloor} \left(n - \frac{Wn^2}{n+L}\right) W. \tag{11}$$

Proof: Let G^i be the number of positive flows of user i in rate allocation x. Let $\mathcal{S}$ denote the set of users with at least $W+1$ positive flows, i.e.,

$$\mathcal{S} = \{i | G^i \ge W+1, i = 1, \ldots, N\}.$$

Note that the set $\mathcal{S}$ contains the users who will be affected by the projection, and N' is just the number of users in $\mathcal{S}$.

If $i \in \mathcal{S}$, then the average of those positive rates of user i in x must be less than or equal to that in y, i.e.,

$$\frac{\|x^i\|_1}{G^i} \le \frac{\|y^i\|_1}{W},$$

because y^i only contains the flows of user i with W largest rates. Now

$$\begin{aligned}
\|x^i\|_1 - \|y^i\|_1 &\le \left(1 - \frac{W}{G^i}\right)\|x^i\|_1 \\
&\le \left(1 - \frac{W}{G^i}\right)\sum_{k=1}^{K^i} \frac{x_k^i}{\hat{c}_k^i}\|c\|_\infty \\
&\le \left(1 - \frac{W}{G^i}\right) W\|c\|_\infty, \qquad (12)
\end{aligned}$$

where (12) holds because x is feasible for (9). If $i \notin \mathcal{S}$, then

$$\|x^i\|_1 - \|y^i\|_1 = 0. \tag{13}$$

Adding up (12) and (13) for all users, we have

$$\begin{aligned}
\sum_{i=1}^{N}\|x^i\|_1 - \sum_{i=1}^{N}\|y^i\|_1 &\le \sum_{i\in\mathcal{S}}\left(1 - \frac{W}{G^i}\right)W\|c\|_\infty \\
&= \left(N' - W\sum_{i\in\mathcal{S}}\frac{1}{G^i}\right)W\|c\|_\infty \\
&\le \left(N' - W\frac{N'^2}{\sum_{i\in\mathcal{S}} G^i}\right)W\|c\|_\infty \\
&\le \left(N' - \frac{WN'^2}{N'+L}\right)W\|c\|_\infty,
\end{aligned}$$

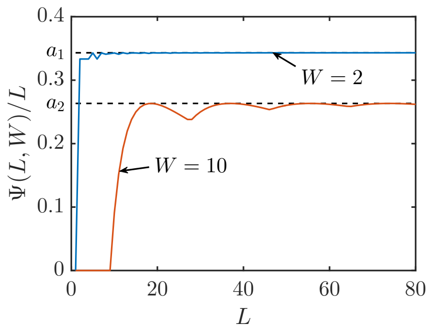

Fig. 4. The plot of $\Psi(L, W)/L$ as a function of L for a given W. When $W = 2$, as $L \to \infty$, the function tends to the upper bound, $a_1 = 2(\sqrt{2}-1)^2 \approx 0.343$, established in Proposition 6. Similarly, for $W = 10$, the function tends to the upper bound $a_2 = 10(\sqrt{10}-3)^2 \approx 0.263$.

in which the second last inequality holds because

$$\frac{N'}{\sum_{i\in\mathcal{S}} 1/G^i} \le \frac{1}{N'}\sum_{i\in\mathcal{S}} G^i,$$

and the last inequality is from Lemma 4.

Since N' is an integer between 0 and L/W, by enumerating all the possibilities for N', we establish the desired result. ■

Because the feasible region of (9) is bounded, the optimal solution to (9) can be attained at one vertex of the feasible region. Applying Theorem 5 to this optimal vertex solution, we can obtain a suboptimal solution to the original problem (1) whose value is at least

$$opt_C - \Psi(L, W)\|c\|_\infty,$$

and thus we have

$$opt_S \le opt_C \le opt_S + \Psi(L, W)\|c\|_\infty.$$

Note that this relaxation gap depends on neither the network topology nor the number of available paths for each user.

For single-path routing ($W = 1$),

$$\left(n - \frac{Wn^2}{n+L}\right)W = \frac{L}{1 + L/n}.$$

The maximizer inside (11) is always attained by $n = L$, so $\Psi(L, 1) = L/2$ and by solving the improved convex relaxation (9) we can find a suboptimal solution to the original problem (1) within the bound of $L\|c\|_\infty/2$. In contrast, Theorem 2 gives a suboptimal solution within the bound of $\sum_{i=1}^{\min\{N,L\}} \rho^i$, where

$$\rho^i = \left(1 - \frac{1}{K^i}\right)\|c\|_\infty$$

is the nonconcavity of the function (4) sorted in decreasing order. When $N \ge L$ and K^i is large, it is easy to see that our improved convex relaxation (9) for the linear utility case can provide a stronger guarantee as compared to the general case in Theorem 2.

The cases for $W > 1$ is complex due to the floor function in (11) (see Fig. 4), but we have the following result.

Proposition 6: If $W \geq 2$,

$$\Psi(L, W) \leq (\sqrt{W} - \sqrt{W-1})^2 WL.$$

Proof: For any integer $n = 1, \ldots, \lfloor L/W \rfloor$,

$$\begin{aligned}
&\left(n - \frac{Wn^2}{n+L}\right) W \\
&\quad= \left((2W-1)L - (n+L)(W-1) - \frac{WL^2}{n+L}\right) W \\
&\quad\leq \left((2W-1)L - 2\sqrt{(W-1)WL^2}\right) W \\
&\quad= (\sqrt{W} - \sqrt{W-1})^2 WL,
\end{aligned}$$

and the result follows. ∎

Remark 3: It is worth emphasizing that Theorem 5 only holds for vertices. To see how it may fail otherwise, consider a network with two nodes and three unit-capacity links from the left node to the right node. Assume three users want to simultaneously use the network to transmit, and each of them is only able to choose a single link. Then

$$x^1 = x^2 = x^3 = \begin{pmatrix} 1/3 & 1/3 & 1/3 \end{pmatrix}^T$$

is an optimal solution to convex relaxation (9). One possible projection of x to a sparse routing configuration is

$$y^1 = y^2 = y^3 = \begin{pmatrix} 1/3 & 0 & 0 \end{pmatrix}^T.$$

But

$$\sum_{i=1}^{3} \|x^i\|_1 - \sum_{i=1}^{3} \|y^i\|_1 = 2 > \Psi(3,1) = \frac{3}{2}.$$

V. Distributed Dual Algorithm

The convex relaxation (9) and Theorem 5 enable us to obtain an approximate solution to the sparse routing problem (1) by a distributed algorithm (Algorithm 1). The idea is first finding a vertex solution $\hat{x}$ to (9) using Lagrange dual decomposition, and then projecting $\hat{x}$ to a sparse routing configuration, i.e., selecting only W paths that have the largest rates in $\hat{x}$ for each user. A novel aspect of our distributed algorithm is its ability to yield a vertex solution to (9) using randomization and a feasible solution to (1) by projection *at the same timescale*.

Algorithm 1 (Distributed Algorithm with Randomization): In the following, $b > 0$ and $\alpha > 0$ are fixed constants. Let t be the algorithm iteration index.

- Source initialization for user i:
  1) Randomly choose nonnegative numbers τ_k^i for $k = 1, \ldots, K^i$ subject to i.i.d. uniform distribution over $[0, b]$.
  2) Set $\bar{x}^{i(0)}$ to be a $K^i \times 1$ zero vector.
- Link initialization for link l:
  1) Set initial link price $p_l^{(0)} \leftarrow 0$.
- Source update for user i in iteration t:
  1) Calculate the aggregate cost for each path:
  $$d_k^i \leftarrow \sum_{l=1}^{L} R_{lk}^i p_l^{(t-1)} + \tau_k^i, \quad k = 1, \ldots, K^i.$$
  2) Find the index k' maximizing $(1 - d_k^i) W \hat{c}_k^i$.
  3) $\check{x}_k^i \leftarrow \begin{cases} W\hat{c}_{k'}^i & \text{if } k = k' \text{ and } d_{k'}^i < 1, \\ 0 & \text{otherwise.} \end{cases}$
  4) Update the running average
  $$\bar{x}^{i(t)} \leftarrow \frac{1}{t}\left((t-1)\bar{x}^{i(t-1)} + \check{x}^i\right).$$
- Link update for link l in iteration t:
  1) Current step size $\alpha^{(t)} \leftarrow \alpha / t$.
  2) $p^{(t)} \leftarrow \left(p^{(t-1)} - \alpha^{(t)}(c - R\check{x})\right)_+$.[2]

The design of Algorithm 1 and its proof of convergence and optimality are summarized in the three main steps below:

1) We add a random perturbation to the objective function of the convex relaxation (9). The resulting problem (14) has a unique optimal solution that is also a vertex.
2) We show that the optimal solution to (14) is near-optimal for (9).
3) We apply dual decomposition to (14) to obtain a decentralized algorithm, establish its convergence and characterize its performance (Theorem 9).

Step 1: Note that Theorem 5 only holds for vertices, so the first step is to ensure that the obtained optimal solution to (9) is a vertex. If the optimal solution is unique, then it is always a vertex. However, we need a method to prevent the problematic case where the solution is not unique (this happens commonly when the network has some symmetry). Inspired by [14], we consider the following randomized version of (9):

$$\begin{aligned}
\max \quad & \sum_{i=1}^{N} \sum_{k=1}^{K^i} (1 - \tau_k^i) x_k^i \\
\text{s.t.} \quad & Rx \leq c, \\
& x \geq 0, \\
& \sum_{k=1}^{K^i} \frac{x_k^i}{\hat{c}_k^i} \leq W, \quad \forall i = 1, \ldots, N.
\end{aligned} \tag{14}$$

Here the vector $\tau = \{\tau_k^i | i = 1, \ldots, N, k = 1, \ldots, K^i\}$ can be randomly chosen according to any continuous distribution whose support is in the box $[0, b]^K$, where $b > 0$ is a fixed constant. In the source initialization part of Algorithm 1, it suffices to choose the uniform distribution. Based on the following well-known theorem, with probability 1, problem (14) has a unique optimal solution.

Theorem 7: Suppose S is a nonempty compact set in $\mathbb{R}^n$, then the set of $u \in \mathbb{R}^n$ for which the maximizer of $u^T x$ over $x \in S$ is not unique has Lebesgue measure 0.

The readers are referred to [11, Theorem 3] for a proof. For convenience, we omit the *with probability 1* condition in the rest of paper.

The unique optimal solution to (14) must be a vertex of its feasible region, and it is also a feasible vertex solution to the convex relaxation (9).

[2] Here $(u)_+ = \max\{u, 0\}$.

Step 2: In the following result, we bound how far the optimal solution to (14) can be from the optimality for (9).

Lemma 8: Let $\hat{z}$ be the unique optimal solution to the randomized problem (14), then

$$opt_C - \sum_{i=1}^{N} \|\hat{z}^i\|_1 \leq b\left(\frac{L}{W} + L\right)\|c\|_\infty.$$

Proof: Assume $\hat{x}$ is an optimal vertex solution to (9), then

$$\begin{aligned} opt_C - \sum_{i=1}^{N}\|\hat{z}^i\|_1 &= \sum_{i=1}^{N}\|\hat{x}^i\|_1 - \sum_{i=1}^{N}\|\hat{z}^i\|_1 \\ &\leq \sum_{i=1}^{N}\|\hat{x}^i\|_1 - \sum_{i=1}^{N}\sum_{k=1}^{K^i}(1-\tau_k^i)\hat{z}_k^i \\ &= \sum_{i=1}^{N}\sum_{k=1}^{K^i}(1-\tau_k^i)(\hat{x}_k^i - \hat{z}_k^i) + \sum_{i=1}^{N}\sum_{k=1}^{K^i}\tau_k^i\hat{x}_k^i \\ &\leq b\left(\frac{L}{W}+L\right)\|c\|_\infty, \end{aligned}$$

in which the inequalities hold because $0 \leq \tau_k^i \leq b$, $0 \leq \hat{x}_k^i \leq \|c\|_\infty$, $\hat{z}$ is an optimal solution to (14), and Lemma 4 implies that $\hat{x}$ has at most $L/W + L$ positive flows. ■

Step 3: We apply the Lagrange dual decomposition to the randomized problem (14). For each user i, let Q^i denote the polyhedron

$$\begin{cases} x^i \geq 0, \\ \sum_{k=1}^{K^i} \frac{x_k^i}{\hat{c}_k^i} \leq W. \end{cases}$$

Let $p \in \mathbb{R}^{L\times 1}$ be the Lagrange dual prices for the link capacity constraints. Consider the (partial) Lagrangian of (14):

$$\begin{aligned} \mathcal{L}(x,p) &= \sum_{i=1}^{N}\sum_{k=1}^{K^i}(1-\tau_k^i)x_k^i + p^T(c - Rx) \\ &= p^Tc + \sum_{i=1}^{N}\left(\sum_{k=1}^{K^i}(1-\tau_k^i)x_k^i - p^TR^ix^i\right) \\ &= p^Tc + \sum_{i=1}^{N}\sum_{k=1}^{K^i}(1-d_k^i)x_k^i, \end{aligned}$$

where

$$d_k^i = \sum_{l=1}^{L} R_{lk}^i p_l + \tau_k^i$$

is the aggregate cost for path k of user i, calculated by the first step of the source update part in Algorithm 1. The domain of $\mathcal{L}(x,p)$ is $p \geq 0$ and $x^i \in Q^i, \forall i = 1,\ldots,N$.

Remark 4: The aggregate cost d_k^i above has two components. The first component is the dynamic price depending on the link congestion. The second component τ_k^i is fixed during the entire iteration, so it can be understood as the static price for this path. Although, from a purely theoretical viewpoint, τ_k^i can be chosen according to any probability distribution satisfying the mentioned requirement, in the networking application we can choose τ_k^i in some meaningful way such as letting τ_k^i be proportional to the measured delay of the path.

Now, the Lagrangian $\mathcal{L}(x,p)$ is separable for different users. For a fixed link price p, let $\breve{x}^i$ be a maximizer for the separated Lagrangian, i.e.,

$$\breve{x}^i = \arg\max_{x^i \in Q^i} \sum_{k=1}^{K^i}(1-d_k^i)x_k^i.$$

$\breve{x}^i$ must be attained at some vertex of the polyhedron Q^i. If $d_k^i > 1$ for all available paths of user i, then $\breve{x}^i = 0$. Otherwise, assume path k' maximizes $(1-d_k^i)W\hat{c}_k^i$, then

$$\breve{x}^i = \begin{pmatrix} 0 & \cdots & W\hat{c}_{k'}^i & \cdots & 0 \end{pmatrix}^T,$$

in which the only nonzero component is the k'th. The third step of the source update part in Algorithm 1 computes $\breve{x}^i$.

The Lagrange dual problem for the randomized problem (14) is

$$\min_{p\geq 0}\max_{x^i\in Q^i} \mathcal{L}(x,p).$$

It is easy to see that strong duality holds for (14) by Slater's constraint qualification. The following iteration gives the subgradient method for solving the dual problem with step size $\alpha^{(t)}$:

$$p^{(t)} = \left(p^{(t-1)} - \alpha^{(t)}(c - R\breve{x}^{(t)})\right)_+,$$

where $\breve{x}^{(t)}$ maximizes the Lagrangian $\mathcal{L}(x,p)$ for fixed $p = p^{(t-1)}$.

In general, during the iterations of subgradient method, $\breve{x}^{(t)}$ may not be primal feasible, and this sequence may not be convergent. However, the running average

$$\bar{x}^{(t)} = \frac{1}{t}\sum_{j=1}^{t}\breve{x}^{(j)}$$

will converge to the optimal solution to (14) [15]. Now we can state and prove the main theorem about Algorithm 1.

Theorem 9: In Algorithm 1, $\bar{x}^{(t)}$ converges to a vertex $\hat{z}$ of convex relaxation (9). Let $\hat{y}$ be the projection of $\hat{z}$ to a sparse routing configuration, then

$$opt_S - \sum_{i=1}^{N}\|\hat{y}^i\|_1 \leq \Psi(L,W)\|c\|_\infty + b\left(\frac{L}{W}+L\right)\|c\|_\infty.$$

Proof: Let $\hat{z}$ be the unique optimal solution to (14), which is also a vertex of convex relaxation (9). Using the result in [15, Theorem 6], we know that every limit point of the sequence $\{\bar{x}^{(t)}\}$ must be $\hat{z}$. Since $\{\bar{x}^{(t)}\}$ is bounded and its limit point is unique, we have

$$\lim_{t\to\infty}\bar{x}^{(t)} = \hat{z}.$$

From Lemma 8, we have

$$opt_C - \sum_{i=1}^{N}\|\hat{z}^i\|_1 \leq b\left(\frac{L}{W}+L\right)\|c\|_\infty. \tag{15}$$

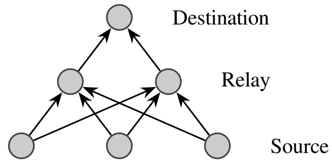

Fig. 5. A sample of relay network with three source nodes, two relay nodes and a single destination node. A user at each source node sends data to the destination node through some relay nodes.

Applying Theorem 5 to the projection $\hat{y}$ of $\hat{z}$, we have

$$\sum_{i=1}^{N} \|\hat{z}^i\|_1 - \sum_{i=1}^{N} \|\hat{y}^i\|_1 \leq \Psi(L, W) \|c\|_\infty. \quad (16)$$

The result follows by considering $opt_S \leq opt_C$ and adding up (15) and (16). ∎

Remark 5: After the projection, some users may lose some rates when restricted to using only W paths. Thus, $\hat{y}$ does not necessarily achieve the optimal utility for this fixed routing configuration. One can further improve the solution $\hat{y}$ by solving a network utility maximization over the sparse routing configuration chosen in the projection of $\hat{z}$.

VI. Numerical Examples

A. Relay Network Topology

In Section V, we have given a worst-case analysis for Algorithm 1. However, in many situations, the performance of Algorithm 1 is much better. In this part, we discuss how our algorithm can be applied to a relay network topology, in which Algorithm 1 is always able to converge to an optimal solution to the sparse routing problem (1).

We consider a network (Fig. 5) consisting of three types of nodes. There is a single node that is the destination of all users. There are R relay nodes with links to the destination node, and there are N users at different source nodes, which do not directly connect to the destination node but connect to all the relay nodes. Assume that each link in the network has unit capacity and $N \geq R$.

If the capacity constraints for the links between the destination and the relay nodes are satisfied, then all the other link capacity constraints are automatically satisfied. Only listing the constraints for up-level links, the sparse routing problem (1) for this example can be written as

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} \|x^i\|_1 \\ \text{s.t.} \quad & \sum_{i=1}^{N} x_r^i \leq 1, \quad \forall r = 1, \dots, R, \\ & x \geq 0, \\ & \|x^i\|_0 \leq W, \quad \forall i = 1, \dots, N, \end{aligned} \quad (17)$$

where x_r^i is the transmission rate of user i toward relay node r. The convex relaxation of (17) can be obtained by replacing the last N constraints with

$$\|x^i\|_1 \leq W, \quad \forall i = 1, \dots, N. \quad (18)$$

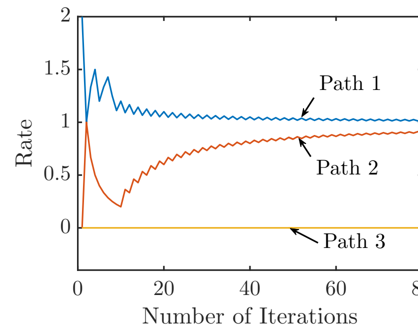

Fig. 6. The rates of one particular user under iterations of Algorithm 1 for the relay network topology example with $N = 4$, $R = 3$ and $W = 2$.

If the random perturbation in Algorithm 1 is sufficiently small, the algorithm will converge to an optimal vertex solution to the convex relaxation, and by the next proposition, which is also an optimal solution to the original sparse routing problem (17). The Algorithm 1 for the relay network topology is demonstrated in Fig. 6.

Proposition 10: For the relay network topology, each vertex of the convex relaxation is feasible for the original sparse routing problem (17).

Proof: The idea is to show that the coefficient matrix of the convex relaxation is totally unimodular and its every vertex is an integral solution. Thus, for such a vertex, the last N constraints (18) are equivalent to the last N path cardinality constraints in (17). ∎

B. Random Graph

The authors in [6] also proposed a heuristic algorithm based on Lagrange duality to solve the single-path routing problem, which is summarized in Algorithm 2. Similar to our Algorithm 1, it involves combining the dynamic price additively with a static component τ_k^i.

Algorithm 2 (Lagrange Dual Heuristic Algorithm in [6]): In iteration t:

1) Every user i calculates the aggregate cost for each path:
$$d_k^i \leftarrow \sum_{l=1}^{L} R_{lk}^i p_l^{(t-1)} + \tau_k^i, \quad k = 1, \dots, K^i.$$
2) For each user i, find the path $\hat{k}^i$ that has the least cost among all available paths of user i.
3) Solve the network utility maximization problem on the paths selected in step 2, i.e., solving problem (2) with the additional restrictions
$$x_k^i = 0, \quad \forall i = 1, \dots, N, \; \forall k \neq \hat{k}^i.$$
Let $(x^{(t)}, p^{(t)})$ be the optimal primal-dual pair.

If τ_k^i is omitted or if this term is small, Algorithm 2 does not converge and may oscillate among several routing configurations [6]. But, if τ_k^i is large, the resulting routing configuration weakly depends on the congestion level of links,

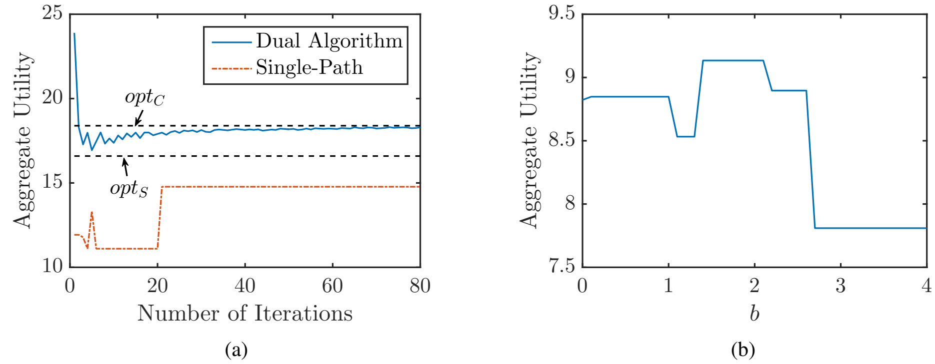

Fig. 7. (a) The aggregate utility of the solution to the relaxed problem (9) and the final single-path solution obtained after t iterations of Algorithm 1 with parameter $b = 0.2$ and $\alpha = 0.1$. (b) The time-average utility achieved by Algorithm 2 under different choices of b.

and it can be far from optimality. On the contrary, the reason for our Algorithm 1 to include the static price is to guarantee a vertex solution with the least violation of path cardinality constraints (a totally different algorithm design consideration as compared to the work in [6]). In Algorithm 1, τ_k^i can be relatively small, so it only slightly affects the performance of our algorithm.

In the following, we compare our Algorithm 1 with Algorithm 2 under a random network setting. The network has 10 links and 4 users. Each user has five available paths but it is only allowed to use one. The routing matrix R and link capacities c are randomly generated, and c is normalized such that $\|c\|_\infty = 10$. The static price τ_k^i is also randomly chosen according to i.i.d. uniform distribution over $[0, b]$.

Fig. 7(a) displays the aggregate utility of the approximate solution to the convex relaxation (9) generated by Algorithm 1 after t iterations. The final aggregate utility of the single-path routing solution, obtained by taking the algorithm output, projecting it to a single-path routing configuration and improving it by solving the network utility maximization problem over the selected paths (see Remark 5), is also shown in Fig. 7(a). We can see that a near-optimal single-path solution can be found after dozens of iterations. However, Algorithm 2 suffers large instability when $b \leq 2.6$. Even measuring the time-average utility (Fig. 7(b)), the performance is not nearly as good as Algorithm 1.

VII. Conclusion

We studied the network utility maximization over joint congestion control and routing with path cardinality constraints, which is in general a nonconvex and NP-hard problem. We proposed a novel convex relaxation that significantly outperformed the standard multipath relaxation, and enabled the design of distributed algorithms that can be interpreted as *sparse multipath* TCP/IP joint congestion control and routing algorithm. Even with path cardinality constraints, it can load balance a large portion of the traffic with a small number of paths. Numerical simulations show that the distributed dual algorithm performs well in the contrived setting as well as a random graph topology in comparison to a state-of-the-art TCP/IP algorithm, which demonstrates the value of our new convex relaxation and its decomposition.

References

[1] J. Qadir, A. Ali, K.-L. A. Yau, A. Sathiaseelan, and J. Crowcroft, "Exploiting the power of multiplicity: A holistic survey of network-layer multipath," *Communications Surveys & Tutorials, IEEE*, vol. 17, no. 4, pp. 2176–2213, Fourthquarter 2015.

[2] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Multi-path TCP: A joint congestion control and routing scheme to exploit path diversity in the Internet," *Networking, IEEE/ACM Transactions on*, vol. 14, no. 6, pp. 1260–1271, Dec. 2006.

[3] X. Lin and N. B. Shroff, "Utility maximization for communication networks with multipath routing," *Automatic Control, IEEE Transactions on*, vol. 51, no. 5, pp. 766–781, May 2006.

[4] C. Raiciu, S. Barre, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley, "Improving datacenter performance and robustness with multipath TCP," in *Proceedings of the ACM SIGCOMM 2011 Conference*, Aug. 2011, pp. 266–277.

[5] Q. Peng, A. Walid, J. Hwang, and S. H. Low, "Multipath TCP: Analysis, design, and implementation," *Networking, IEEE/ACM Transactions on*, 2015, forthcoming.

[6] J. Wang, L. Li, S. H. Low, and J. C. Doyle, "Cross-layer optimization in TCP/IP networks," *Networking, IEEE/ACM Transactions on*, vol. 13, no. 3, pp. 582–595, Jun. 2005.

[7] M. Wang, C. W. Tan, W. Xu, and A. Tang, "Cost of not splitting in routing: Characterization and estimation," *Networking, IEEE/ACM Transactions on*, vol. 19, no. 6, pp. 1849–1859, Dec. 2011.

[8] X. Liu, S. Mohanraj, M. Pioro, and D. Medhi, "Multipath routing from a traffic engineering perspective: How beneficial is it?" In *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*, Oct. 2014, pp. 143–154.

[9] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[10] S. H. Low, "A duality model of TCP and queue management algorithms," *Networking, IEEE/ACM Transactions on*, vol. 11, no. 4, pp. 525–536, Aug. 2003.

[11] M. Udell and S. Boyd. (Oct. 2014). Bounding duality gap for problems with separable objective. arXiv: 1410.4158.

[12] J. P. Aubin and I. Ekeland, "Estimates of the duality gap in nonconvex optimization," *Mathematics of Operations Research*, vol. 1, no. 3, pp. 225–245, Aug. 1976.

[13] A. Tang, J. Wang, and S. H. Low, "Counter-intuitive throughput behaviors in networks under end-to-end control," *Networking, IEEE/ACM Transactions on*, vol. 14, no. 2, pp. 355–368, Apr. 2006.

[14] S. Mehrotra, "On finding a vertex solution using interior point methods," *Linear Algebra and its Applications*, vol. 152, pp. 233–253, Jul. 1991.

[15] K. M. Anstreicher and L. A. Wolsey, "Two 'well-known' properties of subgradient optimization," *Mathematical Programming*, vol. 120, no. 1, pp. 213–220, Aug. 2009.