On Optimal Routing with Multiple Traffic Matrices

Chun Zhang[a], Yong Liu[a], Weibo Gong[b], Jim Kurose[a], Robert Moll[a], and Don Towsley[a]
[a]Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
{czhang, yongliu, kurose, moll, towsley}@cs.umass.edu
[b]Dept. of Electrical & Computer Engineering
University of Massachussetts
Amherst, MA 01003
gong@ecs.umass.edu

Abstract

Routing optimization is used to find a set of routes that minimizes cost (delay, utilization). Previous work has addressed this problem for the case of a known, static end-to-end traffic matrix. In the Internet, it is difficult to accurately estimate a traffic matrix, and the constantly changing nature of Internet traffic makes it costly to maintain optimal routing by responding to traffic changes. Thus, it is of interest to maintain a set of routes that are "good" for a number of different possible traffic scenarios. In this paper, we explore ways to find an optimal set of routes with multiple traffic matrices to minimize expected cost. We focus on two general approaches, source-destination routing and destination routing. In the case of source-destination routing, we extend existing methods with a single traffic matrix to solve the optimization problem with multiple traffic matrices: we extend the convex optimization solution methods for a single traffic matrix to the multiple traffic matrix case; we also extend the gradient-based solution methods for a single traffic matrix to the multiple traffic matrix case. However, the multiple traffic matrix case requires many more control variables. In the case of destination routing, we encounter many more differences from the single traffic matrix case. The loop-free property, which is valid for the single traffic matrix case, is no longer valid for the multiple traffic matrix case, and it is difficult to extend existing methods for a single traffic matrix to solve the optimization problem with multiple traffic matrices. We show that it is NP-complete even to determine the feasibility of multiple traffic matrices. We thus propose and evaluate a heuristic algorithm for this case.

I. Introduction

Routing optimization is used to find a set of routes, i.e., the set of paths along which packets are forwarded in order to optimize a well-defined objective function (such as delay or utilization). Routing approaches are generally divided into source-destination routing (henceforth referred to as flow routing) and destination routing. As a packet travels through a network, a flow routing approach such as MPLS [1] forwards it based on its source and destination addresses. A destination routing approach such as OSPF [2] forwards it only on the basis of its destination address. Destination routing is unable to provide as fine control on routing as flow routing because it uses less information.

A traffic matrix (TM) specifies the data rate between every pair of ingress and egress points. A number of works [3] [4] [5] have focused on calculating an optimal set of routes for a single TM. For a given TM, those works consider minimizing the sum of link costs, each of which is an increasing convex function of link data rate. The problem is then formalized and solved as an optimization problem. With a single TM, methods to solve the problem for flow routing and destination routing are similar, and the optimal costs are identical. In [3], Cantor *et al.* proposed a centralized algorithm. In [4], Gallager proposed a distributed algorithm. To solve the problem more efficiently, the link costs can be approximated as piece-wise linear functions [5], and the problem then formalized and solved by linear programming (LP).

For a large-scale Internet with changing traffic, optimization with multiple TMs is an important problem for several reasons. First, accurate TM estimation is hard to achieve due to scale, as well as due to the inherent challenges in estimating a TM [6] [7]. Without an accurate TM, optimization over multiple TM candidates calculates a set of routes that is more robust to estimation errors. Second, even if the current TM is known, the changing nature of Internet traffic makes it costly to continually maintain optimal routing by responding to traffic changes (Routing convergence normally take seconds, during which packets may be lost, or arrive out of order. Frequent routing updates can make the situation even worse). As routing updates are performed at a slower rate than the change in traffic, it is preferable to implement a set of routes that can perform well for all TMs between routing updates.

In this paper, we explore ways to obtain an optimal set of routes with multiple TMs so as to minimize expected cost. We focus on both flow routing and destination routing. In the case of flow routing, we extend existing solution methods for a single TM to solve the optimization problem with multiple TMs: we show that the optimization problem can still be formalized and solved as a convex optimization problem (as in the single TM case). We also extend Gallager's work with a single TM to solve the problem with multiple TMs using gradient-based methods. However, the multiple TM case requires many more control variables. In the case of destination routing, we encounter many more differences from the single TM case. We find that the loop-free property, which is valid for the single TM case, is no longer valid for the multiple TM case. It is difficult to extend the solution methods used with a single TM to solve the problem with multiple TMs. We show that it is NP-complete even to determine the feasibility of multiple TMs. Thus, we propose and evaluate a heuristic algorithm for this case.

The remainder of this paper is organized as follows. In Section 2, we review related work. In Section 3, we formulate the multiple TM routing optimization problem. In Section 4, for flow routing, we first compute an optimal set of flow routes using convex optimization techniques, and then extend Gallager's work to solve the problem using gradient-based methods. In Section 5, for destination routing, we demonstrate the inherent difficulty of solving the optimization problem, and then propose and evaluate a heuristic algorithm. Section 6 concludes the paper.

II. Context and Related Work

Internet routing protocols are generally classified into two categories: flow routing and destination routing. MPLS, a flexible routing protocol, is normally considered a flow routing protocol [5] [6][8]; OSPF, a commonly used intra-domain Internet routing protocol, falls into the category of destination routing. Specifically, a relaxed version of OSPF, which allows arbitrary routing fractions on the shortest paths to the destination, is a loop-free destination routing protocol [9].

Routing fractions are useful for describing a set of routes along which packets are forwarded. In flow routing, for each source and destination pair, a router maintains a routing fraction for each of its out-going links. Specifically, $\phi_{kl}(i,j)$ denotes the fraction of traffic originating from router i destined to router j at router k forwarded over link (k,l). In Figure 1, router 3 forwards 100% of the traffic originating from router 1 destined to router 6 over outgoing link $(3,4)$ and 100% of the traffic originating from router 2 destined to router 6 over link $(3,5)$. In contrast, destination routing only maintains a routing fraction for each destination. Specifically, $\phi_{kl}(j)$ denotes the fraction of traffic destined to router j at router k forwarded over outgoing link (k,l). In Figure 2, router 3 forwards traffic destined to router 6 evenly over two out-going links: 50% over link $(3,4)$, and 50% over link $(3,5)$. Destination routing can be viewed as a special case of flow routing where the routing fractions to a common destination are identical for all sources. Given a TM, routing fractions determine packet forwarding, the link data rates, and thus the cost. In our optimization problem, we refer to routing fractions as *routing variables*.

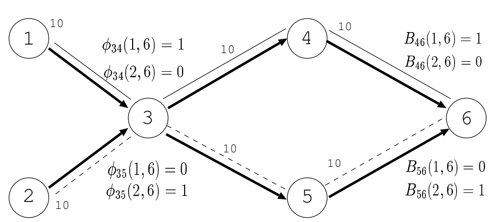

Fig. 1. Flow Routing : traffic originating from different source addresses is forwarded by *different sets of routes*

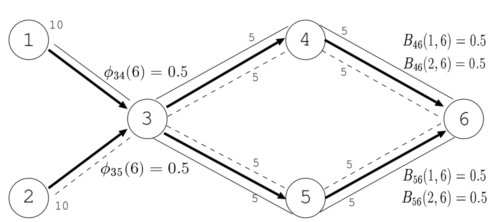

Fig. 2. Destination Routing : traffic originating from different source addresses is forwarded by *single set of routes*

An alternative way to describe a set of routes is through so-called traffic ratios. For each source and destination pair, $B_{kl}(i,j)$ denotes the ratio of the traffic originating from router i destined to router j over link (k,l) to the overall traffic originating from router i destined to router j. In Figure 1, link $(4, 6)$ carries 100% of the traffic originating from router 1 destined to router 6; in Figure 2, link $(4, 6)$ carries 50% of the traffic originating from router 1 destined to router 6. Given a TM, traffic ratios determine packet forwarding, the link data rates, and thus the cost. In our optimization problem, we refer to such traffic ratios as *ratio variables*.

A number of efforts [3][4][5] have investigated the routing optimization problem in the case of a single TM. The methods used for flow routing and destination routing are similar. In general, the problem is formalized and solved as an optimization problem. Using ratio variables as control variables, Cantor *et al.* [3] solved the problem using convex optimization techniques. To increase efficiency, [5] approximates the link costs as piece-wise linear functions, and solves the problem using LP. In [4], Gallager proposed a distributed gradient-based algorithm to solve the problem using routing variables as control variables.

The route optimization problem is relatively new in the case of multiple TMs. While researchers have recently identified the importance of the route optimization problem in the presence of multiple TMs [7][8], they have yet to investigate techniques for solving the problem.

An optimal set of routes is necessarily feasible. With multiple TMs, the set of feasible route-sets fundamentally differs from that with a single TM. A set of TMs is feasible if there exists a set of routes so that the resulting link data rates are always less than or equal to link capacity *for each TM*. The set of routes is then called a feasible set of routes for the set of TMs. With a single TM, [3][4] find an optimal set of routes out of the set of feasible route-sets. The cost of an optimal set of flow routes, an optimal set of destination routes, and an optimal set of loop-free destination routes are identical. With multiple TMs, the set of feasible route-sets is the intersection of the sets of feasible route-sets for each individual TM. As a result, a set of TMs may be infeasible even though each TM in the set is individually feasible. Moreover, the cost of an optimal set of flow routes may be lower than that of an optimal set of destination routes; also, the cost of an optimal set of destination routes with loops may be lower than that of an optimal set of loop-free destination routes.

We will see that, with multiple TMs, the hardness of the optimization problem is closely related to the routing approach. In the case of flow routing, we can extend the solution methods of a single TM to the case of multiple TMs. Using ratio variables as control variables, we extend [3] to solve the route optimization problem using convex optimization techniques, and thus solve the problem using LP when link costs are approximated as piece-wise linear functions. Using routing variables as control variables, we extend [4] to solve the flow routing problem with multiple TMs using gradient-based methods. However, the multiple TM case requires many more control variables compared to the single TM case. In the case of destination routing, we demonstrate the inherent difficulties to solve the problem with multiple TMs. The set of feasible route-sets is not convex when we use ratio variables as control variables. As a result, we cannot solve the problem with multiple TMs as a convex optimization problem. Furthermore, when using routing variables as control variables, we find local minima — making it difficult to solve the problem using gradient-based methods. Finally, we show that it is NP-complete even to determine the feasibility of a set of multiple TMs.

III. Problem Formulation

In this section, we formulate the optimal routing problem with multiple TMs. We first introduce the necessary notation, and then formalize the problem. Finally, we describe the difference between route optimization with a single TM and with multiple TMs.

A. Notation

We first introduce the notation for flow routing, and then the notation needed for destination routing.

Notation for flow routing:

Network topology: $G = (V, E)$ is a strongly connected graph.[1] The network G is composed of a set of nodes V and a set of directed links E. The nodes in V are represented by the integers $1, 2, \ldots, |V|$. The directed links in E are represented by $(k, l) \in V^2$.

Link capacity : $C = \{c_{kl}\}$, where $c_{kl} > 0$ denotes the capacity of link $(k, l) \in E$.

[1] In some cases, we relax this assumption for ease of exposition, and note this relaxation when used.

Traffic Matrices : $R = \{R_1, R_2, \ldots, R_n\}$ is a set of n traffic matrices with associated positive weights $w = \{w_1, w_2, \ldots, w_n\}$, $\sum_y w_y = 1$. In TM $R_y = [R_y(i,j)]$, $i,j \in V$, $y \in \{1,\ldots,n\}$, $R_y(i,j)$ denotes the rate of exogenous traffic, in bits/s, originating from node i destined to node j; w_y is the weight of TM R_y.

Routing variables : $\Phi = \{\phi_{kl}(i,j)\}$, $i,j \in V$, $(k,l) \in E$, where $\phi_{kl}(i,j)$ denotes the fraction of traffic rate from node i to node j at node k forwarded over link (k,l). When Φ are used as control variables in optimization problem formulation, the constraints are,

1) $\phi_{kl}(i,j) \geq 0, i,j \in V, (k,l) \in E$,
2) $\phi_{kl}(i,j) = 0$ if $k = j$,
3) $\sum_l \phi_{kl}(i,j) = 1$ if $k \neq j$,
4) $\forall i,j,k (k \neq j) \in V$, for traffic from i destined to j at node k, there exists at least one path between k and j: there is a sequence of nodes, $k,l,p,\ldots,q,j$ such that $\phi_{kl}(i,j) > 0$, $\phi_{lp}(i,j) > 0$, …, $\phi_{qj}(i,j) > 0$.

Ratio variables : $B = \{B_{kl}(i,j)\}$, $i,j \in V$, $(k,l) \in E$, where $B_{kl}(i,j)$ denotes the ratio of the traffic rate originating from i destined to j that is forwarded over link (k,l) to the overall traffic rate originating from i destined to j. When B are used as control variables in optimization problem formulation, the constraints are,

1) $B_{kl}(i,j) \geq 0, i,j \in V, (k,l) \in E$,
2) $B_{kl}(i,j) = 0$ if $k = j$,
3)
$$\sum_m B_{mk}(i,j) - \sum_l B_{kl}(i,j) = \begin{cases} 1 & k = j \\ -1 & k = i \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$
4) if $B_{kl}(i,j) > 0$, then for traffic from i destined to j at node k, there exists at least one path between k and j: there is a sequence of nodes, $k,l,p,\ldots,q,j$ such that $B_{kl}(i,j) > 0$, $B_{lp}(i,j) > 0$, …, $B_{qj}(i,j) > 0$.

The equivalence of routing variables (Φ) and ratio variables (B) is indicated by [4]. For completeness, we explicitly express it as Theorem III.1 (see below). A set of routes is said to be loop-free if the corresponding set of ratio variables B is loop-free. i.e., there is no sequence of nodes, $i,j,k,l,p,\ldots,q \in V$ such that $B_{kl}(i,j) > 0$, $B_{lp}(i,j) > 0$, …, $B_{qk}(i,j) > 0$.

Theorem III.1: In a strongly connected graph $G = (V,E)$, a set of flow routing variables Φ determines a set of flow ratio variables B; a set of flow ratio variables B can be implemented by a set of flow routing variables Φ.

Proof: Given a set of routing variables Φ, we can compute a set of ratio variables B as follows. Let $b_k(i,j)$ denote the ratio of the traffic rate originating from node i destined to node j at node k to the overall traffic rate originating from node i destined to node j, we have

$$b_k(i,j) = \mathbf{1}(k = i) + \sum_m b_m(i,j)\phi_{mk}(i,j) \tag{2}$$

Here, $\mathbf{1}(P)$ is 1 if the predicate P is true and 0 otherwise. [4] shows that equations (2) must have a unique solution of b. After solving b, we compute B from b,

$$B_{kl}(i,j) = b_k(i,j)\phi_{kl}(i,j) \tag{3}$$

Given a set of ratio variables B, we can construct a set of routing variables Φ to implement B as follows. For each node $j \in V$, we construct a shortest path tree to j. For $i,j \in V$, $(k,l) \in E$, if $\sum_m B_{km}(i,j) > 0$, we set

$$\phi_{kl}(i,j) = \frac{B_{kl}(i,j)}{\sum_m B_{km}(i,j)} \tag{4}$$

If $\sum_m B_{km}(i,j) = 0$, we set $\phi_{kl}(i,j) = 1$ if link (k,l) is on the shortest path tree to node j, and $\phi_{kl}(i,j) = 0$ otherwise. ■

Node data rates : $T_y = \{t_{y,k}(i,j)\}$, $k,i,j \in V$, $y \in \{1,\ldots,n\}$, where $t_{y,k}(i,j)$ denotes the data rate at node k from node i destined to node j under TM R_y. We have,

$$t_{y,k}(i,j) = \mathbf{1}(k = i)R_y(i,j) + \sum_m t_{y,m}(i,j)\phi_{mk}(i,j) \tag{5}$$

$$t_{y,k}(i,j) = R_y(i,j)\Big(\sum_m B_{mk}(i,j) + 1(k = i)\Big) \tag{6}$$

Link data rates : $F_y = \{f_{y,kl}\}$, $(k,l) \in E$, $y \in \{1, \dots, n\}$, where $f_{y,kl}$ denotes the link data rate over link (k,l) under TM R_y. We have,

$$f_{y,kl} = \sum_{i,j} t_{y,k}(i,j)\phi_{kl}(i,j) \tag{7}$$

$$f_{y,kl} = \sum_{i,j} R_y(i,j)B_{kl}(i,j) \tag{8}$$

Feasibility : Given a set of n TMs, if there exists a set of routing variables Φ (or ratio variables B) such that the resulting link data rates are always less than or equal to the link capacity for *each* of the n TMs, then the set of n TMs is feasible, and Φ (or B) is feasible for the set of n TMs. Specifically, given n TMs, we use Ψ (or $\mathcal{B}$) to denote the set of feasible route-sets described by Φ (or B).

Link cost function: $D = \{D_{kl}\}$, $(k,l) \in E$, where D_{kl} denotes the cost function of link (k,l). We assume that the link cost is a convex, increasing function of link data rate. While our analysis can be applied to any function with such properties, we will use,

$$D_{kl}(x) = \frac{x}{c_{kl} - x} \tag{9}$$

This $M/M/1$-like link cost can be approximated by piece-wise linear functions. Specifically, Let (k_i, b_i), $i \in \{1, \dots, 6\}$ be (2^{4i-2}, $-2^{4i-2} + 5 * 2^{2i-2} - 1$). We have,

$$D_{kl}(x) = \max_{1 \le i \le 6} (k_i \frac{x}{c_{kl}} + b_i) \tag{10}$$

Network Cost : Let A_y denote the cost of TM R_y, $y \in \{1, \dots, n\}$ and A the expected cost; we have,

$$A_y = \sum_{(k,l) \in E} D_{kl}(f_{y,kl}) \tag{11}$$

$$A = \sum_{y=1}^{n} w_y A_y \tag{12}$$

The following notation differs in the case of destination routing:

Routing variables : $\Phi = \{\phi_{kl}(j)\}$, $j \in V$, $(k,l) \in E$, where $\phi_{kl}(j)$ denotes the fraction of traffic rate to node j at node k forwarded over link (k,l). Destination routing variables can be viewed as a special case of flow routing variables with the additional constraints,

$$\phi_{kl}(i_1, j) = \phi_{kl}(i_2, j), i_1, i_2, j \in V, (k,l) \in E \tag{13}$$

Ratio variables : Similar to routing variables, combining equations (4) and (13), the destination ratio variables B must satisfy the additional constraints,

$$\frac{B_{kl}(i_1, j)}{\sum_m B_{km}(i_1, j)} = \frac{B_{kl}(i_2, j)}{\sum_m B_{km}(i_2, j)} \quad i_1, i_2, j \in V, (k,l) \in E \tag{14}$$

where $\sum_m B_{km}(i_1, j) > 0$ and $\sum_m B_{km}(i_2, j) > 0$.

All other definitions, theorems for flow routing are the same in the case of destination routing. Theorem III.2 shows that destination routing variables (Φ) and destination ratio variables (B) are equivalent.

Theorem III.2: In a strongly connected graph $G = (V, E)$, a set of destination routing variables Φ determines a set of destination ratio variables B; a set of destination ratio variables B can be implemented by a set of destination routing variables Φ. (Following the same proof of Theorem III.1)

B. The Single TM Problem

With a single TM, the routing optimization problem has several important properties. It was known that, for a single TM, the link data rates implemented by a set of flow routes can also be implemented by a set of destination routes

[10]. This, plus the loop-free property (see Theorem III.3), state that with a single TM, the optimal set of flow routes, destination routes and loop-free destination routes yield the same cost.

Theorem III.3: Loop-free property: in a strongly connected graph $G = (V, E)$, given a feasible TM R_1, the flow route optimization problem always has an optimal solution as a set of loop-free flow routes, and the destination route optimization problem always has an optimal solution as a set of loop-free destination routes. [3][4]

With a single TM, in order to solve the optimal routing problem in a distributed or centralized manner, the problem has been formulated using either routing variables Φ or ratio variables B as control variables.

[4] formulated the problem using routing variables Φ as control variables in the case of destination routing.

Problem Formulation over Φ:

Given: network $G = (V, E)$, link capacity C, a single TM R_1.
Minimize: cost A.
Constraints:
1) Route constraints. F_1 is implemented by a set of destination routes Φ.
2) Feasibility constraints. $F_1 \leq C$. i.e., $\forall (k,l) \in E$, $f_{1,kl} \leq c_{kl}$.

[3] formulated the problem as a convex optimization problem. The problem was formalized using ratio variables B as control variables in the case of flow routing.

Problem Formulation over B:

Given: network $G = (V, E)$, link capacity C, a single TM R_1.
Minimize: cost A.
Constraints:
1) Route constraints. F_1 is implemented by a set of flow routes B.
2) Feasibility constraints. $F_1 \leq C$.

With a single TM, the route optimization problem can also be formulated using a smaller number of control variables when destination-based link data rates F^D (introduced next) are used as control variables [10].

Destination-based link data rates: $F_y^D = \{f_{y,kl}^D(j)\}$, $y \in \{1, \ldots, n\}$, $(k,l) \in E$, $j \in V$, where $f_{y,kl}^D(j)$ denotes the data rate of the traffic destined to j over link (k,l) under TM R_y. When F_y^D are used as control variables, the constraints are,

1) $f_{y,kl}^D(j) \geq 0$, $y \in \{1, \ldots, n\}$, $j \in V$, $(k,l) \in E$,
2) $f_{y,kl}^D(j) = 0$ if $k = j$,
3)
$$\sum_m f_{y,mk}^D(j) - \sum_l f_{y,kl}^D(j) = \begin{cases} \sum_i R_y(i,j) & k = j \\ -R_y(i,j) & k = i \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

The link data rates F_y are expressed by destination-based link data rates F_y^D as follows.

$$f_{y,kl} = \sum_j f_{y,kl}^D(j), \quad y \in \{1, \ldots, n\}, (k,l) \in E, j \in V \quad (16)$$

Problem Formulation over F_1^D:

Given: network $G = (V, E)$, link capacity C, a single TM R_1.
Minimize: cost A.
Constraints:
1) Flow conservation constraints. F_1 is expressed by F_1^D.
2) Feasibility constraints. $F_1 \leq C$.

With a single TM, the set of destination routing variables Φ can be expressed in terms of F_1^D,

$$\phi_{kl}(j) = \frac{f_{1,kl}^D(j)}{\sum_m f_{1,km}^D(j)} \quad (17)$$

where $\sum_m f^D_{1,km}(j) > 0$.

C. The Multiple TM Problem Formulation

We now generalize the problem statement for a single TM to the case of multiple TMs. We use either ratio variables B or routing variables Φ as control variables.

Given: network $G = (V, E)$, link capacity C, n TMs.
Minimize: cost A.
Constraints:
For each TM R_y, $y \in \{1, \ldots, n\}$,
1) Route constraints. F_y is implemented by a set of routes Φ or B.
2) Feasibility constraints. $F_y \leq C$.

When link costs are approximated by piece-wise linear functions, they can be expressed as additional constraints.

3) Piece-wise constraints. For $y \in \{1, \ldots, n\}$,

$$D_{kl}(f_{y,kl}) \geq k_i \frac{f_{y,kl}}{c_{kl}} + b_i, (k,l) \in E, i \in \{1, \ldots, 6\}$$

However, the formulation with destination-based link data rates F^D cannot be easily extended to the case of multiple TMs. Flow conservation (15) only guarantees that for each individual TM in isolation, the demand can be satisfied by some set of destination routes (17). It does not guarantee that *a single* set of destination routes be used to forward packets for *all* TMs.

D. Route Optimization with Multiple TMs: differences from the Single TM case

Properties that hold for a single TM do not necessarily hold for multiple TMs. In particular, with multiple TMs, the cost of an optimal set of flow routes may be lower than that of destination routes, and the cost of an optimal set of destination routes with loops may be lower than that of destination loop-free routes. We demonstrate this through three counter-examples. We show that a set of TMs is not feasible even though each TM in the set is individually feasible; we also show that a set of TMs that is feasible with respect to flow routing may not be feasible with respect to destination routing. Finally, we also show that a set of TMs that is feasible with respect to destination routing may not be feasible with respect to loop-free destination routing. All examples are based on a network[2] G shown in Figure 3. In all cases, traffic is only destined to node 3.

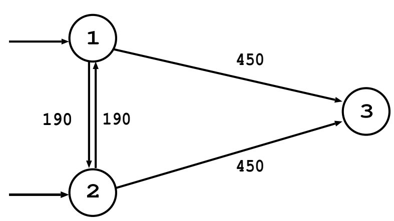

Fig. 3. A topology to illustrate the difference of route optimization between single TM and multiple TMs

First, we present a set of two TMs where each TM is individually feasible but not feasible under flow routing when considered as a set,

$$R_1 = \begin{bmatrix} 0 & 0 & 600 \\ 0 & 0 & 200 \\ 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 200 \\ 0 & 0 & 600 \\ 0 & 0 & 0 \end{bmatrix} \tag{18}$$

In TM R_1, as $R_1(1,3) = 600$ and $c_{13} = 450$, a feasible set of routes of R_1 ($\phi_{13}(3) = 0.7$, $\phi_{23}(3) = 1$, $\phi_{12}(3) = 0.3$, $\phi_{21}(3) = 0$) must forward at least $1/4$ of the traffic originating from 1 over link $(2,3)$. i.e., $B_{23}(1,3) \geq 0.25$. This

[2]We use directed graph for ease of exposition.

results in that, in TM R_2, the rate of the traffic originating from 1 that is forwarded over link $(2,3)$ must be at least 50 bits/s, and the remaining capacity of link $(2,3)$ for the traffic originating from 2 is at most 400. As $c_{21}+400 \le R_2(2,3)$, the set of the two TMs is not feasible.

Second, we present a set of two TMs that is feasible under flow routing but not under destination routing,

$$R_1 = \begin{bmatrix} 0 & 0 & 600 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 600 \\ 0 & 0 & 0 \end{bmatrix} \tag{19}$$

In TM R_1, traffic only originates from node 1, and in TM R_2, traffic only originates from node 2. Since TM R_1 and TM R_2 are individually feasible (with the same feasible set of routes as in the previous example), and with flow routing, the traffic of TM R_1 and TM R_2 are forwarded using routes based on different source and destination pairs, the set of the two TMs is still feasible in the case of flow routing. However, in the case of destination routing, traffic is forwarded without differentiating the source address of the packets. Using similar arguments as in the previous example, we know that when TM R_1 is feasible, the rate of the traffic originating from 1 forwarded over link $(2,3)$ must be at least 150 bits/s under TM R_1. Thus, $\phi_{23}(3) \ge \frac{150}{190}$ where 190 is the capacity of link $(1,2)$. This results in, that in TM R_2, the rate of the traffic originating from 2 that is forwarded over link $(2,3)$, $R_2(2,3)\phi_{23}(3)$, is at least $\frac{9000}{19}$, and thus exceeds the capacity of link $(2,3)$. Consequently, the set of the two TMs is not feasible in the case of destination routing.

Third, we present a set of two TMs that is feasible under destination routing but not under loop-free destination routing,

$$R_1 = \begin{bmatrix} 0 & 0 & 500 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 500 \\ 0 & 0 & 0 \end{bmatrix} \tag{20}$$

In TM R_1, as the traffic rate originating from node 1 exceeds the capacity of link $(1,3)$, a feasible set of routes for TM R_1 must forward part of that traffic through node 2. Similarly, a feasible set of routes for TM R_2 must forward part of the traffic originating from node 2 through node 1. Thus, a feasible set of destination routes for the set of the two TMs must include a loop between node 1 and 2. In fact, we can see that the set of destination routes with loops ($\phi_{13}(3) = \phi_{23}(3) = 0.7$, $\phi_{12}(3) = \phi_{21}(3) = 0.3$) is feasible for the set of the two TMs. The resulting link data rates for TM R_1 are ($f_{1,12} \simeq 165$, $f_{1,21} \simeq 50$, $f_{1,13} \simeq 385$, $f_{1,23} \simeq 115$).

With multiple TMs, the costs of the optimal set of flow routes and destination routes may differ. Therefore, we consider the flow routing and destination routing problems separately in the following two sections.

IV. Optimal Flow Routing with Multiple TMs

In the previous section, we formulated the routing optimization problem, and discussed the differences in optimizing routes with a single TM and with multiple TMs. In this section, we explore ways of computing an optimal set of *flow routes* for multiple TMs. We first solve the problem with routing variables as control variables. Then we solve the problem with ratio variables as control variables.

Using routing variables as control variables, we now extend [4] to the case of flow routing with multiple TMs, and solve the problem using a gradient-based algorithm. Assume that Φ is the set of routing variables used by a set of n TMs. In order to obtain derivative information $\partial A/\partial\phi_{kl}(i,j)$, $(k,l) \in E$, $i,j \in V$, we introduce a set of dummy variables $r_y = \{r_{y,k}(i,j)\}$, $y \in \{1,\ldots,n\}$, $k,i,j \in V$, where $r_{y,k}(i,j)$ can be understood as the rate of the dummy traffic injected at node k destined to node j under TM R_y using the same routing fractions Φ as the traffic originating from i destined to j.

For TM R_y, $y \in \{1,\ldots,n\}$, similar to [4], we have,

$$\frac{\partial A_y}{\partial r_{y,k}(i,j)} = \sum_l \phi_{kl}(i,j)\left[D'_{kl}(f_{y,kl}) + \frac{\partial A_y}{\partial r_{y,l}(i,j)}\right] \tag{21}$$

$$\frac{\partial A_y}{\partial \phi_{kl}(i,j)} = t_{y,k}(i,j)\left[D'_{kl}(f_{y,kl}) + \frac{\partial A_y}{\partial r_{y,l}(i,j)}\right] \tag{22}$$

where $D'_{kl}(f_{y,kl}) = \frac{dD_{kl}(f_{y,kl})}{df_{y,kl}}$.

Combined with equations (6) and (12), we have,

$$\frac{\partial A}{\partial r_{y,k}(i,j)} = \sum_l \phi_{kl}(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] \tag{23}$$

$$\frac{\partial A}{\partial \phi_{kl}(i,j)} = \sum_y t_{y,k}(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] \tag{24}$$

$$= \left(\sum_m B_{mk}(i,j) + 1(k=i)\right)\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] \tag{25}$$

The existence and uniqueness of $\partial A/\partial r_{y,k}(i,j)$ and $\partial A/\partial \phi_{kl}(i,j)$ is given by the following theorem.

Theorem IV.1: Let a network G have n TMs and routing variables Φ, and let each marginal link cost $D'_{kl}(f_{y,kl})$ be continuous in $f_{y,kl}$, $(k,l) \in E$. Then the set of equations (23), $k \neq j$, has a unique (and correct) set of solutions for $\partial A/\partial r_{y,k}(i,j)$. Furthermore, (24) is valid and both $\partial A/\partial r_{y,k}(i,j)$ and $\partial A/\partial \phi_{kl}(i,j)$ for $k \neq j$, $(k,l) \in E$ are continuous in r and Φ.

Proof: See Appendix A. ■

Using Lagrange multipliers for the constraint $\sum_l \phi_{kl}(i,j) = 1$, and taking into account the constraint $\phi_{kl}(i,j) \geq 0$, the necessary conditions for a minimum of A with respect to Φ are, for all $k \neq j$, $(k,l) \in E$,

$$\frac{\partial A}{\partial \phi_{kl}(i,j)} \begin{cases} = \lambda_{kij} & \phi_{kl}(i,j) > 0 \\ > \lambda_{kij} & \phi_{kl}(i,j) = 0. \end{cases} \tag{26}$$

This states that for given i,j,k, all links (k,l) for which $\phi_{kl}(i,j) > 0$ must have the same marginal cost $\partial A/\partial \phi_{kl}(i,j)$, and that this marginal cost must be less than or equal to $\partial A/\partial \phi_{kl}(i,j)$ for the links on which $\phi_{kl}(i,j) = 0$. However, as shown by [4], even for a single TM, (26) is not a sufficient condition to minimize A.

Given i,j,k in (24), if $\sum_m B_{mk}(i,j) + 1(k=i) = 0$, then $\forall l$, we have $\partial A/\partial \phi_{kl}(i,j) = 0$. This means that, if node k is not on any route carrying the traffic from i destined to j, the above conditions would be automatically satisfied. Thus, we hypothesize that (26) would be sufficient to minimize A if the factor $\sum_m B_{mk}(i,j) + 1(k=i)$ were removed from the condition.

Theorem IV.2: For each $(k,l) \in E$, assume that $D_{kl}(f_{y,kl})$ is convex and continuously differentiable for $0 \leq f_{y,kl} < c_{kl}$. Let Υ be the set of Φ for which the link data rates satisfy $f_{y,kl} < c_{kl}$, $y \in \{1,\ldots,n\}$, $(k,l) \in E$. Then (26) is necessary for Φ to minimize A over Υ and (27), for all $k \neq j$, $(k,l) \in E$, is sufficient.

$$\sum_{y=1}^{n} R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] \geq \sum_{y=1}^{n} R_y(i,j)\frac{\partial A}{\partial r_{y,k}(i,j)} \tag{27}$$

Proof: See Appendix B. ■

Based on the above sufficient condition, we developed a gradient-based algorithm for multiple TMs as an extension of [4]. At node k, the algorithm reduces the routing variables $\phi_{kl}(i,j)$ for which the quantity $\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]$ is large, and increases them for which the above quantity is small. The algorithm is described in Appendix C. Furthermore, in Appendix D, we proved that our algorithm converges to an optimal set of flow routes.

Using ratio variables as control variables, we now extend [3] to the multiple TM case by showing that the optimization problem is a convex optimization problem, and then solve it using convex optimization techniques.

With multiple TMs, link data rates F_y are linear combinations of B (see (8)). As a result, $\mathcal{B}$ is a convex polyhedron. As a simple extension from [3], the loop-free property remains valid with multiple TMs for the case of flow routing. When we restrict our consideration to loop-free B, the set of feasible loop-free route-sets is a convex, closed, and bounded set. From (8) and (12), we can see that cost A is a convex function of B. Thus, the problem is a convex optimization problem in the case of multiple TMs. Furthermore, when the link cost functions are approximated by piece-wise linear functions, the problem becomes a LP problem.

Note that with multiple TMs, we require many more control variables ($|V|(|V|-1)|E|$) compared to the single TM case ($|V||E|$ when using destination-based link data rates as control variables).

V. Optimal Destination Routing with multiple TMs

In this section, we explore ways of computing an optimal set of *destination routes* for multiple TMs. It is difficult to extend the existing gradient-based method [4] and convex optimization method [3], from the case of single TM to the case of multiple TMs. We show that it is NP-complete even to determine the feasibility of a set of multiple TMs. Thus we propose and evaluate a heuristic algorithm for computing routes.

Let us begin by considering the case where routing variables Φ are used as control variables. With a single TM, from any feasible set of loop-free routes, the gradient-based algorithm [4] converges to an optimal set of routes. However, with multiple TMs, we find local minima, which makes it hard to solve the problem using gradient-based methods. The following example demonstrates the existence of local minima. The example is based on network[3] G (shown in Figure 4) and two TMs R_1, R_2 (associated with weights $w_1 = w_2 = 0.5$),

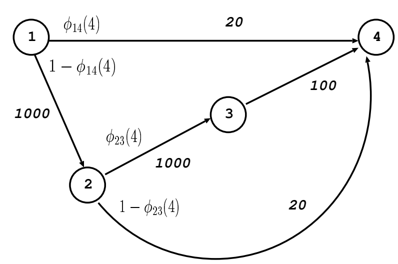

Fig. 4. A topology to illustrate the local-minima and non-convexity

$$R_1 = \begin{bmatrix} 0 & 0 & 0 & 8 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 83 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 5 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{28}$$

Note that routing variables ($\phi_{14}(4)$, $\phi_{23}(4)$) completely determine the set of routes, and thus the cost. With a single TM R_1 (R_2), Figure 5 shows cost A_1 (A_2), as a function of ($\phi_{14}(4)$, $\phi_{23}(4)$). We can see that there are no local minima. Thus, gradient-based methods can be used to solve the problem. However, with two TMs, we find local minima. Figure 6 shows A as a function of ($\phi_{14}(4)$, $\phi_{23}(4)$). We can see that the global optimal is at ($\phi_{14}(4) = 1$, $\phi_{23}(4) = 1$) and local minima is around ($\phi_{14}(4) \approx 0.5$, $\phi_{23}(4) = 0$). Hence, a gradient-based method gets stuck at this local minima point. Additionally, in Appendix E, we show that the ratio between the cost of local minima and that of global optima can be arbitrarily large.

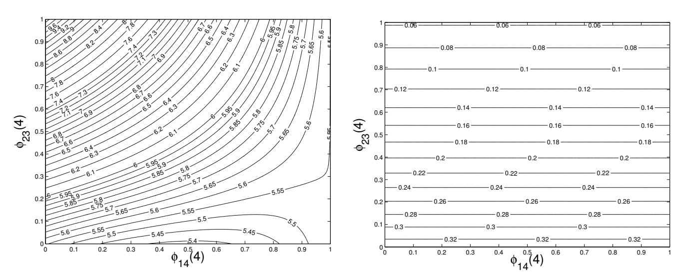

Fig. 5. A_1 (left), A_2 (right) as a function of routing variables

An alternative formulation of the optimization problem is to use ratio variables B as control variables. In the case of destination routing, although the cost A is a convex function of B, we find that the set of feasible route-sets $\mathcal{B}$ is not convex. To demonstrate a counter-example, note that ($B_{14}(1,4)$, $B_{12}(1,4)$, $B_{23}(1,4)$, $B_{23}(2,4)$) completely

[3]We use directed graph for ease of exposition.

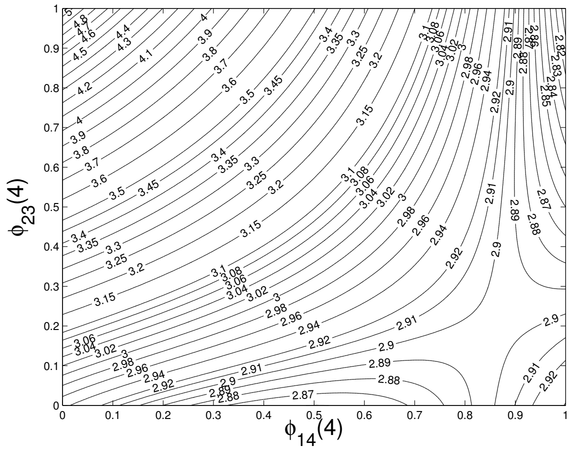

Fig. 6. $A = 0.5A_1 + 0.5A_2$ as a function of routing variables

determines the set of routes, and thus the cost. $(1, 0, 0, 0)$ and $(0, 1, 1, 1)$ are two feasible sets of destination ratio variables. However, the average of the two vectors, $(0.5, 0.5, 0.5, 0.5)$, is not a set of destination ratio variables. Since $\mathcal{B}$ is not convex, we cannot solve the problem as a convex optimization problem.

With multiple TMs, we next prove that it is NP-complete to determine whether the set of feasible destination route-sets Ψ (or $\mathcal{B}$) is empty or not. Our proof is given for the case of a set of two TMs.

Problem Description: Feasibility of a set of Two TMs in the case of Destination Routing (F2TDR).
Instance: Network[4] G, Integer-valued link capacity C, Integer-valued TM R_1, R_2.
Question: Is there a set of destination routes Φ (or B) of rational numbers in the set of feasible route-sets Ψ (or $\mathcal{B}$).

Theorem V.1: F2TDR is NP-complete.

Proof: As Ψ and $\mathcal{B}$ are equivalent (see Theorem III.2), and the relationship between Φ and B is rational (see (4)), we only prove the theorem for the case of Φ.

For TM R_1, R_2, given a set of destination routing variables Φ, we can calculate link data rates F_y, $y \in \{1, 2\}$ using equations (5) and (7), and check the feasibility in polynomial time. Thus, F2TDR belongs to NP. Next, it suffices to show: 3SAT $\propto$ F2TDR.

Let the clauses of the 3SAT problem be $U_1, \ldots, U_l$ and $x_1, \ldots, x_k, \bar{x}_1, \ldots, \bar{x}_k$ be the literals, where $l, k \geq 1$. Network G is constructed as follows. For each variable x_i, we construct a lobe shown in Figure 7. For each clause U_j, we create two nodes (s_j and t_j). s_j is connected to v^i_j, and v^i_{j+1} to t_j if and only if x_i appears in s_j. Also, s_j is connected to $\bar{v}^i_j$, and $\bar{v}^i_{j+1}$ to t_j if and only if $\bar{x}_i$ appears in s_j. The capacity of each link is 1.

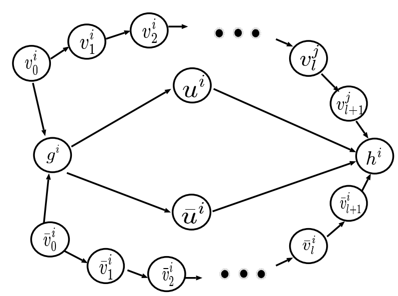

Fig. 7. Lobe for each variable x_i

Next, we construct two TMs R_1 and R_2. In R_1, $R_1(s_j, t_j) = 1$, $j \in \{1, \ldots, l\}$, and $R_1(v^i_0, h^i) = 1$, $R_1(u^i, h^i) = 1$, $i \in \{1, \ldots, k\}$. In R_2, $R_2(s_j, t_j) = 1$, $j \in \{1, \ldots, l\}$, and $R_2(\bar{v}^i_0, h^i) = 1$, $R_2(\bar{u}^i, h^i) = 1$, $i \in \{1, \ldots, k\}$.

(a) Assume that there is a feasible set of destination routing variables Φ satisfying the two TMs. Let $x_i = \lceil \phi_{g^i \bar{u}^i}(h^i) \rceil$. In Φ, assume traffic from s_j to t_j flows through lobe i. If it flows through link (v^i_j, v^i_{j+1}), then $\phi_{g^i \bar{u}^i}(h^i)$

[4]We use directed graph for ease of exposition.

must be 1, otherwise routing Φ is not feasible for R_1. Thus U_j is satisfied. If it flows through link $(\bar{v}_j^i, \bar{v}_{j+1}^i)$, a similar argument holds. This completes the proof that the expression is satisfiable.

(b) If the expression is satisfiable, we set $(\phi_{v_0^i g^i}(h^i), \phi_{\bar{v}_0^i g^i}(h^i), \phi_{g^i \bar{u}^i}(h^i))$ to be $(1,0,1)$ if x_i is 1 and $(0,1,0)$ otherwise. Since each clause U_j contains at least one literal x_i or $\bar{x}_i$ which is 1, traffic from s_j to t_j must be forwarded through either link (v_j^i, v_{j+1}^i) or link $(\bar{v}_j^i, \bar{v}_{j+1}^i)$ of lobe i depending on whether x_i or $\bar{x}_i$ is 1. Φ is then a feasible set of destination routing variables.

Thus, the 3SAT problem is satisfied if and only if there is a feasible set of destination routes Φ (or B) of rational numbers. ■

We have proved that F2TDR is NP-complete. The routing optimization problem is even harder. Consequently, we propose a heuristic algorithm to solve the problem.

As a guideline for our heuristic algorithm, we first obtain the following upper and lower bounds on the optimal cost in the case of a feasible set of n TMs.

Theorem V.2: Let $\hat{R} = \sum_y w_y R_y$ and $\check{R} = \max_y R_y$ where the max is element-wise. If $\check{R}$ is feasible, then the set of n TMs $R = \{R_1, \ldots, R_n\}$ is feasible, and $\hat{R}$ is feasible. Let A_y^O, $y \in \{1, \ldots, n\}$, $\hat{A}^O$, $\check{A}^O$ be the optimal cost for TM R_y, $\hat{R}$ and $\check{R}$ respectively, we have,

$$\hat{A}^O \le \sum_{y=1}^{n} w_y A_y^O \le A^{FO} \le A^{DO} \le \check{A}^O \tag{29}$$

where A^{FO} and A^{DO} are the optimal cost of the n TMs for flow routing and destination routing respectively.

Proof: First, we prove that $\hat{A}^O \le \sum_y w_y A_y^O$. Let B_y be a set of flow routes for TM R_y, $y \in \{1, \ldots, n\}$. We can construct a set of flow routes $\hat{B}$ for TM $\hat{R}$,

$$\hat{B}_{kl}(i,j) = \frac{\sum_y w_y R_y(i,j) B_{y,kl}(i,j)}{\sum_y w_y R_y(i,j)} \tag{30}$$

Let $f_{y,kl}(B_y)$ denote the link data rate for TM R_y given B_y, and let $\hat{f}_{kl}(\hat{B})$ denote the link data rate for TM $\hat{R}$ given $\hat{B}$. Combined with equation (8), we have

$$\hat{f}_{kl}(\hat{B}) = \sum_{y=1}^{n} w_y f_{y,kl}(B_y) \qquad (k,l) \in E \tag{31}$$

Because of the convexity of cost A as a function of link data rates, we have,

$$\hat{A}^O \le \sum_{y=1}^{n} w_y A_y^O \tag{32}$$

Second, we prove that $\sum_y w_y A_y^O \le A^{FO} \le A^{DO}$. Given any set of flow routes B, we have $A_y^O \le A_y$, $y \in \{1, \ldots, n\}$, thus $\sum_y w_y A_y^O \le A$. As a result,

$$\sum_{y=1}^{n} w_y A_y^O \le A^{FO} \tag{33}$$

As the set of destination route-sets is a subset of the set of flow route-sets. we have,

$$A^{FO} \le A^{DO} \tag{34}$$

Finally, we prove $A^{DO} \le \check{A}^O$. Given any set of destination route B, from (8), we have $f_{y,kl} \le \check{f}_{kl}$. Thus,

$$A^{DO} \le \check{A}^O \tag{35}$$

Combining the above steps yields (29). ■

In our heuristic algorithm, we compute an "expected TM" as the element-wise expectation of the TMs based on the perturbed weights. We then compute an optimal set of routes for this single "expected TM" and use this as our solution for the n TM problem. From theorem V.2, we see that the optimal cost of $\check{R}$ provides an upper bound to the problem. To ensure that the solution of our heuristic algorithm is also upper bounded by $\check{A}^O$, we incorporate $\check{R}$ as an extra perturbation dimension when we calculate the "expected TM". Formally, a perturbed weight vector is represented by $\omega = (\omega_1, \omega_2, \ldots, \omega_n, \omega_{n+1})$, $\omega_y \in [0, 1]$, $y \in \{1, \ldots, n+1\}$, $\sum_{y=1}^{n+1} \omega_y = 1$, where ω_y, $y \in \{1, \ldots, n\}$ is the perturbed weight of TM R_y, and ω_{n+1} is the perturbed weight for $\check{R} = \max_y R_y$. We use Ω to represent the set of perturbed weight vectors. Given a perturbed weight vector $\omega \in \Omega$, $a)$ we calculate the "expected TM" $\bar{R} = \sum_{y=1}^{n} \omega_y R_y + \omega_{n+1}\check{R}$, $b)$ and then find an optimal set of destination routes for the "expected TM" $\bar{R}$ (Note, there may be more than one optimal set of destination routes for $\bar{R}$; we randomly select one of them). $c)$ Finally, we evaluate the cost A for the set of n TMs given the set of routes derived in step b.

Let g denote the mapping from Ω to the cost A as described by the above procedure. Our heuristic algorithm finds the perturbed weight vector $\omega \in \Omega$ with the minimum cost $A^{O(\Omega)}$. The set of destination routes achieving $A^{O(\Omega)}$ is then the "good" set of routes for the set of n TMs returned by our heuristic algorithm.

Because of the contribution of $\check{R}$ to $\bar{R}$, we have,

$$A^{O(\Omega)} \leq \check{A}^O, \quad \text{if } \check{R} \text{ is feasible} \tag{36}$$

Our heuristic algorithm consists of two stages: a global stage and a local stage. The global stage examines the perturbed weight vector space Ω and identifies promising perturbed weight vectors. The local stage focuses on the promising perturbed weight vectors and attempts to quickly improve the quality. Similar methods are used in [5] and [11] to solve OSPF routing optimization problems.

In the global stage, uniform searching effectively identifies promising perturbed weight vectors [11]. For function $g(\omega)$, $\omega \in \Omega$ with a range of $[A^{O(\Omega)}, A^{M(\Omega)}]$, the distribution function of g is defined as:

$$\gamma_\Omega(A) = \frac{m(\{\omega \in \Omega \mid g(\omega) \leq A\})}{m(\Omega)} \tag{37}$$

where $A \in [A^{O(\Omega)}, A^{M(\Omega)}]$ and $m(\cdot)$ denotes *Lebesgue measure*, a measure of the size of a set. Assuming $A^{(r)} \in [A^{O(\Omega)}, A^{M(\Omega)}]$ such that $\gamma_\Omega(A^{(r)}) = r$, $r \in [0, 1]$, an *r-percentile* set in Ω is defined as:

$$\tau_\Omega(r) = \{\omega \in \Omega \mid g(\omega) \leq A^{(r)}\} \tag{38}$$

Consider l randomly generated perturbed weight vectors ω^1, ω^2, ..., ω^l, and let $\acute{\omega}^1$, $\acute{\omega}^2$, ..., $\acute{\omega}^l$ be the corresponding perturbed weight vectors ranked in increasing order of g. According to [12], the probability of $\acute{\omega}^k$ in $\tau_\Omega(r)$ is,

$$P(\acute{\omega}^k \in \tau_\Omega(r)) = \int_0^r \frac{l!}{(k-1)!(l-k)!} x^{k-1}(1-x)^{l-k} dx \tag{39}$$

It takes 183 samples for the 10th top ranked sample, $\acute{\omega}^{10}$, to reach $0.1 - percentile$ with a probability of 99%.

During the global stage, we uniformly sample 183 independent perturbed weight vectors through a method given in [13]. The most promising 10 samples are then passed to the local stage to improve the quality.

During the local stage, we use an iterative procedure to make improvement. The perturbed weight vector space Ω is discretized and a neighborhood structure $\mathcal{N}(\omega)$ is defined on it. Starting from a promising perturbed weight vector ω, at each iteration, the neighbor perturbed weight vector with the lowest cost is chosen for the next iteration. In order that our algorithm not become trapped in a local minimal, it allows non-improvement moves so that the search proceeds in a larger neighborhood. The search stops when the number of iteration reaches certain threshold (100 is used for each promising weight vector in the results of this paper) or the quality of result is satisfactory.

We define the neighborhood structure $\mathcal{N}(\omega)$ as follows. First, ω is discretized so that ω_i can only take a value from $\{0, 1/\delta, 2/\delta, \ldots, 1\}$. Second, ω' is a neighbor of ω if they differ in 2 dimensions. The maximum number of neighbors for a perturbed weight vector is thus $n(n+1)$.

We present our results obtained using a synthetic network (50 nodes and 156 links). The synthetic network is produced using the generator GT-ITM [14], based on a model of Calvert *et al.* [15] [16]. This model places nodes in a unit square, thus generating a distance between each pair of nodes. Links are divided into two classes: local access links and long distance links. The capacity is 200 for a local access link and 1000 for a long distance link. We generate TMs using the methods in [5]. For each node $v \in V$, we pick two random numbers $O_v, Q_v \in [0, 1]$. Furthermore, for each node pair (v_i, v_j), we pick a random number $Z_{(v_i,v_j)} \in [0, 1]$. For v_i and v_j with Euclidean distance l, the rate traffic between v_i and v_j is

$$\alpha O_{v_i} Q_{v_j} Z_{(v_i,v_j)} e^{-l/2L} \tag{40}$$

where α is scale parameter and L is the largest Euclidean distance among all pair of nodes. The values O_v, Q_v models the degree to which a node generates or attracts traffic. The distance l models the traffic locality. In this model, there is more traffic between close pairs of nodes.

With $n = 3$ TMs associated with weights $w_1 = w_2 = w_3 = 1/3$, we compare the results of our heuristic algorithm with the lower bound ($\sum_y w_y A_y^O$), the upper bound ($\check{A}^O$), and a baseline algorithm *SINGLE*, which chooses the best set of routes out of the n sets of routes, that optimize for each TM. In our heuristic algorithm, we choose different precisions ($\delta = 1$, and $\delta = 10$). When $\delta = 1$, we choose between the set of routes optimized for $\check{R}$ and the set of routes given by *SINGLE*. Thus, we call it $SINGLE + MAX$. When $\delta = 10$, we are searching a "good" set of routes by mixing the n TMs and $\check{R}$. We call the resulting heuristic $MIX(SINGLE + MAX)$. In our experiments, we use piece-wise linear functions (approximation of $M/M/1$) as link cost functions, and we use AMPL/CPLEX [17] to compute the optimal set of destination routes for a single "expected TM".

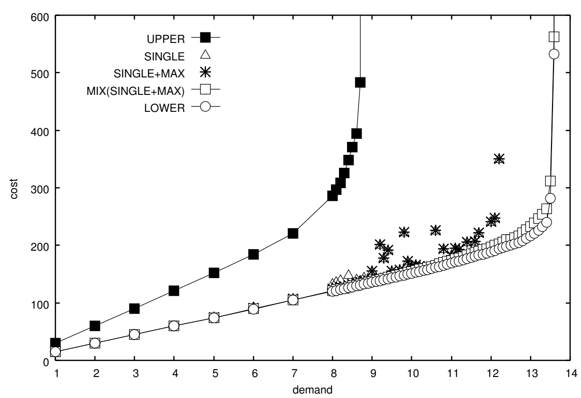

Fig. 8. Experiment results with a synthetic network (50 nodes, 156 links) and 3 TMs

The results of our experiments are presented in Figure 8 with different scalings of the TMs. In the experiments, we see that the cost rises as demand increases. All curves start off flat, and then, start increasing rapidly. And the demand becomes too large to be feasible as link capacity constraints are reached. This behavior is somewhat similar to that of a single link.

We can see that the curve of *SINGLE*, *SINGLE+MAX* and *MIX(SINGLE+MAX)* are upper-bounded by the curve of *UPPER* and lower-bounded by the curve of *LOWER*. We also see that our heuristic algorithm *MIX(SINGLE+MAX)* does very well, always falling within 11% of *LOWER*.

When $\check{R} = max_y R_y$ is feasible (scaled up to 8.8), we can see that the cost generated by *SINGLE+MAX* is mostly lower than *SINGLE*, and is close to *MIX(SINGLE+MAX)*. This indicates that the optimal set of routes for the element-wise max TM $\check{R}$ can be a "good" solution to the problem. When $\check{R}$ is no longer feasible, the cost returned by *SINGLE* is the same as the cost returned by *SINGLE+MAX*, as expected. As demand increases (scaled between 8.9 and 12.2), we can see that *SINGLE* may return high cost solution (70% more than the cost of *MIX(SINGLE+MAX)*) or cannot even find a feasible set of destination routes. When demand approaches the limit that the network can carry (scaled between 12.2 and 13.6), *SINGLE* cannot find a feasible set of destination routes while *MIX(SINGLE+MAX)* can.

VI. Conclusion and Discussions

The key contributions are summarized as follows:

1. We extended the formulation of the route optimization problem from the case of a single TM to the case of multiple TMs. Specifically, we extended the formulation in [4] that uses routing variables as control variables, and we extended the formulation in [3] that uses ratio variables as control variables.

2. We identified the fundamental difference in the route optimization problem between the case of a single TM and the case of multiple TMs. We showed that unlike the single TM case, with multiple TMs, the optimal cost of flow routing may be lower than that of destination routing, and the optimal cost of destination route-sets with loops may be lower than that of loop-free destination route-sets.

3. In the case of flow routing, we extended the solution methods for a single TM to the case of multiple TMs. With routing variables as control variables, we extended [4] to solve the problem with multiple TMs using gradient-based methods; with ratio variables as control variables, we extended [3] to solve the problem using convex optimization techniques, and thus solve the problem using LP when link costs are piece-wise linear functions.

4. In the case of destination routing, we demonstrated the inherent difficulties of the problem with multiple TMs. We identified local minima when routing variables are used as control variables. Local minima make it difficult to solve the problem using gradient-based methods. We also demonstrated that the set of feasible route-sets is not convex when ratio variables are used as control variables. i.e., the optimization problem is not a convex optimization problem. Finally, we proved that it is NP-complete even to determine the feasibility of a set of multiple TMs.

5. In the case of destination routing, we proposed and evaluated a heuristic algorithm.

With multiple TMs in the case of flow routing, although we have shown that the problem can be solved by extending existing methods, the extremely large number of control variables can hardly be handled by a single computer (We have 3 million control variables for a 100-node network with 300 links). As a result, distributed computation might be desirable to solve the problem. The extension of Gallager's work may be a useful starting point although it is necessary for all routers to get derivative information for all TMs.

REFERENCES

[1] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus/IETF, "Requirements for traffic engineering over MPLS (RFC 2702)," 1999.

[2] J. Moy/IETF, "OSPF version 2 (RFC 2328)," 1998.

[3] David G. Cantor and Mario Gerla, "Optimal routing in a packet-switched computer network," *IEEE Transactions on Computer*, vol. C-23, no. 10, pp. 1062–1069, 1974.

[4] Robert G. Gallager, "A minimum delay routing algorithm using distributed computation," in *IEEE Transaction on Communications*, January 1977, pp. 73–85.

[5] Bernard Fortz and Mikkel Thorup, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM*, 2000.

[6] David Applegate and Edith Cohen, "Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs," in *ACM SIGCOMM*, 2003.

[7] M. Roughan, M. Thorup, and Y. Zhang, "Traffic engineering with estimated traffic matrices," in *Internet Measurement Conference*, 2003.

[8] Bernard Fortz and Mikkel Thorup, "Optimizing OSPF/IS-IS weights in a changing world," in *IEEE Journal on Selected Areas in Communications*, 2002, pp. 756–767.

[9] Yufei Wang, Zheng Wang, and Leah Zhang, "Internet traffic engineering without full mesh overlaying," in *INFOCOM*, 2001.

[10] Ashwin Sridharan, Roch Guerin, and Christophe Diot, "Achieving near-optimal traffic engineering solutions for current OSPF/IS-IS networks," in *INFOCOM*, 2003.

[11] Tao Ye and Shivkumar Kalyanaraman, "A recursive random search algorithm for large-scale network parameter configuration," in *ACM Sigmetrics*, 2003.

[12] George Casella and Roger L. Berger, *Statistical Inference*, Duxbury, 2 edition, 2002.

[13] JR. Marshall Hall, *Combinatorial Theory*, John Wiley & Sons, 2 edition, 1986.

[14] E.W. Zegura, "GT-ITM: Georgia tech internetwork topology models (software)," .

[15] E.W. Zegura, K.L. Calvert, and S. Bhattacharjee, "How to model an internetwork," in *Proc. 15th IEEE Conf. on Computer Commnications*, 1996.

[16] K. Calvert, M. Doar, and E.W. Zegura, "Modeling internet topology," in *IEEE Communication Magazine*, 1997, vol. 35, pp. 160–163.

[17] "AMPL/CPLEX," http://www.ilog.com.

[18] David G. Luenberger, *Linear and Nonlinear Programming*, Kluwer Academic Publishers, 2003.

VII. APPENDIX A

Proof of Theorem IV.1:

As we introduce the set of dummy variables $r_y, y \in \{1, \dots, n\}$, (5) becomes,

$$t_{y,k}(i,j) = \mathbf{1}(k=i)R_y(i,j) + r_{y,k}(i,j) + \sum_m t_{y,m}(i,j)\phi_{mk}(i,j) \quad (41)$$

Without loss of generality, for source-destination pair (i, j), take the source node i to be the 1st of the $|V|$ nodes, take the destination node j to be the $|V|$th of the $|V|$ nodes, and drop the argument (i, j), we thus rewrite (41) as

$$t_{y,k} = \mathbf{1}(k=1)R_y + r_{y,k} + \sum_{m=1}^{|V|-1} t_{y,m}\phi_{mk} \tag{42}$$

Let $\mathcal{T}_y = (t_{y,1}, \ldots, t_{y,|V|-1})$, $\mathcal{R}_y = (R_y + r_{y,1}, r_{y,2}, \ldots, r_{y,|V|-1})$, and let φ be the $|V|-1 \times |V|-1$ matrix with components ϕ_{kl} $(1 \le k, l \le |V|-1)$. Equation (42) for $1 \le k \le |V|-1$ is then $\mathcal{T}_y(I-\varphi) = \mathcal{R}_y$. From [4], we know that $I - \varphi$ must have an inverse. Following the same procedure as [4], we have,

$$\mathcal{T}_y = \mathcal{R}_y(I-\varphi)^{-1} \tag{43}$$

$$\frac{\partial t_{y,k}}{\partial \mathcal{R}_{y,m}} = [(I-\varphi)^{-1}]_{mk} \tag{44}$$

$$t_{y,k} = \sum_m \frac{\partial t_{y,k}}{\partial \mathcal{R}_{y,m}} \mathcal{R}_{y,m} \tag{45}$$

$$\frac{\partial t_{y,k}}{\partial \phi_{mq}} = \frac{\partial t_{y,k}}{\partial \mathcal{R}_{y,q}} t_{y,m} \tag{46}$$

Take into account that r is a set of dummy variables. i.e., only derivative information when $r_y = \{0, \ldots, 0\}$ is useful. For fixed R_y, equation (44), (45) and (46) become,

$$\frac{\partial t_{y,k}}{\partial r_{y,m}} = \frac{\partial t_{y,k}}{\partial \mathcal{R}_{y,m}} = [(I-\varphi)^{-1}]_{mk} \tag{47}$$

$$t_{y,k} = \frac{\partial t_{y,k}}{\partial r_{y,1}} R_{y,1} \tag{48}$$

$$\frac{\partial t_{y,k}}{\partial \phi_{mq}} = \frac{\partial t_{y,k}}{\partial r_{y,q}} t_{y,m} = \frac{\partial t_{y,k}}{\partial r_{y,q}} \frac{\partial t_{y,m}}{\partial r_{y,1}} R_{y,1} \tag{49}$$

Next, we show that (23), repeated below with the source node and destination node again taken to be 1 and $|V|$, has a unique solution.

$$\frac{\partial A}{\partial r_{y,k}} = \sum_l \phi_{kl} \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}} \right] \tag{50}$$

Following the same proof as in [4], we solve the above equations and get $\frac{\partial A}{\partial r_{y,k}}$ as a continuous function in Φ,

$$\frac{\partial A}{\partial r_{y,k}} = \sum_m \frac{\partial t_{y,m}}{\partial r_{y,k}} \sum_q w_y \phi_{mq} D'_{mq}(f_{y,mq}) = \sum_{m,q} w_y \frac{\partial f_{y,mq}}{\partial r_{y,k}} D'_{mq}(f_{y,mq}) \tag{51}$$

Differentiating A directly with (7), (11) and (12), we get the same unique solution.

Finally we calculate $\frac{\partial A}{\partial \phi_{y,k}}$ directly using (7), (11) and (12),

$$\frac{\partial A}{\partial \phi_{kl}} = \sum_{y,m,q} w_y D'_{mq}(f_{y,mq}) \phi_{mq} \frac{\partial t_{y,m}}{\partial \phi_{kl}} + \sum_y w_y D'_{kl}(f_{y,kl}) t_{y,k} \tag{52}$$

$$= \sum_y t_{y,k} \left\{ \left[\sum_{m,q} w_y D'_{mq}(f_{y,mq}) \phi_{mq} \frac{\partial t_{y,m}}{\partial r_{y,l}} \right] + w_y D'_{kl}(f_{y,kl}) \right\} \tag{53}$$

$$= \sum_y t_{y,k} \left[\frac{\partial A}{\partial r_{y,l}} + w_y D'_{kl}(f_{y,kl}) \right] \tag{54}$$

We have used (49) and (51) to derive (54), which is the same as (24). This is clearly continuous in Φ given the continuity of $t_{y,k}$ and $\frac{\partial A}{\partial r_{y,l}}$, and the proof is complete. ■

VIII. APPENDIX B

Proof of Theorem IV.2:

First we show that (26) is a necessary condition to minimize A by assuming that Φ does not satisfy (26). This means that there is some i, j, k, l, and m such that

$$\phi_{kl}(i,j) > 0, \quad \frac{\partial A}{\partial \phi_{kl}(i,j)} > \frac{\partial A}{\partial \phi_{km}(i,j)} \tag{55}$$

Since these derivatives are continuous, a sufficiently small increase in $\phi_{km}(i,j)$ and corresponding decrease in $\partial\phi_{kl}(i,j)$ will decrease A, thus establishing that Φ does not minimize A.

Next we show that (27), repeated below, is a sufficient condition to minimize A.

$$\sum_{y=1}^{n} R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] \geq \sum_{y=1}^{n} R_y(i,j)\frac{\partial A}{\partial r_{y,k}(i,j)} \tag{56}$$

Suppose that Φ satisfies (56) and has link data rates F_y and node data rates T_y, $y \in \{1,\dots,n\}$. Let Φ^* be any other set of routing variables with link data rates F^*_y and node data rates T^*_y, $y \in \{1,\dots,n\}$. Define

$$f_{y,kl}(\lambda) = (1-\lambda)f_{y,kl} + \lambda f^*_{y,kl} \tag{57}$$

$$A(\lambda) = \sum_{y=1}^{n} w_y \sum_{(k,l)\in E} D_{kl}(f_{y,kl}(\lambda)) \tag{58}$$

Since each link cost D_{kl} is a convex, non-decreasing function of the link data rate, therefore A, is convex in λ, and hence

$$\left.\frac{dA}{d\lambda}\right|_{\lambda=0} \leq A(\Phi^*) - A(\Phi) \tag{59}$$

Since Φ^* is arbitrary, proving that $dA/d\lambda \geq 0$ at $\lambda = 0$ will complete the proof. From (57) and (58),

$$\left.\frac{dA}{d\lambda}\right|_{\lambda=0} = \sum_{y,(k,l)} w_y D'_{kl}(f_{y,kl})\Big(f^*_{y,kl} - f_{y,kl}\Big) \tag{60}$$

We now show that

$$\sum_{y,(k,l)} w_y D'_{kl}(f_{y,kl})f^*_{y,kl} \geq \sum_{y,i,j} R_y(i,j)\frac{\partial A}{\partial r_{y,i}(i,j)} \tag{61}$$

Note from (56) that

$$\sum_{y,l} w_y R_y(i,j) D'_{kl}(f_{y,kl})\phi^*_{kl}(i,j) \geq \sum_{y} R_y(i,j)\left[\frac{\partial A}{\partial r_{y,k}(i,j)} - \sum_{l}\frac{\partial A}{\partial r_{y,l}(i,j)}\phi^*_{kl}(i,j)\right] \tag{62}$$

From Theorem III.1 and equation (6), we know that for source routes Φ^*, $\exists B^*$, so that

$$t^*_{y,k}(i,j) = R_y(i,j)\left[\sum_{m} B^*_{mk}(i,j) + 1(k=i)\right] \tag{63}$$

Multiplying both sides of (62) by $(\sum_m B^*_{mk}(i,j) + 1(k=i))$, summing over i, j, k and recalling that $f^*_{y,kl} = \sum_{i,j} t^*_{y,k}(i,j)\phi^*_{kl}(i,j)$ (see (7)) , we obtain

$$\sum_{y,k,l} w_y D'_{kl}(f_{kl})f^*_{kl} \geq \sum_{y,i,j,k} t^*_{y,k}(i,j)\frac{\partial A}{\partial r_{y,k}(i,j)} - \sum_{y,i,j,k,l} t^*_{y,k}(i,j)\phi^*_{kl}(i,j)\frac{\partial A}{\partial r_{y,l}(i,j)} \tag{64}$$

From (5), $\sum_k t^*_{y,k}(i,j)\phi^*_{kl}(i,j) = t^*_{y,l}(i,j) - \mathbf{1}(i=l)R_y(i,j)$. Substituting this into the rightmost term of (64) and canceling, we get (61). Note that the only inequality used here was (62), and that if Φ is substituted for Φ^*, this becomes an equality for the equation (23). Thus

$$\sum_{y,(k,l)} w_y D'_{kl}(f_{y,kl}) f_{y,kl} = \sum_{y,i,j} R_y(i,j) \frac{\partial A}{\partial r_{y,i}(i,j)} \tag{65}$$

Substituting (61) and (65) into (60), we see that $dA/d\lambda \geq 0$ at $\lambda = 0$, completing the proof. ■

IX. Appendix C

We present a distributed gradient-based algorithm to minimize cost A. In our algorithm, a node k is implemented by a process k. All processes have a copy of TMs R_y and weights w_y, $y \in \{1,\ldots,n\}$. Besides, process k maintains a set of local routing variables $\{\phi_{kl}(i,j),\ (k,l) \in E,\ i,j \in V\}$. Next, we use the term "node" and "process" interchangeably.

The algorithm breaks into two parts: a protocol between processes to calculate the $\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]$ and an algorithm for modifying the local routing variables. We discuss the protocol part first.

In order to see how process k can calculate $\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]$, for traffic originating from i to j, define node m be the *downstream* from node q if there is a routing path from q to j through m (i.e., a path with positive routing variables on each link). Similarly, we define q as the *upstream* from m if m is downstream from q. In our algorithm, we focus on the set of loop-free route-sets.

The protocol used to calculate the $\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]$ is further divided into two rounds: in the forward round, for each source and destination pair i,j, each process k waits until it has received the value $t_{y,m}(i,j)\phi_{mk}(i,j)$, $y \in \{1,\ldots,n\}$ from each of its upstream neighbor processes m. Then process k calculates the node traffic rate $t_{y,k}(i,j)$ using (5) and broadcasts this to all of its downstream neighbor processes; in the backward round, for each source and destination pair i,j, each process k waits until it has received the value $\frac{\partial A}{\partial r_{y,l(i,j)}}$, $y \in \{1,\ldots,n\}$ from each of its downstream neighbors l. Then process k calculates the $\frac{\partial A}{\partial r_{y,k(i,j)}}$ and the $\sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]$ using (7) and using (23). and broadcasts the former to all of its upstream neighbor processes. It is easy to see that the protocol is free of deadlocks because Φ is loop-free.

The algorithm Γ, on each iteration, maps the current routing variable set Φ into a new set $\Phi^1 = \Gamma(\Phi)$. In order to maintain loop-free property, for traffic originating from i destined to j and each node k, the algorithm maintains a set $H_k(i,j)$ of blocked nodes k for which $\phi_{kl}(i,j) = 0$ and the algorithm is not permitted to increase $\phi_{kl}(i,j)$ from 0. For notation convenience we include l such that $(k,l) \notin E$ in the set $H_k(i,j)$. We first define and discuss the algorithm and then define the sets $H_k(i,j)$.

The algorithm is described as follows. For $l \in H_k(i,j)$,

$$\phi^1_{kl}(i,j) = 0, \Delta_{kl}(i,j) = 0. \tag{66}$$

For $l \notin H_k(i,j)$, define

$$a'_{kl}(i,j) = \sum_y R_y(i,j)\left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right] - \min_{m \notin H_k(i,j)} \left\{\sum_y R_y(i,j)\left[w_y D'_{km}(f_{y,km}) + \frac{\partial A}{\partial r_{y,m}(i,j)}\right]\right\} \tag{67}$$

$$a_{kl}(i,j) = \frac{a'_{kl}(i,j)}{\sum_y w_y R_y(i,j)} \tag{68}$$

$$\Delta_{kl}(i,j) = \min\{\phi_{kl}(i,j), \frac{\eta a_{kl}(i,j)}{\sum_y w_y t_{y,k}(i,j)}\} \tag{69}$$

where η is a scale parameter of Γ to be discussed later. Let $l_k^M(i,j)$ be a value of m that achieves the minimization in (69). Then

$$\phi_{kl}^1(i,j) = \begin{cases} \phi_{kl}(i,j) - \Delta_{kl}(i,j) & l \neq l_k^M(i,j) \\ \phi_{kl}(i,j) + \sum_{l \neq l_k^M(i,j)} \Delta_{kl}(i,j) & otherwise. \end{cases} \quad (70)$$

The algorithm reduces the fraction of traffic sent on non-optimal links and increases the fraction on the best link. The amount of reduction, given by $\Delta_{kl}(i,j)$, is proportional to $a_{kl}(i,j)$, with the restriction that $\phi_{kl}^1(i,j)$ cannot be negative. In turn $a_{kl}(i,j)$ is the difference between the marginal cost of the traffic originating from node i to node j at node k using link (k,l) and using the best link. Note that as the sufficient condition (27) is approached, the changes get small, as desired. The amount of reduction is also inversely proportional to $\sum_y w_y t_{y,k}(i,j)$. The reason for this is that the change in link traffic under TM R_y is related to $w_y \Delta_{kl}(i,j) t_{y,k}(i,j)$. Thus when $\sum_y w_y t_{y,k}(i,j)$ is small, $\Delta_{kl}(i,j)$ can be changed by a large amount without greatly affecting the marginal link cost. Finally the changes depend on the scale factor η. For η very small, convergence of the algorithm is guaranteed, as shown in Theorem IX.2, but rather slow. As η increases, the speed of convergences increases but the danger of no convergence increases.

We now complete the definition of algorithm Γ by defining the sets $H_k(i,j)$. First define a routing variable $\phi_{kl}(i,j)$ to be *improper* if $\phi_{kl}(i,j) > 0$ and $\sum_y R_y(i,j) \partial A / \partial r_{y,k}(i,j) \leq \sum_y R_y(i,j) \partial A / \partial r_{y,l}(i,j)$. We have already said that $H_k(i,j)$ includes only k for which $\phi_{kl}(i,j) = 0$, and thus, from (23),

$$\min_{l \notin H_k(i,j)} \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)} \right] \leq \sum_y R_y(i,j) \frac{\partial A}{\partial r_{y,k}(i,j)} \quad (71)$$

Assuming weighted positive marginal link costs, $\sum_y R_y(i,j) \partial A / \partial r_{y,k}(i,j) < \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \partial A / \partial r_{y,l}(i,j) \right]$ if $\phi_{kl}(i,j)$ is improper, and we see that the algorithm always reduces improper routing variables. In fact, since $\sum_y R_y(i,j) \frac{\partial A}{\partial r_{y,k}(i,j)}$ is the weighted marginal cost for the traffic originating from i to j at node k, we would expect weighted marginal cost to decrease as we move downstream, and improper routing variables should be rather atypical.

For a given source and destination pair (i,j), the set of weighted marginal costs $\sum_y R_y(i,j) \frac{\partial A}{\partial r_{y,k}(i,j)}$ forms an ordering of the nodes k. Note that if there are no improper routing variables, this ordering is consistent with the downstreaming partial ordering.

Similar to [4], if Φ is loop-free and $\Phi^1 = \Gamma(\phi)$ contains a loop for some source destination pair i,j, then the following two conditions must hold.

1) The loop contains some link (k,l) for which $\phi_{kl}(i,j) = 0$, $\phi_{kl}^1(i,j) > 0$, and $\sum_y R_y(i,j) \partial A / \partial r_{y,k}(i,j) > \sum_y R_y(i,j) \partial A / \partial r_{y,l}(i,j)$.
2) The loop contains some link (l,m) for which $\phi_{lm}(i,j)$ is improper and for which $\phi_{lm}^1(i,j) > 0$.

The first condition reiterates that some routing variables must be increased from 0 to form a loop and that the algorithm only increases routing variables on links to nodes with smaller marginal cost. The second make use of the fact that if nodes are ranked by marginal cost, $\sum_y R_y(i,j) \partial A / \partial r_{y,k}(i,j)$, then it is impossible to move around a loop of nodes and have marginal cost monotonically decrease.

Definition: The set $H_k(i,j)$ is the set of nodes l for which either $\phi_{kl}(i,j) = 0$ and l is blocked relative to (i,j) or $(k,l) \notin L$. A node l is blocked relative to (i,j) if for traffic originating from i destined to j, node l has a routing path to j containing some link (m,q) for which $\phi_{mq}(i,j)$ is improper and

$$\phi_{mq}(i,j) \geq \eta \frac{\sum_y R_y(i,j) [w_y D'_{mq}(f_{y,mq}) + \frac{\partial A}{\partial r_{y,q}(i,j)} - \frac{\partial A}{\partial r_{y,m}(i,j)}]}{[\sum_y w_y t_{y,m}(i,j)][\sum_y w_y R_y(i,j)]} \quad (72)$$

The reason for (72) can be seen from (69) and (71). $\Delta_{mq}(i,j) = \phi_{mq}(i,j)$ and $\phi_{mq}^1(i,j) = 0$, so that (m,q) can not be part of a loop for traffic originating from i destined to j.

Theorem IX.1: If the marginal link costs D'_{kl} are positive and Φ is loop-free, then $\Phi^1 = \Gamma(\Phi)$ is loop-free.

Proof: Following similar proof in [4]. ■

The protocol required for a process k to determine the set $H_k(i,j)$ is as follows. Each process k, when it calculates $\partial A/\partial r_{y,k}$, determines, for each downstream l, if ϕ_{kl} is improper and satisfies (72). If any downstream neighbor satisfies these conditions, node k adds a special tag to its broadcast of $\partial A/\partial r_{y,k}$. The node k also adds the special tag if the received value $\partial A/\partial r_{y,l}$ from any downstream l contained a tag. In this way all nodes upstream of k also send the tag. The set $H_k(i,j)$ is then the set of nodes l for with either $(k,l) \notin E$ or the received $\partial A/\partial r_{y,l}$ was tagged.

Theorem IX.2: Assume that for all $(k,l) \in E$, $D_{kl}(f_{y,kl})$ has a positive first derivative and nonnegative second derivative for $0 \leq f_{y,kl} < c_{kl}$ and that $\lim_{f_{y,kl}\to c_{kl}} = \infty$. For every positive number A_0 there exists a scale factor η for Γ such that if Φ^0 satisfies $A(\Phi^0) \leq A_0$, then

$$\lim_{m\to\infty} A(\Phi^m) = \min_{\Phi}(A(\Phi)) \tag{73}$$

This is proved in Appendix D. Note that η depends on some upper bound A_0 to A; this is natural, since when the link data rates are very close to capacity, small changes in the link data rates cause large changes in marginal cost. The proof use a ridiculously small value of η to guarantee convergence under all conditions and experimental work is necessary to determine practical values for η.

X. APPENDIX D

Proof of Theorem IX.2:

We prove Theorem IX.2 through a sequence of seven lemmas. The first five establish the descent properties of the algorithm, the sixth establishes a type of continuity condition, showing that if Φ does not minimize A, the for any Φ^* in a neighborhood of Φ, $A(\Gamma^m(\Phi^*)) < A(\Phi)$ for some m. The seventh lemma is a new global convergence theorem which does not require continuity in the algorithm Γ; Lemmas X.6 and X.7 together establish Theorem IX.2.

Let Φ be an arbitrary set of routing variables satisfying $A(\Phi) < A_0$ for some A_0. Let $\Phi^1 = \Gamma(\Phi)$ and let T_y, F_y, T_y^1, F_y^1, $y \in \{1,\ldots,n\}$ be the node and link data rates corresponding to Φ and Φ^1, respectively. Let F_y^λ, $(0 \leq \lambda \leq 1)$,$y \in \{1,\ldots,n\}$ be defined by $f_{y,kl}^\lambda = (1-\lambda)f_{y,kl} + \lambda f_{y,kl}^1$, and let

$$A(\lambda) = \sum_{y,k,l} w_y D_{y,kl}(f_{y,kl}^\lambda) \tag{74}$$

From the Taylor remainder theorem,

$$A(\Phi^1) - A(\Phi) = \frac{dA(\lambda)}{d\lambda}\bigg|_{\lambda=0} + \frac{1}{2}\frac{d^2A(\lambda)}{d\lambda^2}\bigg|_{\lambda=\lambda^*} \tag{75}$$

where λ^* is some number between 0 and 1. The continuity of the second derivative above will be obvious from the proof of Lemma X.4, which upper bounds that term. The first three lemmas deal with $\frac{dA(\lambda)}{d\lambda}|_{\lambda=0}$.

Lemma X.1:

$$\frac{dA(\lambda)}{d\lambda}\bigg|_{\lambda=0} = \sum_{i,j,k,l} -\Delta_{kl}(i,j)a_{kl}(i,j)\sum_y w_y t_{y,k}^1(i,j) \tag{76}$$

Proof: Using the definitions of $a_{kl}(i,j)$ and $\Delta_{kl}(i,j)$ in (68) and (69),

$$\begin{aligned}
\sum_l \Delta_{kl}(i,j)a_{kl}(i,j) &= \frac{1}{\sum_y w_y R_y(i,j)} \sum_{l\neq l_k^M(i,j)} [\phi_{kl}(i,j) - \phi_{kl}^1(i,j)]\left\{\sum_y R_y(i,j)\left[w_y D_{kl}'(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]\right.\\
&\quad \left. - \min_{m\notin H_k(i,j)}\left\{\sum_y R_y(i,j)\left[w_y D_{km}'(f_{y,km}) + \frac{\partial A}{\partial r_{y,m}(i,j)}\right]\right\}\right\}\\
&= \frac{1}{\sum_y w_y R_y(i,j)} \sum_l [\phi_{kl}(i,j) - \phi_{kl}^1(i,j)]\left\{\sum_y R_y(i,j)\left[w_y D_{kl}'(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]\right\} \qquad (77)\\
&= \frac{1}{\sum_y w_y R_y(i,j)}\left\{\sum_y R_y(i,j)\frac{\partial A}{\partial r_{y,k}(i,j)} - \sum_{l,y}\phi_{kl}^1(i,j)R_y(i,j)\left[w_y D_{kl}'(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)}\right]\right\} \qquad (78)
\end{aligned}$$

In (77), we have used (70) to extend the sum over all l and in (78), we have used (23). Multiplying both sides of (78) by $\sum_y w_y t^1_{y,k}(i,j)$, summing over i, j, k, and using (5), (6) and (7), we get

$$\begin{aligned}\sum_{y,i,j,k,l} \Delta_{kl}(i,j) a_{kl}(i,j) w_y t^1_{y,k}(i,j) &= \sum_{y,i,j,k} t^1_{y,k}(i,j) \frac{\partial A}{\partial r_{y,k}(i,j)} - \sum_{y,k,l} f^1_{y,kl} w_y D'_{kl}(f_{y,kl}) \\ &\quad - \sum_{y,i,j,l} \left(t^1_{y,l}(i,j) - 1(l=i) R_y(i,j) \right) \frac{\partial A}{\partial r_{y,l}(i,j)} \\ &= -\sum_{y,k,l} f^1_{y,kl} w_y D'_{kl}(f_{y,kl}) + \sum_{y,i,j} R_y(i,j) \frac{\partial A}{\partial r_{y,i}(i,j)} \quad (79) \\ &= \sum_{y,k,l} (f_{y,kl} - f^1_{y,kl}) w_y D'_{kl}(f_{y,kl}) \quad (80) \\ &= -\left.\frac{dA(\lambda)}{d\lambda}\right|_{\lambda=0} \quad (81)\end{aligned}$$

We have used (65) to get (80), and (81) from (74), completing the proof. ■

Lemma X.2:

$$\left.\frac{dA(\lambda)}{d\lambda}\right|_{\lambda=0} \le -\frac{1}{\eta(|V|-1)^3} \sum_{i,j,k} \Delta^2_k(i,j) \left[\sum_y w_y t_{y,k}(i,j) \right]^2 \le -\frac{1}{\eta(|V|-1)^3} \sum_{y,i,j,k} w^2_y \Delta^2_k(i,j) t^2_{y,k}(i,j) \quad (82)$$

where

$$\Delta_k(i,j) = \sum_l \Delta_{kl}(i,j) \quad (83)$$

Proof: From the definition of Δ in (69), $-a_{kl}(i,j) \le -\sum_y w_y t_{y,k}(i,j) \Delta_{kl}(i,j)/\eta$. Substituting this into (76) yields

$$\begin{aligned}\left.\frac{dA(\lambda)}{d\lambda}\right|_{\lambda=0} &\le -\frac{1}{\eta} \sum_{i,j,k,l} \Delta^2_{kl}(i,j) \sum_y w_y t_{y,k}(i,j) \sum_y w_y t^1_{y,k}(i,j) \\ &\le -\frac{1}{(|V|-1)\eta} \sum_{i,j,k} \Delta^2_k(i,j) \sum_y w_y t_{y,k}(i,j) \sum_y w_y t^1_{y,k}(i,j) \quad (84)\end{aligned}$$

where (84) follows from Cauchy's inequality, $(\sum_l \alpha_l \beta_l)^2 \le (\sum \alpha^2_l)(\sum \beta^2_l)$, with $\alpha_l = 1$, $\beta_l = \Delta_{kl}(i,j)$, and the sum over $l \ne k$.

Now define $t^*_{y,k}(i,j)$ as the node traffic rate originating from i destined to j at node k if the routing variables $\phi_{kl}(i,j)$ (for $l \ne l^M_k(i,j)$) are reduced by $\Delta_{kl}(i,j)$ but $\phi_{kl}(i,j)$ for $l = l^M_k(i,j)$ is not increased. Mathematically $t^*_{y,k}(i,j)$ satisfies

$$t^*_{y,k}(i,j) = \sum_l t^*_{y,l}(i,j)[\phi_{lk}(i,j) - \Delta_{lk}(i,j)] + 1(k=i) R_y(i,j) \quad (85)$$

This has a unique solution because of the loop-freedom of Φ. Subtracting (85) from (5) results in

$$t_{y,k}(i,j) - t^*_{y,k}(i,j) = \sum_l [t_{y,l}(i,j) - t^*_{y,l}(i,j)] \phi_{lk}(i,j) + \sum_l t^*_{y,l}(i,j) \Delta_{lk}(i,j) \quad (86)$$

From (45), using $\sum_l t^*_{y,l}(i,j)\Delta_{lk}(i,j)$ for $\mathcal{R}_{y,m}(i,j)$, and from (47),

$$t_{y,k}(i,j) - t^*_{y,k}(i,j) = \sum_l \frac{\partial t_{y,k}(i,j)}{\partial r_{y,l}(i,j)} \sum_m t^*_{y,m}(i,j) \Delta_{ml}(i,j) \quad (87)$$

Since Φ is loop-free, $\partial t_{y,k}(i,j)/\partial r_{y,l}(i,j) \le 1$. Also if $\partial t_{y,k}(i,j)/\partial r_{y,l}(i,j) > 0$, then l is upstream of k for traffic originating from i destined to j and $\phi_{kl}(i,j)$ (and hence $\Delta_{kl}(i,j)$) is zero. Thus

$$t_{y,k}(i,j) - t^*_{y,k}(i,j) \le \sum_l \sum_{m \ne k} t^*_{y,m}(i,j) \Delta_{ml}(i,j) = \sum_{m \ne k} t^*_{y,m}(i,j) \Delta_m(i,j) \quad (88)$$

Multiplying the left side by $\Delta_k(i,j) \le 1$ preserves the inequality, yielding

$$t_{y,k}(i,j)\Delta_k(i,j) \le \sum_m t^*_{y,m}(i,j)\Delta_m(i,j) \tag{89}$$

Multiplying w_y and summing up over y,

$$\sum_y w_y t_{y,k}(i,j)\Delta_k(i,j) \le \sum_m \sum_y w_y t^*_{y,m}(i,j)\Delta_m(i,j) \tag{90}$$

Since the right-hand side of (87) is nonnegative, we also have $t_{y,k}(i,j) \ge t^*_{y,k}(i,j)$. Multiplying w_y and summing up over y,

$$\sum_y w_y t_{y,k}(i,j) \ge \sum_y w_y t^*_{y,k}(i,j) \tag{91}$$

The following lemma proved by [4] will be used for further proof.

Lemma X.3: Let α_k, $\beta_k (1 \le k \le m)$ be nonnegative numbers satisfying $\alpha_k \le \sum_l \beta_l$; $\alpha_k \ge \beta_k$ for $1 \le k \le m$. Then

$$\sum_{k=1}^{m} \alpha_k \beta_k \ge \frac{1}{m^2} \sum_k \alpha_k^2 \tag{92}$$

Now let $\alpha_k = \sum_y w_y t_{y,k}(i,j)\Delta_k(i,j)$ and $\beta_k = \sum_y w_y t^*_{y,k}(i,j)\Delta_k(i,j)$. Since these terms are nonzero only for $k \ne j$, we can take $m = |V| - 1$. Since the conditions of the lemma are satisfied for this choice,

$$\sum_k \Delta_k^2(i,j) \sum_y w_y t_{y,k}(i,j) \sum_y w_y t^*_{y,k}(i,j) \ge \frac{1}{(|V|-1)^2} \sum_k \Delta_k^2(i,j) \left[\sum_y w_y t_{y,k}(i,j)\right]^2. \tag{93}$$

Since $t^1_{y,k}(i,j) \ge t^*_{y,k}(i,j)$, we can substitute (93) into (84), and proved the first inequality of Lemma 2.

$$\left.\frac{dA(\lambda)}{d\lambda}\right|_{\lambda=0} \le -\frac{1}{\eta(|V|-1)^3} \sum_{i,j,k} \Delta_k^2(i,j) \left[\sum_y w_y t_{y,k}(i,j)\right]^2 \tag{94}$$

Additionally, since $w_y t_{y,k}(i,j) \ge 0$, we have

$$\left[\sum_y w_y t_{y,k}(i,j)\right]^2 \ge \sum_y w_y^2 t_{y,k}^2(i,j) \tag{95}$$

The second inequality of Lemma X.2 is derived by substituting (95) into (94). Proof is then completed. ■

Lemma X.4: Let M be an upper bound of $D''_{kl}(f^\lambda_{y,kl})$ over all y, k, l and over $0 \le \lambda \le 1$. Then for any λ, $0 \le \lambda \le 1$,

$$\frac{d^2 A(\lambda)}{d\lambda^2} \le M(|V|+2)(|V|-1)^2|V| \sum_{y,i,j,k} w_y \Delta_k^2(i,j) t_{y,k}^2(i,j) \tag{96}$$

Proof: The bound M must exist because $D''_{kl}(f^\lambda_{y,kl})$ is a continuous function of λ over the compact region $0 \le \lambda \le 1$. Taking the second derivative, we get

$$\frac{d^2 A(\lambda)}{d\lambda^2} = \sum_{y,k,l} w_y D''_{kl}(f^\lambda_{y,kl})[f^1_{y,kl} - f_{y,kl}]^2 \le \sum_{y,k,l} w_y M[f^1_{y,kl} - f_{y,kl}]^2 \tag{97}$$

We now upper bound $|f^1_{y,kl} - f_{y,kl}|$ by first upper bounding $|t^1_{y,k}(i,j) - t_{y,k}(i,j)|$. As in the proof of Lemma X.2, we have

$$\begin{aligned} t^1_{y,k}(i,j) - t_{y,k}(i,j) &= \sum_l [t^1_{y,l}(i,j) - t_{y,l}(i,j)]\phi^1_{lk}(i,j) + \sum_l t_{y,l}(i,j)[\phi^1_{lk}(i,j) - \phi_{lk}(i,j)] \\ &= \sum_l \frac{\partial t^1_{y,k}(i,j)}{\partial r_{y,l}(i,j)} \sum_m t_{y,m}(i,j)[\phi^1_{ml}(i,j) - \phi_{ml}(i,j)] \end{aligned} \tag{98}$$

Since $0 \le \partial t^1_{y,k}(i,j)/\partial r_{y,l}(i,j) \le 1$, we can upper bound this by

$$t^1_{y,k}(i,j) - t_{y,k}(i,j) \le \sum_m t_{y,m}(i,j)\Delta_m(i,j) \tag{99}$$

We can lower bound (99) in the same way, considering only terms in which $\phi^1_{ml}(i,j) - \phi_{ml}(i,j) < 0$, and this leads to

$$|t^1_{y,k}(i,j) - t_{y,k}(i,j)| \le \sum_m t_{y,m}(i,j)\Delta_m(i,j) \tag{100}$$

$$f^1_{y,kl} - f_{y,kl} = \sum_{i,j}[t^1_{y,k}(i,j) - t_{y,k}(i,j)]\phi^1_{kl}(i,j) + t_{y,k}(i,j)[\phi^1_{kl}(i,j) - \phi_{kl}(i,j)]$$

$$|f^1_{y,kl} - f_{y,kl}| \le \sum_{i,j}\sum_m t_{y,m}(i,j)\Delta_m(i,j)\phi^1_{kl}(i,j) + \sum_{i,j} t_{y,k}(i,j)|\phi^1_{kl}(i,j) - \phi_{kl}(i,j)| \tag{101}$$

The double sum in (101) has at most $(|V|-1)^3$ nonzero terms ($i \ne j$, $k \ne j$, $m \ne j$) and the second sum at most $(|V|-1)^2$ terms. Using Cauchy's inequality on both terms together, we get

$$|f^1_{y,kl} - f_{y,kl}|^2 \le |V|(|V|-1)^2 \left\{ \sum_{i,j,m} t^2_{y,m}(i,j)\Delta^2_m(i,j)[\phi^1_{kl}(i,j)]^2 + \sum_{i,j} t^2_{y,k}(i,j)[\phi^1_{kl}(i,j) - \phi_{kl}(i,j)]^2 \right\}$$

$$\sum_l |f^1_{y,kl} - f_{y,kl}|^2 \le |V|(|V|-1)^2 \left\{ \sum_{i,j,m} t^2_{y,m}(i,j)\Delta^2_m(i,j) + 2\sum_{i,j} t^2_{y,k}(i,j)\Delta^2_k(i,j) \right\} \tag{102}$$

Summing over k, we get

$$\sum_{k,l} |f^1_{y,kl} - f_{y,kl}|^2 \le |V|(|V|-1)^2(|V|+2) \sum_{i,j,k} t^2_{y,k}(i,j)\Delta^2_k(i,j) \tag{103}$$

Multiplying by w_y, summing over y, and substituting the result in (97),we get (96) completing the proof. ■

Lemma X.5: For given A_0, define

$$M = \max_{k,l} \max_{f:D_{kl}(f) \le A_0/\min w_y} D''_{y,kl}(f) \tag{104}$$

$$\eta = [M|V|^7]^{-1} \min_y w_y. \tag{105}$$

Then for all Φ such that $A \le A_0$,

$$A(\Phi^1) - A(\Phi) \le -\frac{M|V|^7}{2(|V|-1)^3} \sum_{y,i,j,k} w_y \Delta^2_k(i,j) t^2_{y,k}(i,j). \tag{106}$$

Proof: Temporarily let M be as defined in Lemma X.4. Combining Lemma X.2 and Lemma X.4,

$$A(\Phi^1) - A(\Phi) \le \left[-\frac{M|V|^7}{(|V|-1)^3} + \frac{M(|V|+2)(|V|-1)^2|V|}{2} \right] \sum_{y,i,j,k} w_y \Delta^2_k(i,j) t^2_{y,k}(i,j). \tag{107}$$

The second term in brackets above is less than half the magnitude of the first term, yielding (106). It follows that $A(\Phi^1) \le A(\Phi) \le A_0$. By convexity then $D_{kl}(f^\lambda_{y,kl}) \le A_0/w_y \le A_0/\min w_y$ for $0 \le \lambda \le 1$. Thus M as given in (104) satisfies the condition on M in Lemma X.4, completing the proof. ■

Lemma X.6: Let the scale factor η satisfy (105) for a given A_0 and let Φ be an arbitrary set of routing variables that does not minimize A and satisfies $A(\Phi) \le A_0$. Given this Φ, $\exists \epsilon > 0$ and an m, $1 \le m \le |V|$, such that for all Φ^* satisfying $|\Phi - \Phi^*| < \epsilon$,

$$A(\Gamma^m(\Phi^*)) < A(\Phi) \tag{108}$$

Proof: We consider three cases. The first is the typical case in which no blocking occurs and $A(\Gamma(\Phi)) < A(\Phi)$, the second is the case in which blocking occurs, and the third is the case in which $A(\Gamma(\Phi)) = A(\Phi)$.

Case 1: No blocking; $\Delta_k(i,j)t_k(i,j) > 0$ for some y, i, j, k. If no nodes are blocked for Φ, then by definition of blocking (72), there is a neighborhood of Φ^* around Φ for which no blocking occurs. In this neighborhood,

$$a'_{kl}(i,j) = \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)} \right] - \min_{m \notin H_k(i,j)} \left\{ \sum_y R_y(i,j) \left[w_y D'_{km}(f_{y,km}) + \frac{\partial A}{\partial r_{y,m}(i,j)} \right] \right\} \quad (109)$$

which is continuous in Φ. It follows from (68) and (69) that $\Delta_{kl}(i,j)$ is continuous in Φ, and the upper bound to $A(\Gamma(\Phi)) - A(\Phi)$ in (106) is continuous in Φ. Since by assumption the bound in (106) is strictly negative, there is a neighborhood of Φ^* around Φ for which

$$A(\Gamma(\Phi^*)) - A(\Phi^*) < -\frac{M|V|^7}{4(|V|-1)^3} \sum_{y,i,j,k} w_y \Delta_k^2(i,j) t_{y,k}^2(i,j) \quad (110)$$

where $\Delta_k(i,j)$ and $t_{y,k}(i,j)$ correspond to the given Φ. Choose ϵ small enough so that (110) is satisfied for $|\Phi - \Phi^*| < \epsilon$ and also so that

$$|A(\Phi^*) - A(\Phi)| < \frac{M|V|^7}{4(|V|-1)^3} \sum_{y,i,j,k} w_y \Delta_k^2(i,j) t_{y,k}^2(i,j) \quad (111)$$

Combining this with (110), we have (108) for $m = 1$.

Case 2: Blocking occurs. For any Φ, we can use (23) to lower bound $a'_{kl}(i,j)$ by

$$a'_{kl}(i,j) \geq \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)} - \frac{\partial A}{\partial r_{y,k}(i,j)} \right] \quad (112)$$

$$\Delta_{kl}(i,j) \sum_y w_y t_{y,k}(i,j) \geq \min \left\{ \phi_{kl}(i,j) \sum_y w_y t_{y,k}(i,j), \frac{\eta \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)} - \frac{\partial A}{\partial r_{y,k}(i,j)} \right]}{\sum_y w_y R_y(i,j)} \right\} \quad (113)$$

The lower bounds above are continuous functions of Φ. Since blocking occurs in Φ, there is some i, j, k, l such that both

$$\sum_y R_y(i,j) \left[\frac{\partial A}{\partial r_{y,l}(i,j)} - \frac{\partial A}{\partial r_{y,k}(i,j)} \right] \geq 0 \quad (114)$$

and

$$\phi_{kl}(i,j) \sum_y w_y t_{y,k}(i,j) \geq \frac{\eta \sum_y R_y(i,j) \left[w_y D'_{kl}(f_{y,kl}) + \frac{\partial A}{\partial r_{y,l}(i,j)} - \frac{\partial A}{\partial r_{y,k}(i,j)} \right]}{\sum_y w_y R_y(i,j)} \quad (115)$$

Combining (113) to (115)

$$\Delta_{kl}(i,j) \sum_y w_y t_{y,k}(i,j) \geq \frac{\eta \sum_y R_y(i,j) w_y D'_{kl}(f_{y,kl})}{\sum_y w_y R_y(i,j)} \quad (116)$$

Since the right-hand side of (113) is continuous in Φ, there is a neighborhood of Φ^* around Φ for which

$$\Delta^*_{kl}(i,j) \sum_y w_y t^*_{y,k}(i,j) \geq \frac{\eta}{2} \frac{\sum_y R_y(i,j) w_y D'_{kl}(f_{y,kl})}{\sum_y w_y R_y(i,j)} \quad (117)$$

And thus, for some y,

$$\Delta^*_{kl}(i,j) t^*_{y,k}(i,j) \geq \frac{\eta}{2} \frac{\sum_y R_y(i,j) w_y D'_{kl}(f_{y,kl})}{\sum_y w_y R_y(i,j)} \quad (118)$$

Equation (108), for $m = 1$, now follows in the same way as in case 1.

Case 3: $\Delta_{kl}(i,j)t_{y,k}(i,j) = 0$ for all y,i,j,k,l. Let Φ_3 be the set of Φ for which $\Delta_{kl}(i,j)t_{y,k}(i,j) = 0$ for all i,j,k,l. Let $\Phi^{(l)} = \Gamma^l(\Phi)$ for the given Φ and let $m \geq 2$ be the smallest integer such that $\Phi^{(m-1)} \notin \Phi_3$. We first show that $m \leq |V|$. Note first that for any $\Phi \in \Phi_3$, Γ changes $\phi_{kl}(i,j)$ only for y,i,j,k such that $t_{y,k}(i,j) = 0$ and thus the node data rates and link data rates cannot change. $\sum_y R_y(i,j)\partial A/\partial r_{y,k}(i,j)$ can change, however, and as we shall see later, must change for some y,i,j,k if Φ does not minimize A.

Now consider $\Phi^{(q)}$ $(0 \leq q \leq m-2)$, where $\Phi^{(0)}$ denotes the original Φ. Since $\Phi^{(q)} \in \Phi_3$, $\Delta_{kl}^{(q)}(i,j) > 0$ implies that $t_{y,k}(i,j) = 0$. From (69), $\phi_{kl}^{(q)}(i,j) = \Delta_{kl}^{(q)}(i,j)$ and $\phi_{kl}^{(q+1)}(i,j) = 0$. For a given i,j,k, all $\phi_{kl}^{(q)}(i,j)$ are reduced to 0 except for the l which minimizes $\sum_y R_y(i,j)[w_y D_{kl}^{'}(f_{y,kl}) + \partial A/\partial r_{y,l}(i,j)]$. Thus, using (23),

$$\sum_y R_y(i,j)\frac{\partial A(\Phi^{(q+1)})}{\partial r_{y,k}(i,j)} = \min_l \left(\sum_y R_y(i,j)\left[w_y D_{kl}^{'}(f_{y,kl}) + \frac{\partial A(\Phi^{(q)})}{\partial r_{y,l}(i,j)}\right]\right) \leq \sum_y R_y(i,j)\frac{\partial A(\Phi^{(q)})}{\partial r_{y,k}(i,j)} \tag{119}$$

Since this equation is satisfied for all q, $0 \leq q \leq m-2$, we see that $\sum_y R_y(i,j)\partial A(\Phi^{(q)})/\partial r_{y,k}(i,j)$ can be reduced on iteration q only if $\sum_y R_y(i,j)\partial A(\Phi^{(q-1)})/\partial r_{y,l}(i,j)$ is reduced on iteration $q-1$ for some l such that $\sum_y R_y(i,j)\partial A(\Phi^{(q-1)})/\partial r_{y,l}(i,j) < \sum_y R_y(i,j)\partial A(\Phi^{(q)})/\partial r_{y,k}(i,j)$. This reduction at node l however implies a reduction at some node l' of smaller differential cost at iteration $q-2$ and so forth. Since this sequence of differential cost is decreasing with decreasing q and since (from (119)) the differential cost at a given node is nondecreasing with decreasing q, each node in the sequence must be distinct. Since there are $|V|-1$ nodes other than the given destination available for such a sequence, the initial q in such a sequence satisfies $q \leq |V|-2$. On the other hand, if $\sum_y R_y(i,j)\partial A(\Phi^{(q)})/\partial r_{y,k}(i,j)$ is unchanged for all y,i,j,k, we see from (119) that $\Phi^{(q)}$ satisfies the sufficient conditions to minimize A and then Φ also minimizes A contrary to our hypothesis; thus we must have $m \leq |V|$.

Now observe that the middle expression in (119), for $q = 0$, is a continuous function of Φ and consequently $\partial A(\Phi^{(1)})/\partial r_{y,k}(i,j)$ is a continuous function of Φ for all y,i,j,k. It follows by induction that $\partial A(\Phi^{(l)})/\partial r_{y,k}(i,j)$ is a continuous function of Φ for all y,i,j,k and for $l \leq m-1$. Finally $\Phi^{(m-1)} \notin \Phi_3$, so it must satisfy the conditions of case 1 or 2; it will be observed that the analysis there apply equally to $\Phi^{(m-1)}$ because of the continuity of $\partial A(\Phi^{(m-1)})/\partial r_{y,k}(i,j)$ as a function of Φ. This completes the proof. ∎

Our last lemma will be stated in greater generality than required since it is a global convergence theorem for algorithms that avoids the usual continuity constraint on the algorithm. (See Luenberger [18]) for a good discussion of global convergence).

Lemma X.7: Let Υ be a compact region of Euclidean N space. Let Γ be a mapping from Υ into Υ and let A be a continuous real valued function in Υ. Assume that $A(\Gamma(\Phi)) \leq A(\Phi)$ for all $\Phi \in \Upsilon$. Let A^O be the minimum of A over Υ and let Υ^O be the set of $\Phi \in \Upsilon$ such that $A(\Phi) = A^O$. Assume that for every $\Phi \in \Upsilon - \Upsilon^O$, there is an $\epsilon > 0$ and an integer $m \geq 1$ such that for all $\Phi^* \in \Upsilon$ satisfying $|\Phi - \Phi^*| \leq \epsilon$, we have $A(\Gamma^m(\Phi^*)) \leq A(\Phi)$. Then for all $\Phi \in \Upsilon$,

$$\lim_{m\to\infty} A(\Gamma^m(\Phi)) = A^O. \tag{120}$$

Proof: See [4]. ∎

Proof of Theorem IX.2: Let Υ be the set of loop-free routing variable Φ such that $A(\Phi) \leq A_0$. We have verified that Γ maps loop-free routing variables into loop-free routing variables, and from Lemma X.5, $A(\Gamma(\Phi)) \leq A(\Phi)$ for $\Phi \in \Upsilon$. Thus Γ is mapping from Υ into Υ. It is obvious that Υ is bounded and easy to verify that any limit of loop-free variables with $A(\Phi) \leq A_0$ is also loop-free with $A(\Phi) \leq A_0$. Thus Υ is compact. The final assumption of Lemma X.7 is established by Lemma X.6. Thus Lemma X.7 asserts the conclusion of Theorem IX.2. ∎

XI. Appendix E

We show an example where the ratio between the cost of local minima and the cost of global optima can be arbitrarily large.

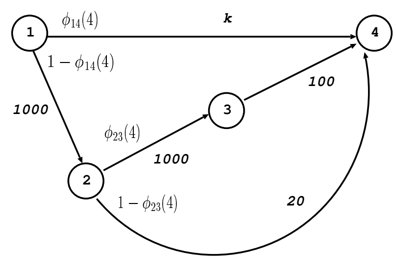

Fig. 9. A topology to illustrate the arbitrary ratio between local minima and global optima

Given integer $k \geq 100$, Figure 9 shows a network G. There are two TMs R_1 and R_2 with weights $w_1 = w_2 = 0.5$,

$$R_1 = \begin{bmatrix} 0 & 0 & 0 & k-(100k)^{-0.5} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 100-k^{-1} \\ 0 & 0 & 0 & 0 \end{bmatrix} \qquad R_2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{121}$$

Note that $\phi_{14}(4)$ and $\phi_{23}(4)$ determine packet forwarding, thus determine cost A. Given ($\phi_{14}(4)$, $\phi_{23}(4)$), the link data rate $f_{y,kl}$ is,

$$\begin{aligned} f_{1,12} &= R_1(1,4)(1-\phi_{14}(4)) \\ f_{1,14} &= R_1(1,4)\phi_{14}(4) \\ f_{1,23} &= R_1(1,4)(1-\phi_{14}(4))\phi_{23}(4) \\ f_{1,24} &= R_1(1,4)(1-\phi_{14}(4))(1-\phi_{24}(4)) \\ f_{1,34} &= R_1(1,4)(1-\phi_{14}(4))\phi_{23}(4)+R_1(3,4) \\ f_{2,12} &= f_{2,14} = 0 \\ f_{2,23} &= f_{2,34} = R_2(2,4)\phi_{23}(4) \\ f_{2,24} &= R_2(2,4)(1-\phi_{23}(4)) \end{aligned} \tag{122}$$

When $\phi_{14}(4) = 1$, $\phi_{23}(4) = 1$, we have $\frac{\partial A}{\partial \phi_{14}(4)} > 0$, $\frac{\partial A}{\partial \phi_{23}(4)} > 0$. Therefore, ($\phi_{14}(4) = 1$, $\phi_{23}(4) = 1$) is a local minimal. The cost of the relative minimal ($\phi_{14}(4) = 1$, $\phi_{23}(4) = 1$), A^{REL}, satisfies

$$A^{REL} > p_2 D_{14}(f_{1,14}) = 5k^{1.5} - 0.5 \tag{123}$$

When $\phi_{14}(4) = \frac{1}{k-(100k)^{-0.5}}$, $\phi_{23}(4) = 0$, we have

$$\begin{aligned} f_{1,12}, f_{2,12} &\leq 1 \\ f_{1,14}, f_{2,14} &< k-1 \\ f_{1,23}, f_{2,23} &= 0 \\ f_{1,24}, f_{2,24} &= 1 \\ f_{1,34}, f_{2,34} &\leq 100-k^{-1} \end{aligned}$$

Therefore,

$$\begin{aligned} A &< D_{12}(1) + D_{14}(k-1) + D_{23}(0) + D_{24}(1) + D_{34}(100-k^{-1}) \\ &= 1/999 + (k-1) + 0 + 1/19 + (100k-1) < 101k-2 \end{aligned} \tag{124}$$

Combining (123) and (124),

$$\lim_{k\to\infty} \frac{A^{REL}}{A^{OPT}} = \infty \tag{125}$$