Optimal Routing with Multiple Traffic Matrices Tradeoff between Average Case and Worst Case Performance

Chun Zhang[a], Zihui Ge[b], Jim Kurose[a], Yong Liu[c] and Don Towsley[a]

[a]Dept. of Computer Science, University of Massachusetts, Amherst, MA 01003 {czhang, kurose, towsley}@cs.umass.edu

[b]AT&T Research, 180 Park Avenue, P.O. Box 971, Florham Park, NJ 07932 gezihui@research.att.com

[c]Dept. of Electrical & Computer Engineering Polytechnic University Brooklyn, NY 11201 yongliu@poly.edu

Abstract

In this paper, we consider the problem of finding an "efficient" and "robust" set of routes in the face of changing/uncertain traffic. The changes/uncertainty in exogenous traffic are characterized by multiple traffic matrices. Our goal is to find a set of routes that results in good average case performance over the set of traffic matrices, while avoiding bad worst case performance for any single traffic matrix. With multiple traffic matrices, previous work aims solely to optimize the average case performance [1], or the worst case performance [2]. For a given set of traffic matrices, different sets of routes offer a different tradeoff between the average case and the worst case performance. In this paper, we quantify the performance of a routing configuration at both network level and link level. We propose a simple metric - a weighted sum of the average case and the worst case performance - to control the tradeoff between these two considerations. Despite of its simple form, this metric is very effective. We prove that optimizing routing using this metric has desirable properties, such as the average case performance being a decreasing, convex and differentiable function to the worst case performance. By extending previous work [1][3], we derive methods to find the optimal routes with respect to the proposed metric for two classes of intra-domain routing protocols: MPLS and OSPF/IS-IS. We evaluate our approach with data collected from an operational tier-1 ISP. For MPLS, we find that there exists significant tradeoff (e.g., $15\% - 23\%$ difference) between optimizing solely on the average case performance and solely on the worst case performance. Our approach can identify solutions that can dramatically improve the worst case performance ($13\% - 15\%$) while only slightly sacrificing the average case performance ($2.2\% - 3\%$), in comparison to that by optimizing solely on the average case performance. For OSPF/IS-IS, we still find a significant difference between the two optimization objectives, however, a fine-grained tradeoff is difficult to achieve due to the limited control that OSPF/IS-IS provide.

Index Terms

Traffic Engineering, Routing Optimization, Tradeoff, Average Case Performance, Worst Case Performance

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 cost function (such as link utilization or packet delay). The most commonly used intra-domain Internet routing protocols today are Open Shortest Path First (OSPF) [4] and intermediate system-intermediate system (IS-IS) [5]. In OSPF/IS-IS, each link is associated with a positive weight. Traffic is routed along the shortest paths. In case of ties where several outgoing links are on shortest paths to the destination, the flow is split evenly among them. As shown by [6], OSPF/IS-IS does not support arbitrary distribution of flow between source and destination, and may incur high cost. Multi-protocol Label Switching (MPLS) [7] has been introduced as a more flexible routing protocol. In MPLS, the routing path and splitting fraction can be arbitrarily chosen for traffic flow based on source and destination addresses.

A traffic matrix (TM) specifies the data rate between every pair of ingress and egress points. A number of works [8][6][9] have focused on calculating an optimal set of routes for a single TM. For a given TM, those works consider minimizing link utilization of the most congested link [9], or minimizing the *network cost* [8][6], characterized by 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 OSPF/IS-IS, the problem is shown to be NP-hard [6]. With MPLS, the problem can be formalized and solved as a convex optimization problem [8] or a linear programming (LP) problem [6].

Recently, for a large-scale Internet with bursty demands, multiple TMs have been used to characterize the change of traffic (different traffic rates at different times of the day) or traffic estimation uncertainty in the exogenous traffic [2][1]. The routing problem is to find a single set of routes to optimize the performance over multiple TMs. With multiple TMs, existing work either focuses on the average case performance, or the worst case TM performance. [1] aims to minimize the expected link/network cost. [2] aims to minimize the worst case link utilization. However, optimizing the average case performance and worst case performance often conflict with each other. Solely minimizing the average case performance may lead to high worst case performance. Also, if we are too conservative, focusing too much on the worst case performance, we may substantially decrease the average case performance as well.

In this paper, we consider the routing optimization problem of finding an "efficient" and "robust" solution in the face of changing/uncertain traffic. Intuitively, we may prefer a set of routes that results in good average case performance with multiple TMs, and not bad worst case performance for any single TM. In order to achieve this, we focus on the tradeoff between the average case and worst case performance. We examine this tradeoff at both link-level and network-level. The link-level focus is on the congestion level of individual links, whereas the network-level focus is on the combined congestion level at all links.

At the link level, we use *expected link cost* to quantify the average case performance, and *worst case link cost* the worst case performance. At the network level , we use *expected network cost* to quantify the average case performance, and *worst case network cost* the worst case performance. In this paper, we propose a simple metric: the weighted sum of average case and worst case performance, to control the tradeoff between these two considerations. We show our metric's effectiveness by proving that the average case performance is a convex, decreasing, continuous, and differentiable function of the worst case performance. We solve our proposed metric for two general inter-area routing approaches, MPLS and OSPF/IS-IS by extending work [1][3]. We evaluate our metric using data derived from an operational tier-1 ISP. For MPLS, we find a considerable tradeoff between average case and worst case performance ($15\% - 23\%$) of the optimal average case/worst case performance. Those tradeoff provides us a solution with both good average case performance and good worst case performance. Compared to the solution with best average case performance, but bad worst case performance, the tradeoff solution dramatically improves the worst case performance ($13\% - 15\%$ of the optimal worst case performance), while only sacrificing the average case performance a little ($2.2\% - 3\%$ of the optimal average case performance). For OSPF, we still find a significant difference between the two optimization objectives, however, a fine-grained tradeoff is difficult to achieve due to the limited control that OSPF/IS-IS provide.

Motivated by a commonly used robust optimization framework [10], we also investigate another routing optimization approach which aims to reduce the performance sensitivity relative to the change of TMs. We use the variance of network/link cost to model the performance variance. We find that it is problematic to adopt this approach because the variance term is not tied to an absolute performance measure, and thus minimizing the performance variance may leads to bad performance for all TMs.

The remainder of this paper is organized as follows. In section II, we review related work. In section III, we introduce tradeoff metric. In section IV, we demonstrate the tradeoff metric effectiveness by showing that average case performance is a convex, decreasing, continuous and differentiable function of worst case performance. We also show how to compute the optimal set of routes according to this metric by extending [1][3]. In section V, we show result using data derived from an operational tier-1 ISP. In section VI, we discuss the deficit of the performance variance metric. Section VII concludes the paper.

II. Context and Related Work

In this section, we first discuss the different performance metrics used in routing optimization with multiple TMs. Then we review the existing routing optimization solutions for those metrics.

A. Performance Metrics

We first review the performance metrics used for routing optimization with a single TM. Then we review the performance metrics in the case of multiple TMs. With a single TM, performance is quantified at either link level or network level [8][9]. Link level performance metrics focus on the congestion level of individual links, whereas network level metrics focus on the combined congestion of all links. More precisely, at the link level, the TM performance is characterized by *the link utilization of the most congested link* [9]. At the network level, performance for a given TM is characterized by *the network cost*, computed as the sum of all link costs, each of which is an increasing convex function of link data rate; or characterized by *the expected link cost*, which differs from *the network cost* by a constant ratio (the number of links of the network) [8].

Recently, for a large-scale Internet with bursty demands, multiple TMs have been used to characterize the change of traffic (different traffic rates at different time of the day) or traffic estimation uncertainty in the exogenous traffic. With multiple TMs, existing work either focuses on the average case performance over all TMs [1], or the worst case TM performance [2]. With multiple TMs, the average case performance is characterized by the *expected TM cost*, or *expected link cost*, which differs from the former by a constant ratio (the number of links of the network) [1]. With multiple TMs, the worst case performance can be addressed at either the link level or the network level. On one hand, link level worst case performance is characterized by the *the link utilization of the most congested link over all TMs* [2] [11]. On the other hand, network level worst case performance is characterized by the *the worst case network cost over all TMs*. To our knowledge, we are the first to focus on network level worst case performance.

In existing works, average case performance and worst case performance are individually addressed. However, focusing exclusively on one may result in bad result for the other. In order to find a set of routes with good average case performance and not bad worst case performance, we focus on the tradeoff between the average case performance and worst case performance. We examine this tradeoff at both link level and network level. At link level, we use *expected link cost* to quantify the average case performance, and *worst case link cost* the worst case performance. At network level, we use *expected network cost* to quantify the average case performance, and *worst case network cost* the worst case performance.

B. Routing Optimization Solutions

In this section, we discuss the existing routing optimization solutions. Specifically, we review the solution methods for routing optimization under MPLS and OSPF/IS-IS routing protocols.

MPLS provides the most flexible routing capability. Under MPLS, traffic from a source to a destination may be split arbitrarily over all possible paths. With multiple TMs, we say a set of MPLS routes is feasible if the resulting link data rates are less than or equal to the link capacity for all TMs. From [1], we know that the feasible set of MPLS routes is convex. When the optimization objective is a convex or linear function of routing variables, the routing optimization problem is a linear/convex optimization problem. Specifically, if the objective is to minimize the worst case link utilization, the optimization problem is a LP problem. If the objective is to minimize the worst case network cost, or to minimize the expected network/link cost, the optimization problem is a convex optimization problem. In the latter case, the link cost function is approximated as a piece-wise linear function to expedite the computation. The convex optimization problem is then converted to and solved as a LP problem.

OSPF/IS-IS are two most commonly used intra-domain Internet routing protocols. In OSPF/IS-IS, each link is associated a positive weight. The traffic is routed along the shortest paths. In case of ties where several outgoing links are on the shortest paths to the destination, the flow is split evenly among them. Weight assignments determine packet forwarding, and thus determine the average case and worst case performance. In [6], Bernard *et al.*showed that the routing optimization problem under OSPF/IS-IS is a NP-hard problem. They used local search techniques to iteratively improve the quality of link weights, changing one or a few link weights in each iteration. The procedure ends after 5000 iterations. The solution is not guaranteed to be optimal. Instead, the solution quality is affected by random choice made through iterations.

III. Tradeoff Metric

In this section, we first introduce notation. Then through a simple example, we show the possible significant tradeoff between average case performance and worst case performance. Finally, we introduce our metric to tradeoff between those two considerations.

A. Notation

Network topology: The network $G = (V, E)$ 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 $(k, l) \in E$.

Traffic Matrices : $R = \{R_1, R_2, \ldots, R_n\}$ is a set of n traffic matrices with associated positive weights $\omega = \{\omega_1, \omega_2, \ldots, \omega_n\}$, $\sum_y \omega_y = 1$. 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.

Ratio variables : $B = \{B_{kl}(i, j)\}$, $i, j \in V$, $(k, l) \in E$, where $B_{kl}(i, j) \geq 0$ 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. A set of ratio variables B should satisfy the flow conservation constraints. i.e.,

$$\sum_m B_{mk}(i,j) - \sum_l B_{kl}(i,j) = \begin{cases} 1 & k = j \\ -1 & k = i \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Let ζ denote the set B that satisfies (1). As shown in [1], $B \in \zeta$ can always be implemented by MPLS. Therefore, we use $B \in \zeta$ and a set of MPLS routes interchangeably.

Link weights : $W = \{w_{kl}\}$, where $w_{kl} > 0$ is the weight for link (k, l). In OSPF/IS-IS, traffic is forwarded along the shortest paths determined by w_{kl} from source to destination. In case of ties where several outgoing links are on shortest paths to the destination, the flow is split evenly among them. A set of weights W completely determines the packet routes. OSPF/IS-IS is not as flexible as MPLS. As shown by [6], in general, the set of ratio variable sets implemented by OSPF/IS-IS is a subset of ζ.

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

$$f_{y,kl} = \sum_{i,j} R_y(i,j) B_{kl}(i,j) \tag{2}$$

Link cost : $D = \{D_{kl}\}$, $(k, l) \in E$, 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. In the context of routing optimization, link cost normally represents the overall packet delay on a link [8]. While our analysis can be applied to any function with such properties, we will use,

$$D_{kl}(x) = \frac{x}{c_{kl} - x} \tag{3}$$

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

$$D_{kl}(x) = \max_{1 \leq i \leq 6} (k_i \frac{x}{c_{kl}} + b_i) \tag{4}$$

Network cost : A_y denotes the network cost of TM R_y, $y \in \{1, \ldots, n\}$. It is the sum of the cost of all of the links when the TM is R_y.

$$A_y = \sum_{(k,l) \in E} D_{kl}(f_{y,kl}) \tag{5}$$

B. Average Case and Worst Case Performance under Multiple TMs

Let P^D and P^A denote the expected link cost and expected network cost.

$$P^D = \frac{1}{|E|} \sum_{y\in\{1,\ldots,n\},(k,l)\in E} \omega_y D_{kl}(f_{y,kl}) \quad (6)$$

$$P^A = \sum_{y\in\{1,\ldots,n\}} \omega_y A_y = |E| \bullet P^D \quad (7)$$

Note that P^D and P^A differ only by a ratio of $|E|$.

Let F^D and F^A denote the worst case link cost and worst case network cost.

$$F^D = \max_{y\in\{1,\ldots,n\},(k,l)\in E} D_{kl}(f_{y,kl}) \quad (8)$$

$$F^A = \max_{y\in\{1,\ldots,n\}} A_y \quad (9)$$

C. Tradeoff between Average Case Performance and Worst Case Performance

It is easy to see that the expected network cost is equivalent to the worst case network cost under a single TM. However, in the case of multiple TMs, the two metrics capture different characteristics of a routing configuration, and thus may introduce a tradeoff between each other. We illustrate this tradeoff with a small example. In this example, the worst case performance is high if we only focus on the average case performance. However, we will observe that the worst case performance can be significantly reduced by sacrificing the average case performance by just a very small amount. The example is based on a network G and a $2-$piece link cost function shown in Figure 1.

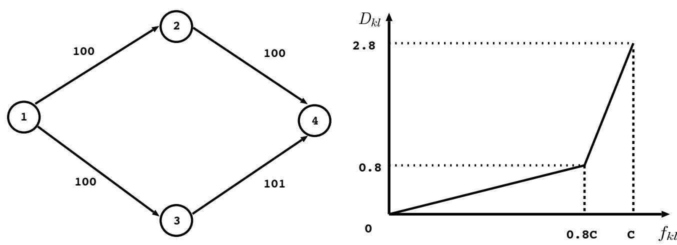

Fig. 1. A topology and a 2-piece link cost function

We consider the case of two TMs $R = \{R_1, R_2\}$ with weights, $w = \{0.5, 0.5\}$.

$$R_1 = \begin{bmatrix} 0 & 0 & 0 & 20 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 80.8 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 0 & 20 \\ 0 & 0 & 0 & 80 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (10)$$

Figure 2 shows the optimization result for average case performance and worst case performance. If we only optimize the average case performance, we incur both high link level and network level worst case performance. The demand between node 3 and node 4 in TM R_1 requires 80% of the capacity of link $(3, 4)$, and the demand between node 2 and node 4 in TM R_2 requires 80% of the capacity of link $(2, 4)$. Beyond 80% link utilization, $\frac{\partial D_{34}}{\partial f_{34}}$ $(= \frac{10}{101})$ is slightly cheaper than $\frac{\partial D_{24}}{\partial f_{24}}$ $(= \frac{10}{100})$. As a result, the average case performance is optimized when all traffic from node 1 to node 4 are forwarded through link $(1, 3)$, and then link $(3, 4)$. i.e., $B_{13}(1, 4) = B_{34}(1, 4) = 1$. However, this optimal solution for average case performance incurs high worst case performance. As all traffic from node 1 to 4 are forwarded through link $(1, 3)$, and then $(3, 4)$, in TM R_1, link $(3, 4)$ is heavily congested, with a high utilization of 99.8%, and a high link cost of 2.78 . Therefore, TM R_1 incurs a high network cost of 2.98 compared to TM R_2 network cost of 1.2.

However, if we can sacrifice the average case performance a little bit, the worst case performance might be substantially improved at both network level and link level. This is achieved by routing the traffic from node 1 to node 4 about evenly via link $(1, 2)$ and link $(1, 3)$. The cost of link $(3, 4)$ under TM R_1 is substantially decreased, from above 2.78 to about 1.8; and the network cost of TM R_1, from above 2.98 to about 2.1. Note that this substantial improvement in worst case performance is achieved by only sacrificing a little bit in average case performance — the expected network cost is slightly increased from above optimal value 2.09 to a value less than 2.10.

D. Our Metric: Tradeoff between Average Case Performance and Worst Case Performance

We have shown the potential significant tradeoff between average case performance and worst case performance. In this paper, we use a simple metric: weighted sum of average case performance and worst case performance to control the tradeoff

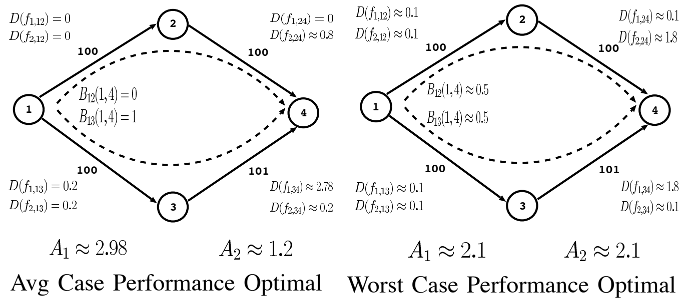

Fig. 2. An example : tradeoff between average case performance and worst case performance

between these two considerations. i.e., given $\alpha \in (0, 1)$,

$$(1-\alpha)P^D + \alpha F^D \tag{11}$$

$$(1-\alpha)P^A + \alpha F^A \tag{12}$$

When α approaches 0, we are optimizing the average case performance; and when α approaches 1, we are optimizing the worst case performance. We control the tradeoff between average case and worst case performance by varying α between 0 to 1. The effectiveness of our metric will be shown in next section.

IV. Tradeoff Metric Effectiveness and Solution

In this section, we first show our metric effectiveness by proving that the average case performance is a convex, decreasing, continuous and differentiable function of worst case performance. Then, we propose solution methods for MPLS and OSPF/IS-IS to solve our metric.

A. Metric Effectiveness

When varying α between 0 to 1, as we shall see, our metric nicely controls the tradeoff between average case performance and worst case performance. We demonstrate this by proving Theorem 4.1.

Theorem 4.1: Under MPLS, assume that the link cost is a continuous, non-decreasing function of link data rate, and $\forall (k,l) \in E$, $\frac{\partial D_{kl}^2}{\partial^2 f_{kl}} > 0$. For the routing optimization problem defined in (11) and (12), the average performance P^A and P^D are non-decreasing continuous function of α; the worst case performance F^A and F^D are non-increasing continuous function of α. The average performance and worst case performance tradeoff curve $\{(P(\alpha), F(\alpha)), \alpha \in (0,1)\}$ is a differentiable decreasing convex curve on P-F plane and

$$\left.\frac{dF(\alpha)}{dP(\alpha)}\right|_{\alpha} = 1 - \frac{1}{\alpha}$$

Proof: See Appendix A. ■

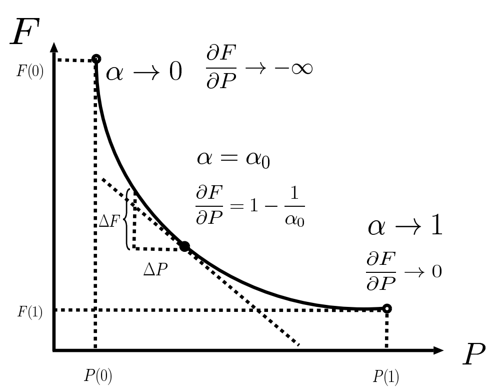

Fig. 3. Average case performance is a convex, decreasing, and differentiable function of worst case performance

Theorem 4.1 reveals the fundamental tradeoff between the average case and the worst case performance: as shown in Figure 3, average case performance is a convex, decreasing, and differentiable function of worst case performance. Our metric, the weighted sum of average case and worst case performance, nicely controls the tradeoff between average case performance and worst case performance. When we prioritize average case performance over worst case performance to the extreme, we end up with the left corner point $(P(0), F(0))$. i.e., a result derived when we first optimize the average case performance, and then optimize the worst case performance given that the average case performance is minimized. As a property of this dual optimization, $\forall\alpha \in (0,1)$, the MPLS solution generated by our metric is always loop-free. Existing work solely focusing on worst case performance may get a solution with loop and bad average case performance. [11] relieves this problem using a loop removal algorithm. However, their solution may still incur a bad average case performance. When we prioritize worst case performance over average case performance to the extreme, we end up with the right corner point $(P(1), F(1))$. In this paper, we use 0.0001 as a value of α to be sufficiently close to 0, and 0.9999 as a value to be sufficiently close to 1.

At the same time, all points on the curve would be of interest given that we intend to tradeoff between those two considerations. For a given α_0, there is a point on the tradeoff curve with average case performance $P(\alpha_0)$ and worst case performance $F(\alpha_0)$. The derivative is $\frac{dF(\alpha)}{dP(\alpha)}|_{\alpha=\alpha_0} = 1 - \frac{1}{\alpha_0}$. The derivative at $(P(0), F(0))$ is dramatic — which indicates that the worst case performance can be substantially reduced at minor cost of average case performance. The derivative at $(P(1), F(1))$ is close to 0 — which indicates that the average case performance can be substantially improved at minor cost of worst case performance. Due to the convexity of the curve, the derivative $\frac{dF(\alpha)}{dP(\alpha)}|_{\alpha=\alpha_0}$ not only describes the local trend of the tradeoff between $F(\alpha)$ and $P(\alpha)$ when the weight is in the neighborhood of α_0, but also serves as a bound for global trend. More specifically, at α_0, in order to reduce the the average case performance from $P(\alpha_0)$ to $P(\alpha_0) - \Delta_p$, $0 \leq \Delta_p \leq P(\alpha_0) - P(0)$, one has to suffer an increase in the worst case performance of $\Delta_F > (\frac{1}{\alpha_0} - 1)\Delta_P$. Symmetrically, to reduce the the worst case performance from $F(\alpha_0)$ to $F(\alpha_0) - \Delta_F$, $0 \leq \Delta_F \leq F(\alpha_0) - F(1)$, one has to suffer an increase in the average case performance of $\Delta_P > \frac{\alpha_0}{1-\alpha_0}\Delta_F$.

B. Solution method for our tradeoff metric

We derive the optimal solution with respect to our tradeoff metric under two class of routing protocols – MPLS and OSPF/IS-IS. Given the tradeoff parameter α, we optimize our tradeoff metric among all feasible sets of routes. (A set of routes is feasible if the resulting link data rates is less than or equal to capacity for all TMs.)

As we mentioned earlier, $B \in \zeta$ and MPLS routes can be used interchangeably. Therefore, we formalize the MPLS routing optimization problem using ratio variables B as control variables.

Given: network $G = (V, E)$, link capacity C, n TMs.
Minimize: $(1-\alpha)P^D + \alpha F^D$, or $(1-\alpha)A^P + \alpha F^A$.
Constraints:
For each TM R_y, $y \in \{1, \ldots, n\}$,
1) Route constraints. F_y is implemented by a set of routes $B \in \zeta$.
2) Feasibility constraints. $\forall y \in \{1, \ldots, n\}, (k,l) \in E, f_{y,kl} \leq c_{kl}$.

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

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

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

From the above problem formulation, we see that our tradeoff metric can be solved as a convex optimization problem. When the link cost function is approximated as piece-wise linear functions, our metric can be solved as a LP problem.

Under OSPF/IS-IS, [6] shows that the optimization problem for average case performance is an NP-hard problem even for a single TM. Our problem is even harder. Therefore, we solve our tradeoff metric under OSPF/IS-IS using a so-called local search heuristic introduced in [6]. In this approach, we iteratively improve the solution quality by adjusting one or a few link weights. Here the quality of a solution is evaluated using our proposed metric defined over the multiple input TMs. Since the problem is NP-hard, the solution is not guaranteed to be optimal. Instead, the solution quality is affected by random choice made through iterations.

V. RESULT

A. Network Topology and TMs

Our evaluations are based on real traffic data collected from a large operational IP network – AT&T's North American commercial backbone network. The network consists of tens of Point of Presence (PoPs), hundreds of routers, thousands of links, and carries over one petabyte of traffic per day.

We use the PoP-level network topology on February 14, 2005 in our evaluations. We first obtain the router-level topology using the methods of Feldmann *et al.* [12]. We then reduce the router-level topology into a PoP-level topology by collapsing the router-level links between the same pair of PoP into a single PoP-level link. The capacity of the PoP-level link is computed as the sum of the capacities of all the underlying router-level links.

The TMs are estimated from SNMP link load measurements using the *tomo-gravity* method [13], which has been shown to yield accurate estimates, especially for large TM elements. We use hourly TMs, as they are commonly used in network engineering applications. The data collection in our study contains two weeks of hourly TMs (from February 13, 2005 to February 26, 2005). The TMs in our original dataset are at the router level. We aggregate them into PoP-level TMs by mapping the demand between each pair of routers to the corresponding pair of PoPs. Each PoP-level TM contains over 400 origin-destination flows at rates ranging from tens of Kbps to tens of Gbps.

As the true value of the traffic demand and link utilization is considered proprietary, we normalize the TMs in our data collection so that the maximum link utilization with Cisco's default routing (OSPF/IS-IS weight of each link being inversely proportional to its capacity) under the first TM in our collection is 1. In the rest of the paper, we will use these normalized TMs for our evaluation.

B. Evaluation Methodology

From the small example in Figure 2, we have observed that there can be a significant tradeoff between average case performance and worst case performance for a given set of demand scenarios. Our evaluation in this section is to explore such a tradeoff with real network demands and configuration. We base our analysis primarily on a set of peak demand scenarios, which we will refer to as the *Peak TMs*. We look at the TMs at the peak hours on weekdays. For each day, we select the six consecutive hours that appear to have the highest total traffic volume. Intuitively, those TMs should capture the traffic variability due to the wide geographic range of the network – the network we considered spans three time zones, and the traffic on the west tends to lag a few hours behind in reaching its peak compared to that on the east. Due to space limit, we only report the result of one day – February 14, 2005. The result of the other days are quantitatively similar. Besides *Peak TMs*, we also consider scenarios with a mixture of *Peak TMs* and *Low TMs*, which represents the demand at hours when the total traffic volume appears to be low. We pick, from each day, the twelve consecutive hours with the lowest total volume as the set of *Low TMs*.

For each set of input TMs, we use equal weight for each TM. We consider the two classes of routing protocols: MPLS and OSPF/IS-IS. We identify the optimal routes under MPLS by constructing a LP system (as described in Section IV-B) and then solve it with the AMPL/CPLEX [14] toolkit. Under OSPF/ISIS, we explore the "best" weight setting by implementing a local search strategy similar to that in [6]. We limit our search to 5000 iterations, which appear to have converged to an optimal (at least locally) solution. Next, we first present our evaluation result for MPLS, and then for OSPF/IS-IS.

C. MPLS

1) Tradeoff grows as load increases: We first explore the extent of the performance tradeoff under a variable level of offered work load. To do so, we take the *Peak TMs*, and scale every demand in each of the TMs by a constant factor. We vary the scaling factor from 0.6 to 1.8. For each scaling factor, we evaluate the performance of two extreme routing configurations. On one extreme, we focus on optimizing the average case performance (by setting $\alpha = 0.0001$). On the other extreme, we focus on optimizing the worst case performance (by setting $\alpha = 0.9999$). In either case, we compute the average case performance cost (P^A or P^D, depending on which of the performance measures being considered), and the worst case performance cost (F^A or F^D).

Figure 4 plots the difference between the two extreme routing optimizations in terms of the average case performance and worst case performance: those when $\alpha = 0.9999$ subtract those when $\alpha = 0.0001$. Figure 6 (a) presents the result of using the network-wide cost as the performance measure and Figure 6 (b) presents the result for the link-level cost. From Figure 6 (a), we observe that optimizing only on the average case performance ($\alpha = 0.0001$) incurs a higher cost for the worst case than the lowest worst-case-cost achievable ($\alpha = 0.9999$), and the difference grows as the offered workload increases – manifested by the increasing curve on the positive side of y-axis. The counterpart is also true: optimizing only on the worst case performance would result in a higher average-case-cost than the optimal, with the difference increasing with load. This is manifested by the decreasing curve (increasing in magnitude) on the negative side of y-axis. This result is expected since when the load is very low, all demand flows can be routed through the least expensive route, which is optimal for each individual TM. In this case, there is no difference between optimizing either the worst case or the average case. However, when the load becomes sufficiently large in comparison to the capacity, contention arises among different TMs – a routing configuration works best for one TM may perform poorly for another TM. Thus, finding a routing configuration that performs well both in average case and in worst case becomes very important.

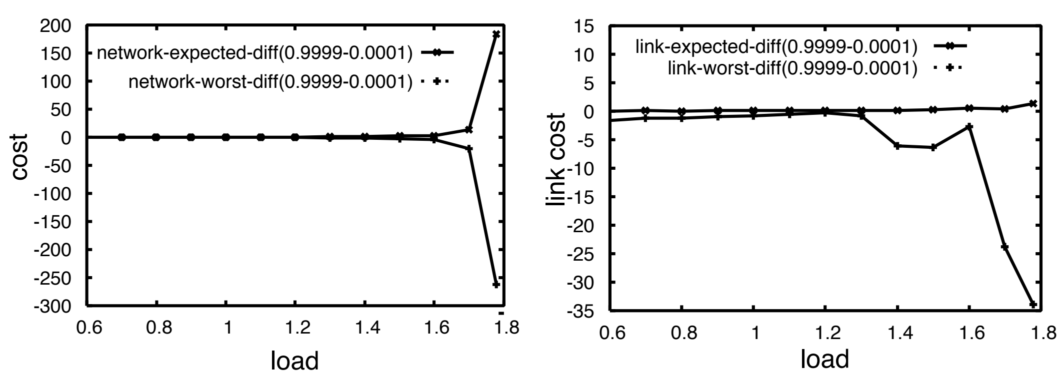

(a) Network-level Performance (b) Link-level Performance

Fig. 4. Tradeoff grows as load increases (*Peak TMs*)

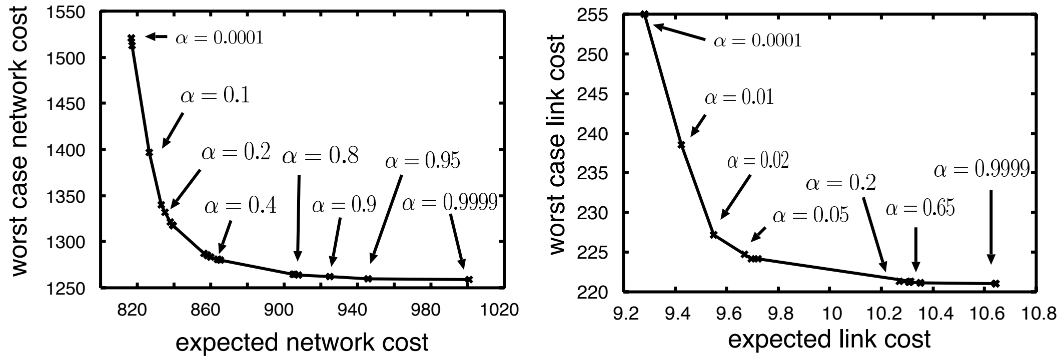

(a) Network-level Performance

(b) Link-level Performance

Fig. 5. Tradeoff between performance and worst case performance

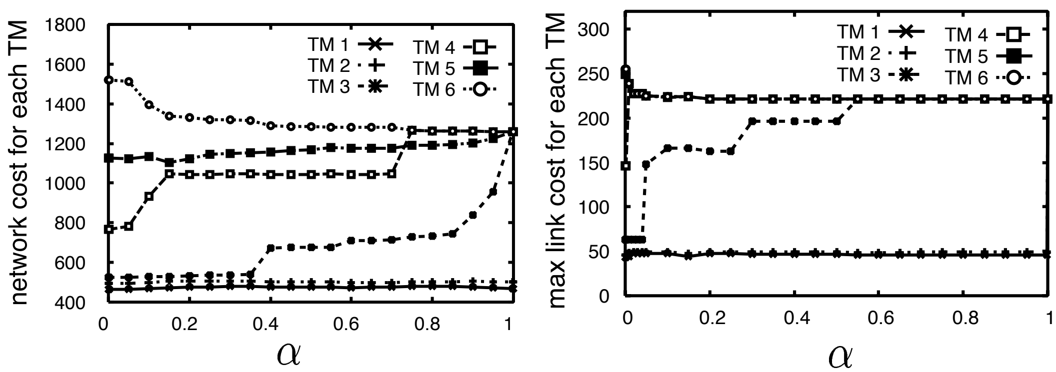

(a) Network-level Performance

(b) Link-level Performance

Fig. 6. Change of each TM in network level cost and link level cost

For link level performance measure (Figure 4 (b)), we find that the general trend matches that in Figure 4 (a). However, the magnitude of the difference is not monotonically increasing as load increases. This is an artifact introduced by the piece-wise linear function that we used to approximate the continuous link cost. Consider the case where a flow can be routed through two paths. In order to optimize the average link cost, it is advantageous to route traffic through the one with higher capacity. However, as load increases, the utilization of the higher-capacity-path reaches a next level on the piece-wise linear function, it may become advantageous to route the additional traffic through the lower-capacity-path. This continues until a next level of piece-wise function is reached. Meanwhile, in order to optimize the worst link cost, it is always advantageous to balance the traffic so that the two paths have equal cost. Now we look again at the process as load increases: at the beginning, the two objectives contradict to each other, which means the difference between the optimization result should increase; at the second phase, however, both optimization lead to a more balanced routing, i.e., their difference should decrease; and this repeats until a next level of piece-wise linear function is hit, producing the non-monotonic behavior of Figure 4 (b).

2) α provides us control on the level of desired tradeoff: Now we have seen that the two optimization objectives may have significant difference. In this subsection, we examine how much control we can achieve by optimizing routing with respect to our metric in the tradeoff between the average case and worst case performance. To do so, we fix the set of offered load (*Peak TMs* with a scaling factor of 1.78) and vary the tradeoff parameter α from 0.0001 to 0.9999.

Figure 5 plots such a tradeoff curve – the worst case performance as a function of the average case performance. As α changes from 0.0001 to 0.9999, we are trading off average case performance with worst case performance. The amount of possible tradeoff is considerable. The range of tradeoff for worst case network cost is 262, about 21% of the optimal worst case network cost. The range of tradeoff for expected network cost is 184, about 23% of the optimal expected network cost. The range of tradeoff for worst case link cost is 34, about 15% of the optimal worst case link cost. The range of tradeoff for expected link cost is 1.4, about 15% of the optimal expected link cost.

We also observe that by changing α, we have very fine-grained control in determining the desirable level of tradeoff. With either network level performance measure or link level performance measure, we can identify solutions with both good average case performance and good worst case performance. In particular, with network level performance measure (Figure 5 (a)), when we choose $\alpha = 0.2$, the worst case network cost drops dramatically, from 1520 to 1332 (15% of the optimal worst case cost), with minor increase of average case cost, from 817 to 835 (2.2% of the optimal average cost). Similarly, with link level performance measure (Figure 5 (b)), when we choose $\alpha = 0.02$, the worst case link cost drops dramatically, from 255 to 227 (13% of the optimal worst case cost), with minor increase of average case cost, from 9.28 to 9.55 (3% of the optimal average case cost).

Figure 6 provides us insight on how is the tradeoff achieved by plotting the performance cost under each individual TM when α changes from 0.0001 to 0.9999. Figure 6 (a) uses the network-wide cost and Figure 6 (b) uses the link-level cost as the performance measure. We can easily observe the contention among the different TMs. In either performance measure, TM 6 dominates the other TMs, in that it represents the worst-cast performance. When we increase α, we observe that the performance cost associated with TM 6 indeed decreases. However, this is achieved by sacrificing the performance under other TMs. In particular, when we uses network level performance measure, we find that TM 3, 4, 5 suffer considerable performance degradation, especially when α gets close to 1. This implies that TM 3, 4, 5 and 6 have competing demand such that their individually preferable routing configurations differ. On the other hand, TM 1 and TM 2 share more conformity with TM 6 in that they appear having contributed little in the contention.

3) Tradeoff in case of peak TMs and low TMs: We have seen that the for *Peak TMs*, the tradeoff can be considerable. Now, we turns to the tradeoff in case of a combination of peak TMs and low TMs. We focus on two TMs with a scaling factor of 1.78, one TM from *Peak TMs*, and one TM from *Low TMs*. Table I shows the result for network level optimization tradeoff. For both expected network cost and worst case network cost, the tradeoff is very limited — no more than 0.1% of the optimal expected network cost or worst case network cost. This is because the peak TM dominants in the expected network cost, as well as the worst case network cost. No matter which one we are focusing on, the routing solution is always close to the optimal solution of the peak TM. As a result, we get very limited network level tradeoff.

α	A^P	A^F
0.0001	483.68	876.59
0.9999	484.01	875.85

TABLE I

LIMITED NETWORK LEVEL TRADEOFF FOR A PEAK TM AND A LOW TM

At link level, Figure 7 (a) shows the result of tradeoff for a peak TM and a low TM. We also plot the link level tradeoff curve for the peak TM as Figure 7 (b). For the worst case link cost, the peak TM determines the range of tradeoff: the tradeoff of two TMs (a) and the tradeoff of the peak TM (b) are equivalent (from 84 to 201). For the expected link cost, the peak TM also determines the range of tradeoff: the range of the tradeoff for two TMs (a) (from 5.5 to 6.24) is equivalent to the 50% of the tradeoff range for the peak TM (from 9.95 to 11.44). The tradeoff curve of two TMs (a) is sharper than the peak TM (b). Therefore, to achieve the same worst case link cost, the corresponding α in (a) is smaller than (b). This α would continue to decrease as we increases the number of low TMs. In the case of a large portion of low TMs, the body of the tradeoff curve corresponds to α with small values. And thus we need to focus on small value α to exploit the benefit of tradeoff for such input TMs.

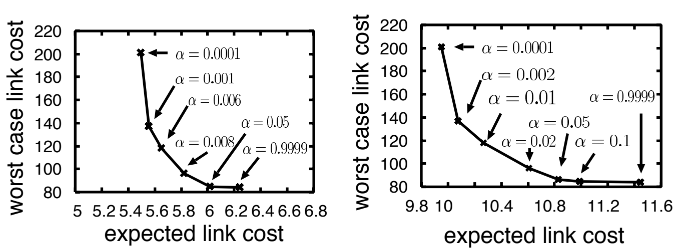

(a) Link level tradeoff for a peak TM and a low TM

(b) Link level tradeoff for peak TM only

Fig. 7. Link cost tradeoff dominated by peak TM, peak TM and low TM

D. OSPF/IS-IS

Under OSPF/IS-IS protocol, we fix the set of offered load (*Peak TMs* with scaling factor 1.59) and examine the tradeoff between average case performance and worst case performance. We solve the tradeoff metric using our implemented heuristic algorithm described in IV-B. Table II shows the optimization result when α takes the value of 0.0001 and 0.9999. The performance difference between optimization result of 0.0001 and 0.9999 is considerable. At link level, the difference in expected link cost D^P is 0.5, about 18% of the best found expected link cost. The difference in worst case link cost D^F is 166.54, about 290% of the best found worst case link cost. At network level, the difference in expected network cost A^P is 44.64, about 14% of the best found expected network cost. The differnce in worst case network cost A^F is 45.54, about 10% of the best found worst case network cost. However, when we vary α from 0.0001 to 0.9999, We could not find more fine-grid tradeoff solution — the solution is always one of the two listed in the table. However, this result might not reflect the existing fine-grid tradeoff solution as OSPF/IS-IS optimizer may not return optimal solution due to the NP-hardness of the problem. On the other hand, OSPF/IS-IS, a coarse grid routing protocol compared to MPLS, may inherently have few fine-grid tradeoff solutions.

α	D^P	D^F	A^P	A^F
0.0001	3.74	223.73	328.85	521.96
0.9999	4.24	57.19	373.49	476.42

TABLE II
TRADEOFF BETWEEN AVERAGE AND WORST CASE PERFORMANCE UNDER OSPF/IS-IS

E. Comparison with average case performance only or worst case performance only approaches

We compare our approach with existing approaches which focus on only average case performance or only worst case performance. i.e. $\alpha = 0$ and $\alpha = 1$. However, completely focusing on one metric and ignoring the other one may lead to catastrophically high cost for the ignored metric. We illustrate this using *Peak TMs*. Table III shows the link level optimization result when α are 0, 0.0001, 0.9999, and 1. When $\alpha = 0$, we are only optimizing the expected link cost. When $\alpha = 0.0001$, we are first optimizing the expected link cost, and then optimizing the worst case link cost given that the expected link cost is minimized. Since $\alpha = 0$ and $\alpha = 0.0001$ both optimize the expected link cost in first priority, we get the same expected link cost D^P. For the same reason, when $\alpha = 1$ and $\alpha = 0.9999$, we get the same worst case link cost D^F. However, when $\alpha = 0.9999$, since we also optimize the expected link cost given that the worst case link cost is minimized. We get a much lower expected link cost 10.64, compared to 83.56, which is achieved when we only optimize the worst case link cost. When $\alpha = 1$, the routing solution may have loop for traffic not been forwarded through the most congested link. Previous work [11] address this problem by applying an additional loop removal algorithm. However, their solution may still incur high expected link cost. In contrast, our algorithm not only guarantee to have loop-free routing solution, but also achieve the optimal expected link cost.

α	D^P	D^F
0	9.28	255.00
0.0001	9.28	255.00
0.9999	10.64	221.03
1	83.56	221.03

TABLE III
BAD EXPECTED LINK COST IF ONLY MINIMIZING WORST CASE LINK COST

VI. REVISIT: OPTIMIZING THE PERFORMANCE VARIANCE

Throughout this paper, we focus on the tradeoff between average case performance and worst case performance. However, motivated by a commonly used robust optimization framework [10], we now investigate another routing optimization approach which aims to reduce the performance sensitivity relative to the change of TMs. More specifically, let D^V and A^V to quantify the link-level and the network-level performance variance.

$$D^V = \frac{1}{|E|} \sum_{y \in \{1,\ldots,n\}, (k,l) \in E} \omega_y (D_{kl}(f_{y,kl}) - D^P)^2 \tag{13}$$

$$A^V = \sum_{y \in \{1,\ldots,n\}} \omega_y (A_y - A^P)^2 \tag{14}$$

However, because the variance-based performance variance measure is not tied to the absolute performance, minimizing variance-based metrics may result in bad performance for all TMs — the performance of each TM is worse than or equal to those of the optimization results focusing on the average case performance (In Appendix B, we illustrate this by a small example). Because of the deficit of this performance variance metric, in this paper, we have chosen to focus on the worst case performance, as well as the tradeoff between it and average case performance.

VII. CONCLUSION

In this paper, we have studied the routing optimization problem in the presence of multiple TMs. Previous work in this field has been focusing either on the average case performance or the worst case performance. Our work, however, has focused on obtaining an understanding of the tradeoff between optimizing the average case performance and optimizing the worst case performance. We have first discovered, through examining both a small constructed example and data from a real operational network, that solely optimizing one of the metrics (average case or worst case performance) can sometimes lead to a relatively poor performance of the other. In particular, we found cases where a $15\% - 23\%$ relative difference exists in their performance cost when optimizing one versus the other.

Since it is often desirable to find a set of routes that results in good average case performance over all TMs considered while avoiding bad worst case performance for any single TM, we have proposed a simple metric, weighted sum of the average case and the worst case performance, to control the tradeoff between the two consideration. Despite of its simple form, this metric is very effective. We have proved that routing optimization over this metric have nice properties such as the average case performance being a decreasing convex and differentiable function to the worst cast performance.

Equipped with this tradeoff metric, we have further developed solution approaches, by extending the methods in [1] and [3], to identifying the routing configuration that represents the desired level of tradeoff under two classes of widely used intra-domain routing protocols – MPLS and OSPF/IS-IS. We have evaluated our approach with data collected from an operational tier-1 ISP. In the case of MPLS, we found that our approach can identify solution that dramatically improves the worst case performance ($13\% - 15\%$) while slightly sacrificing the average case performance ($2.2\% - 3\%$), in comparison to that by optimizing solely on the average case performance. In the case of OSPF/IS-IS, although there exists significant difference in optimizing the average case versus the worst case performance, a fine-grained tradeoff is difficult to achieve due to the limited control that OSPF/IS-IS provide.

VIII. Appendix

A. Convexity of the Trade-Off Curve in Section III-D

We prove Theorem 4.1 in two steps. We first prove Theorem 8.1 and then Theorem 8.2 for MPLS. We assume that the link cost is a continuous, non-decreasing function of link data rate, and $\forall (k,l) \in E$, $\frac{\partial D_{kl}^2}{\partial^2 f_{kl}} > 0$.

Theorem 8.1: For the routing optimization problem defined in (11) and (12), the average performance P^A and P^D are non-decreasing continuous function of α; the worst case performance F^A and F^D are non-increasing continuous function of α.

Proof: We prove the case for network cost P^A and F^A. The case for link cost follows the same procedure.

Notation: let $B^*(\alpha)$ be the optimal routing solution at α and construct $F(\alpha) = F^A(B^*(\alpha))$, $P(\alpha) = P^A(B^*(\alpha))$. Denote the link rate vectors resulted from routing $B^*(\alpha)$ by $f^*(\alpha) = [f_1^*(\alpha), \cdots, f_i^*(\alpha), \cdots, f_n^*(\alpha)]$, where $f_i^*(\alpha)$ is a row vector of link traffic rates for TM i when routing $B^*(\alpha)$ is implemented.

Continuity:

- *$\alpha F(\alpha) + (1-\alpha)P(\alpha)$ is continuous in α.*
  Proof by contradiction: if there is a discontinuity between some α_0 and α_0^+, without lose of generality, assume $\alpha_0 F(\alpha_0^+) + (1-\alpha_0)P(\alpha_0^+) < \alpha_0 F(\alpha_0) + (1-\alpha_0)P(\alpha_0)$, then $B^*(\alpha_0^+)$ is a better solution than $B^*(\alpha_0)$ for the optimization problem at α_0. This contradicts with the definition of $B^*(\alpha_0)$. Similarly, we can rule out the discontinuity between α_0 and α_0^-.
- *Both $F(\alpha)$, $P(\alpha)$ are continuous in α.*
  Proof by contradiction: if there is a discontinuity in $P(\alpha)$ between some α_0 and α_0^+, $P(\alpha_0) \neq P(\alpha_0^+)$. Then we have $B^*(\alpha_0^+) \neq B^*(\alpha_0)$ and $f^*(\alpha_0^+) \neq f^*(\alpha_0)$. Since $\alpha F(\alpha) + (1-\alpha)P(\alpha)$ is continuous in α, then

$$\alpha_0 F(\alpha_0^+) + (1-\alpha_0)P(\alpha_0^+) = \alpha_0 F(\alpha_0) + (1-\alpha_0)P(\alpha_0), \tag{15}$$

We can construct another routing

$$\bar{B} = 0.5B^*(\alpha_0) + 0.5B^*(\alpha_0^+).$$

The corresponding link rate vector is

$$\bar{f} = 0.5f^*(\alpha_0) + 0.5f^*(\alpha_0^+).$$

Since link cost function is a strict convex function of link rate, we will have

$$P^A(\bar{B}) < \frac{P^A(B^*(\alpha_0)) + P^A(B^*(\alpha_0^+))}{2} = \frac{P(\alpha_0) + P(\alpha_0^+)}{2}$$
$$F^A(\bar{B}) < \frac{F^A(B^*(\alpha_0)) + F^A(B^*(\alpha_0^+))}{2} = \frac{F(\alpha_0) + F(\alpha_0^+)}{2},$$

Together with (15), we have

$$\alpha_0 F^A(\bar{B}) + (1-\alpha_0)P^A(\bar{B}) < \alpha_0 F(\alpha_0) + (1-\alpha_0)P(\alpha_0)$$

This contradicts with the definition of $F(\alpha_0)$ and $P(\alpha_0)$.

Monotonicity: Suppose $\alpha_2 > \alpha_1$, then by definition,

$$\alpha_2 F(\alpha_2) + (1-\alpha_2)P(\alpha_2) \leq \alpha_2 F(\alpha_1) + (1-\alpha_2)P(\alpha_1) \tag{16}$$
$$\alpha_1 F(\alpha_1) + (1-\alpha_1)P(\alpha_1) \leq \alpha_1 F(\alpha_2) + (1-\alpha_1)P(\alpha_2) \tag{17}$$

Equivalently,

$$\alpha_2(F(\alpha_2) - F(\alpha_1)) \leq (1-\alpha_2)(P(\alpha_1) - P(\alpha_2)) \tag{18}$$
$$\alpha_1(F(\alpha_2) - F(\alpha_1)) \geq (1-\alpha_1)(P(\alpha_1) - P(\alpha_2)) \tag{19}$$

Therefore,

- if $F(\alpha_2) = F(\alpha_1)$, then $P(\alpha_2) = P(\alpha_1)$
- if $F(\alpha_2) > F(\alpha_1)$, then from (18) we have

$$\frac{F(\alpha_2) - F(\alpha_1)}{P(\alpha_1) - P(\alpha_2)} \leq \frac{1-\alpha_2}{\alpha_2}; \tag{20}$$

From (19) we have

$$\frac{F(\alpha_2)-F(\alpha_1)}{P(\alpha_1)-P(\alpha_2)} \geq \frac{1-\alpha_1}{\alpha_1} \tag{21}$$

Combining (20) and (21), we have $\frac{1-\alpha_1}{\alpha_1} \leq \frac{1-\alpha_2}{\alpha_2}$, which contradicts with $\alpha_2 > \alpha_1$, therefore it is impossible to have $F(\alpha_2) > F(\alpha_1)$.

- if $F(\alpha_2) < F(\alpha_1)$, then $P(\alpha_2) > P(\alpha_1)$ and

$$1-\frac{1}{\alpha_1} \leq \frac{F(\alpha_2)-F(\alpha_1)}{P(\alpha_2)-P(\alpha_1)} \leq 1-\frac{1}{\alpha_2} < 0 \tag{22}$$

Therefore, the average performance is a non-decreasing function of α; the worst case performance is a non-increasing function of α.

■

Theorem 8.2: The average performance and worst case performance tradeoff curve $\{(P(\alpha), F(\alpha)), \alpha \in (0,1)\}$ is a differentiable decreasing convex curve on P-F plane and

$$\left.\frac{dF(\alpha)}{dP(\alpha)}\right|_\alpha = 1-\frac{1}{\alpha}$$

Proof: Look at two arbitrary points on the P-F curve, $(P(\alpha_1), F(\alpha_1))$, $(P(\alpha_2), F(\alpha_2))$ and $P(\alpha_2) > P(\alpha_1)$, then we must have $\alpha_2 > \alpha_1$. Let $(P(\alpha_\beta), F(\alpha_\beta))$ be the point on the P-F curve s.t. $P(\alpha_\beta) = \beta P(\alpha_1) + (1-\beta)P(\alpha_2)$, then we know $\alpha_2 > \alpha_\beta > \alpha_1$. From (22) we will have,

$$\frac{F(\alpha_\beta)-F(\alpha_1)}{P(\alpha_\beta)-P(\alpha_1)} \leq 1-\frac{1}{\alpha_\beta} \leq \frac{F(\alpha_2)-F(\alpha_\beta)}{P(\alpha_2)-P(\alpha_\beta)}$$

Therefore

$$\begin{aligned}&(P(\alpha_2)-P(\alpha_1))F(\alpha_\beta)\\ &\leq (P(\alpha_2)-P(\alpha_\beta))F(\alpha_1) + (P(\alpha_\beta)-P(\alpha_1))F(\alpha_2)\end{aligned}$$

Equivalently,

$$F(\alpha_\beta) \leq \beta F(\alpha_1) + (1-\beta)F(\alpha_2).$$

Therefore, on the P-F tradeoff plane, F is a convex function of P. Furthermore, from (22), since $P(\alpha)$ is continuous in α, then P-F curve is differentiable with $\frac{dF}{dP}|_\alpha = 1-\frac{1}{\alpha} < 0$. Therefore, the average performance and worst case performance tradeoff curve $\{(P(\alpha), F(\alpha)), \alpha \in [0,1]\}$ is a differentiable decreasing convex curve on P-F plane. ■

B. An example to show the deficit of the performance variance metric

In this example, through a simple example, we demonstrate the deficit of the performance variance metric. i.e., under the performance variance measure, the network costs of each TM may be worse than or equal to those of the optimization results given by the average case performance measure — in which performance variance metric provides us no gain over the average performance measure for any TM.

This example is based on a topology shown in the left plot of Figure 1. For simplicity, we use $2-$piece link cost function shown in the right plot of Figure 1. We consider the network cost of two TMs $R = \{R_1, R_2\}$ with equal weights $\omega = \{0.5, 0.5\}$.

$$R_1 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 96 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad R_2 = \begin{bmatrix} 0 & 0 & 0 & 84 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{23}$$

Under TM R_1, the network cost, as well as the the cost of link $(3,4)$, is fixed (2.4). The optimal network cost of TM R_2, is about 1.67. This is achieved when the demand from node 1 to node 4 is evenly split and forwarded through link $(1,2)$ and link $(1,3)$. However, if we minimize A^V, the "variance-based network-level worst case performance", the network cost of R_2 must increase in order to reduce the network cost difference from R_1. We end up with a set of routes which forwards all packets from node 1 to node 4 via link $(1,2)$, resulting $A_1 = A_2 = 2.4$, and then $A^V = 0$.

References

[1] Chun Zhang, Yong Liu, Weibo Gong, Jim Kurose, Robert Moll, and Don Towsley, "On routing optimization with multiple traffic matrices," in *IEEE Infocom*, 2005.

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

[3] 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.

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

[5] R.Callon/IETF, "Use of osi is-is for routing in tcp/ip and dual environments(RFC 1195)," 1990.

[6] Bernard Fortz and Mikkel Thorup, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM (2)*, 2000, pp. 519–528.

[7] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus/IETF, "Requirements for traffic engineering over mpls (RFC 2328)," 1999.

[8] 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.

[9] F. Lin and J. Wang, "Minimax open shortest path first routing algorithms in networks supporting the smds services," in *Proceedings of IEEE International Conference on Commmunications*, 1993, vol. 2, pp. 666–670.

[10] John M. Mulvey, Robert J. Vanderbei, and Stavros A. Zenios, "Robust optimization of large-scale system," in *Operations Research*, 1995, pp. 264–281.

[11] Matthw Roughan, Mikkel Thorup, and Yin Zhang, "Traffic engineering with estimated traffic matrices," in *Proceedings of Internet Measurement Conference*, 2005.

[12] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: Methodology and experience," *IEEE/ACM Transactions on Networking*, vol. 9, no. 3, pp. 265–279, 2001.

[13] Yin Zhang, Matthew Roughan, Carsten Lund, and David Donoho, "An information-theoretic approach to traffic matrix estimation," in *Proceedings of ACM SIGCOMM*, 2003.

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