

Dual-Based Link Weight Determination Towards Single Shortest Path Solutions for OSPF Networks

Shekhar Srivastava, Gaurav Agrawal, Deep Medhi *

University of Missouri–Kansas City, Kansas City, MO 64110 USA

Abstract: An important traffic engineering problem for OSPF networks is the determination of optimal link weight system for a given objective function such as the Fortz-Thorup (FT) piece-wise linear convex function. In this paper, we show that the optimal link weights are obtained from the slopes of the FT function when approached from the dual space. This also suggests that the slopes can be adjusted to obtain better weights, for example to reduce number of equal-cost multi-paths (ECMP), i.e., to achieve single shortest paths for different demand pairs, where possible, without unduely increasing the cost. We present an iterative objective function approach where we modify the FT function with a convex (still piece-wise linear) envelope that generates additional slope values. We present results for experimental and random networks. The results show that the iterative heuristic is fairly powerful and robust and often gives unique shortest paths for most demands in the network.

1 Introduction


Traffic engineering of IP OSPF/IS-IS networks has received much attention in the past several years (for example, see [2,3,5,7]). The heart of the problem is in determining an optimal link weight system for shortest-path first routing of traffic in an OSPF network while minimizing a given objective function. Our work is primarily motivated by the following aspects:

- The choice of the objective function: while many different objectives have been considered for the IP traffic engineering problem, the one that has received considerable attention is a piece-wise linear convex envelope of a non-linear link cost function induced by queueing delay, as proposed by Fortz and Thorup [2,3]; for brevity, we refer to this function as the FT function which is shown below (see also Figure 1(a)) for flow amount y and capacity c :


$$\phi(y; c) = \begin{cases} y & \text{for } 0 \leq \frac{y}{c} < \frac{1}{3} \\ 3y - \frac{2}{3}c & \text{for } \frac{1}{3} \leq \frac{y}{c} < \frac{2}{3} \\ 10y - \frac{16}{3}c & \text{for } \frac{2}{3} \leq \frac{y}{c} < \frac{9}{10} \\ 70y - \frac{178}{3}c & \text{for } \frac{9}{10} \leq \frac{y}{c} < 1 \\ 500y - \frac{1468}{3}c & \text{for } 1 \leq \frac{y}{c} < \frac{11}{10} \\ 5000y - \frac{16318}{3}c & \text{for } \frac{11}{10} \leq \frac{y}{c} < \infty. \end{cases} \quad (1)$$

- Multi-commodity flow and duality: for a linear multi-commodity flow problem, the dual solution associated with the “capacity” constraints of the multi-commodity flow problem is related to the optimal link weight for shortest path routing. This result was described in [1, § 17.5] for a linear multi-commodity flow problem, and has recently gained attention in the context of OSPF routing [5,7].

* Corresponding Author: dmedhi@umkc.edu



(a)



(b)

Fig. 1. (a) FT and M/M/1 link-cost function (for $c = 1$); (b) Splitting process of the cost function

The connection between multi-commodity flow and duality has not been done when the FT function is used as the objective function. In this paper, we prove that the use of FT function as the objective has an interesting relation in terms of the dual solution and the optimal link metric system; specifically, the optimal link weight takes the value from the *slopes* of the FT function.

The above result has taken us to look at the FT function closely. It may be noted that the slopes of the FT function are increasingly steep from 1 to 5000. Through preliminary numerical studies, we found that the FT function pushes flows so that the traffic is fairly load balanced in the network which, in turn, leads to picking link weight from the dual solution which can get “stuck” in one of the slope values. As for example, by inspecting (1), we can see that if loads in the network are evenly distributed, say, at around 40% utilization, for all links, then the slope value will be 3 for all links. Having all links have the same link weight value is equivalent to saying that the network operates as if in shortest-hop routing mode. Another related observation we have noticed is that, due to specific values of slopes of the FT function and load balancing, the network tends to have several shortest paths for the same origin/destination node pairs. In the context of OSPF routing, this leads to the use of equal-cost multi-path (ECMP) feature where the traffic is supposedly evenly distributed among the shortest paths. Recently, Thorup and Roughan [6] have reported that in practice ECMP does not exactly lead to equal load distribution among the shortest paths; furthermore, from a network troubleshooting point of view, it is often desirable to have single shortest paths, if possible.

In light of the above, we have taken a fresh look at the FT function to see if we can do anything to avoid multiple shortest paths, where possible, since the optimal link weight is obtained from the slopes of the FT function. This is where we take the liberty to explore on a statement in [2] that says that it is not necessary to use the FT function in its proposed form, but a different piece-wise linear increasing and convex function, presumably in the same spirit, can be used. Thus, we have developed a successive modification of the FT function that remains as a piece-wise linear convex envelope of the FT function; the approach can be considered as a successive LP-based approach since the objective function is changed to determine additional slope values that can then lead to new link weight values (due to our duality result), hopefully leading to as many pairs with unique shortest paths as possible. Through our computational work, we have found that this is indeed possible in many cases, thus providing a promising way to compute link weight while keep congestion low.

2 Traffic Engineering Preliminaries

The traffic engineering problem for the OSPF network consists in determining a link weight system for a given objective. We will denote the link weight system by the vector, $\mathbf{w} = \{w_\ell, \ell \in \mathcal{L}\}$; here, w_ℓ is the link weight of link ℓ , and \mathcal{L} denotes the set of links. The space of the link weight system is defined as \mathcal{W} , where $\mathcal{W} = \{\mathbf{w} | w_\ell \in \{0, 1, 2, \dots, \bar{w}\}\}$. Here \bar{w} is an upper bound on the allowable value of the link weight.

Now consider the demand volume (traffic) h_d for demand identifier d for a pair of nodes, $d \in \mathcal{D}$ (where \mathcal{D} is the set of demand pairs with positive demand volumes). We will denote the set of possible paths for demand d by \mathcal{P}_d , the flow induced by the link weight system \mathbf{w} on a path $p \in \mathcal{P}_d$ by $x_{dp}(\mathbf{w})$, the link flow on link ℓ by y_ℓ , and the capacity by c_ℓ . If we now use δ_{dp}^ℓ as the indicator that takes the value 1 if path p for demand d uses link ℓ (0, otherwise), we can also write the link flow on link ℓ as $\sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{P}_d} \delta_{dp}^\ell x_{dp}(\mathbf{w})$. Given this notation, the traffic engineering problem, denoted by \mathbf{P} (or sometimes as $\mathbf{P}(\phi)$), that determines \mathbf{w} while minimizing the FT function can be written as the following general representation (see [4,5]):

$$\min_{\mathbf{w}} \sum_{\ell \in \mathcal{L}} \phi(y_\ell; c_\ell) \quad (2a)$$

subject to

$$\sum_{p \in \mathcal{P}_d} x_{dp}(\mathbf{w}) = h_d, \quad d \in \mathcal{D} \quad (2b)$$

$$\sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{P}_d} \delta_{dp}^\ell x_{dp}(\mathbf{w}) = y_\ell, \quad \ell \in \mathcal{L} \quad (2c)$$

$$x_{dp}(\mathbf{w}) \geq 0, \quad d \in \mathcal{D}, p \in \mathcal{P}_d \quad (2d)$$

$$\mathbf{w} \in \mathcal{W}. \quad (2e)$$

3 Duality, Link Weight, and the FT function

Consider again (2); note that it is not a mathematical programming problem due to implicit functional dependency of flow on a path based on the link weight (i.e., $x_{dp}(\mathbf{w})$). Now if we ignore the dependency on link weight \mathbf{w} , (i.e., x_{dp} instead of $x_{dp}(\mathbf{w})$), and take note of the fact that $\phi(y_\ell; c_\ell)$ is a piece-wise linear increasing convex function which can be replaced with a set of inequalities through the introduction of a set of variables, z_ℓ , we then arrive at the following linear programming problem (see [2]):

$$\min_{\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}} \sum_{\ell \in \mathcal{L}} z_\ell \quad (3a)$$

subject to

$$\sum_{p \in \mathcal{P}_d} x_{dp} = h_d, \quad d \in \mathcal{D} \quad (3b)$$

$$\sum_{d \in \mathcal{D}} \sum_{p \in \mathcal{P}_d} \delta_{dp}^\ell x_{dp} = y_\ell, \quad \ell \in \mathcal{L} \quad (3c)$$

$$z_\ell \geq a_i y_\ell - b_i c_\ell, \quad i = 1, 2, \dots, I, \quad \ell \in \mathcal{L} \quad (3d)$$

$$\mathbf{x}, \mathbf{y}, \mathbf{z} \geq \mathbf{0}. \quad (3e)$$

Here, index i indicates the break points with $a_1 = 1, b_1 = 0; a_2 = 3, b_2 = -\frac{2}{3}$, and so on, from the FT function and I is the total number of break points. Note that a_i satisfies the property $a_{i+1} > a_i$, for $i = 1, 2, \dots, I - 1$.

It is easy to see that problem (3) remains the same if 1) equality in (3b) is replaced by the greater-than-equal-to requirement, and 2) equality in (3c) is replaced by a less-than-equal-to requirement. This implies that dual multipliers associated with constraints (3b) and (3c) are really non-negative (rather than unrestricted in sign). We associate dual multipliers $\lambda = \{\lambda_d, d \in \mathcal{D}\}$, $\pi = \{\pi_\ell, \ell \in \mathcal{L}\}$, $\gamma = \{\gamma_{i\ell}, \ell \in \mathcal{L}, i = 1, 2, \dots, I\}$ to constraints (3b), (3c), (3d), respectively. We can write the dual of (3) as:

$$\max_{\{\lambda, \pi, \gamma\}} \sum_{d \in \mathcal{D}} h_d \lambda_d - \sum_{\ell \in \mathcal{L}} \sum_{i=1}^I b_i c_\ell \gamma_{i\ell} \quad (4a)$$

subject to

$$\lambda_d \leq \sum_{\ell \in \mathcal{L}} \delta_{dp}^\ell \pi_\ell, \quad d \in \mathcal{D}, p \in \mathcal{P}_d \quad (4b)$$

$$\sum_{i=1}^I a_i \gamma_{i\ell} \geq \pi_\ell, \quad \ell \in \mathcal{L} \quad (4c)$$

$$\sum_{i=1}^I \gamma_{i\ell} \leq 1, \quad \ell \in \mathcal{L} \quad (4d)$$

$$\lambda, \pi, \gamma \geq 0. \quad (4e)$$

For brevity, the optimal solution of Problem (3) will be denoted by x^*, y^*, z^* while the optimal solution of dual (4) by $\lambda^*, \pi^*, \gamma^*$.

For this set of primal-dual problems, the optimal solutions, in addition to satisfying the set of primal and dual constraints, need to satisfy the following set of complementary slackness conditions:

$$x_{dp}^* (-\lambda_d^* + \sum_{\ell \in \mathcal{L}} \delta_{dp}^\ell \pi_\ell^*) = 0, \quad d \in \mathcal{D}, p \in \mathcal{P}_d \quad (5a)$$

$$y_\ell^* (-\pi_\ell^* + \sum_{i=1}^I a_i \gamma_{i\ell}^*) = 0, \quad \ell \in \mathcal{L} \quad (5b)$$

$$z_\ell^* (1 - \sum_{i=1}^I \gamma_{i\ell}^*) = 0, \quad \ell \in \mathcal{L} \quad (5c)$$

$$\gamma_{i\ell}^* (z_\ell^* - a_i y_\ell^* + b_i c_\ell) = 0, \quad i = 1, 2, \dots, I, \ell \in \mathcal{L}. \quad (5d)$$

An important relation between the dual solution and finding shortest paths, apparently first reported in [1, § 17.5] (see also, recent work in the context of OSPF routing, [5], [7]) is first stated below; the proof that follows [1] is included for completeness.

THEOREM 1 λ_d^* is the shortest path distance for demand d with respect to link weight $w_\ell^* = \pi_\ell^*$, $\ell \in \mathcal{L}$, and in the optimal solution, every path that carries a positive flow must be a shortest path with respect to this link cost.

Proof: From constraint (4c), we know that the path cost for demand d is at least $\sum_{\ell \in \mathcal{L}} \delta_{dp}^\ell \pi_\ell^*$ for each path p .

If flow on a path p for demand d is positive, i.e., $x_{dp}^* > 0$, then from complementary slackness condition (5a), we have for the same (d, p) pair,

$$\lambda_d^* = \sum_{\ell \in \mathcal{L}} \delta_{dp}^\ell \pi_\ell^*.$$

Thus, multipliers π_ℓ^* can be interpreted as link weights for determining the path cost λ_d for demand d over the links ℓ used by path p . If there are multiple paths with positive flow x_{dp}^* for the same demand d , then each path results in the shortest path cost λ_d^* with respect to the appropriate link cost π_ℓ^* for the links used by the paths. ■

If multiple paths for the same demand pair have positive flows, they are identifiable as shortest paths. On the other hand, the flow on each shortest path can be of different amount in the case of the linear program (3).

Before we use the above result to arrive at our main results, we first need to prove a Lemma in regard to constraint binding for Problem (3).

LEMMA 1 *At optimality, constraint (3d) is binding at at least one i ($i = 1, 2, \dots, I$) for each link $\ell \in \mathcal{L}$.*

Proof: We prove this Lemma by contradiction. Suppose that constraint (3d) in Problem (3) is not binding at any of the indices i for each link $\ell \in \mathcal{L}$. Due to complementary slackness condition (5d), this implies that $\gamma_{i\ell}^* = 0$, $i = 1, 2, \dots, I$, $\ell \in \mathcal{L}$. In turn, this would mean that constraint (4d) is non-binding for each $\ell \in \mathcal{L}$. Now using complementary slackness condition (5c), we find that $z_\ell^* = 0$, for each $\ell \in \mathcal{L}$. Consequently, this will force y_ℓ^* for all $\ell \in \mathcal{L}$ to be zero which is not possible for demand volume, $h_d > 0$, $d \in \mathcal{D}$. Thus, constraint (3d) is binding at at least one i for each $\ell \in \mathcal{L}$. ■

For the rest of the discussion, we assume that the network does not have any link that has zero link load but non-zero capacity; this is a reasonable assumption in real-life networks. In our first main result, we start with assuming the uniqueness of the binding on a particular i for each link $\ell \in \mathcal{L}$.

THEOREM 2 *For each link $\ell \in \mathcal{L}$, assume that constraint (3d) is binding for a unique i (denote by $i'(\ell)$). Then an optimal link metric system for the OSPF routing problem is given by:*

$$w_\ell^* = \pi_\ell^* = a_{i'(\ell)}, \quad \ell \in \mathcal{L}. \quad (6)$$

Proof: From Lemma 1, we know that at least one constraint in (3d) is binding for a specific i for each link $\ell \in \mathcal{L}$. Thus, at optimality, we have

$$z_\ell^* = a_{i'(\ell)} y_\ell^* - b_{i'(\ell)} c_\ell, \quad \ell \in \mathcal{L}. \quad (7)$$

Since $i'(\ell)$ is unique, then for all other i for each ℓ , we have

$$z_\ell^* > a_i y_\ell^* - b_i c_\ell, \quad i \neq i'(\ell), \quad \ell \in \mathcal{L}.$$

By complementary slackness condition (5d), this implies that

$$\gamma_{i\ell} = 0, \quad i \neq i'(\ell), \quad \ell \in \mathcal{L}. \quad (8)$$

In turn, this observation reduces (4d) to

$$\gamma_{i'(\ell), \ell}^* \leq 1, \quad \ell \in \mathcal{L}.$$

If we assume that strict inequality holds here, i.e., $\gamma_{i'(\ell), \ell}^* < 1$, then this would imply (due to complementary slackness condition (5c)) that $z_\ell^* = 0$ for each $\ell \in \mathcal{L}$ which is not possible. Thus, we must

have

$$\gamma_{i'(\ell),\ell}^* = 1, \quad \ell \in \mathcal{L}. \quad (9)$$

Using optimal γ^* as given by (8) and (9), constraint (4c) can be reduced to

$$\pi_\ell^* \leq a_{i'(\ell)}, \quad \ell \in \mathcal{L}.$$

Now suppose that strict inequality holds here, i.e., $\pi_\ell^* < a_{i'(\ell)}$. This time, using complementary slackness condition (5b), we find that $y_\ell^* = 0$ for each $\ell \in \mathcal{L}$ which is not possible for any positive demand volume h_d in the network. Thus, we arrive at

$$\pi_\ell^* = a_{i'(\ell)}, \quad \ell \in \mathcal{L}. \quad (10)$$

Since the optimal link weight is $w_\ell^* = \pi_\ell^*$ from Theorem 1, combining with (10), we have arrived at our claim. \blacksquare

Finally, we present the more general result when uniqueness is relaxed.

THEOREM 3 *For each link $\ell \in \mathcal{L}$, constraint (3d) can be binding at at most two consecutive i 's (denote by $i'(\ell)$ and $i'(\ell) + 1$). Then the optimal link-weight system for the OSPF routing problem is given by:*

$$w_\ell^* = \pi_\ell^* = a_{i'(\ell)}\gamma_{i'(\ell),\ell}^* + a_{i'(\ell)+1}\gamma_{i'(\ell)+1,\ell}^*, \quad \ell \in \mathcal{L}, \quad (11)$$

where $\gamma_{i'(\ell),\ell}^* + \gamma_{i'(\ell)+1,\ell}^* = 1$, $\gamma_{i'(\ell),\ell}^*, \gamma_{i'(\ell)+1,\ell}^* \geq 0$.

Proof: From Lemma 1, we know that (3d) is binding at at least one i . Due to the properties of a_i 's, it is trivial to see that the constraint can be binding in at most two consecutive i 's. Since $y_\ell^* > 0$, we then have (11) due to complementary slackness condition (5b). Similarly, since $z_\ell^* > 0$, we have the condition $\gamma_{i'(\ell),\ell}^* + \gamma_{i'(\ell)+1,\ell}^* = 1$. \blacksquare

4 Iterative Heuristic

Our main results presented in Theorems 2 and 3 provide the impetus for the heuristic proposed below. Note from Theorem 3 that the link weight at optimality can be a convex combination of two adjacent a_i 's; this is possible when the solution is binding on the intersection of two piece-wise linear part of the FT function. However, when a typical linear programming solver that implements simplex method is used for solving such a problem, only one of the slopes will show up in the dual solution (in fact, it will be the smaller value of $a_{i'(\ell)}$). On the other hand, a proper convex combination of two adjacent slope values can conceivable give a better link weight, especially if our goal is to minimize number of ECMP paths for different demand pairs.

4.1 Basis of the Heuristic

Consider the following two observations. Firstly, the value of π (weight of links) can be altered by adjusting the slopes (gradients) of the segments of the link cost function ϕ . Secondly, to achieve unique shortest paths for each demand, where possible, it would suffice to ensure that few chosen links have different weights.

Consider the set of shortest paths found in the solution of $(\mathbf{LP}(\phi))$ problem. For ease of tracking, we will write $\phi_0 = \phi$ to indicate the (starting) FT function (with $i = 0$). Now, if we carefully choose two links such that if these links have different weights then there will possibly be fewer number of multiple shortest paths in the final solution. A possible criterion could be to have links which are on

two different shortest paths for a demand. The difference in weights can be enforced by constructing a new link cost function (ϕ_{i+1}) such that the load of two links corresponds to different active segments. Now, if we solve ($\mathbf{LP}(\phi_{i+1})$) and assuming that the allocation of flows has not changed drastically, the solution would have tendency to have lesser number of demands with multiple shortest paths. Following such steps in a recursive fashion could lead us to solutions with unique shortest paths.

Next, we formally present the algorithm in **Algorithm 1** (to be referred to as *MLP*).

Algorithm 1 *MLP: Iterative Heuristic*

- 1: $i = 0, \phi_0 = \phi$
 - 2: Solve Problem ($\mathbf{LP}(\phi_i)$) based on the cost function ϕ_i
 - 3: Derive dual solution π and utilization u
 - 4: Check if all demands have unique shortest paths, then STOP, else continue
 - 5: Identify links m and n using **Algorithm 2**; if no such links available, stop.
 - 6: Compute ϕ_{i+1} , see subsection 4.2
 - 7: $i \leftarrow i + 1$, and i is less than the maximum number of iterations, Go to Step 2; else, stop.
-

Note that Algorithm MLP can also stop while some demands have multiple shortest paths. This can happen because it could not find two links m and n (**Algorithm 2** returns -1). This can also happen when we can not find the point mn such that the three restrictions (discussed later) are honored.

We next present the implementation of the procedure used for choosing links m and n in **Algorithm 2**. Let \mathcal{P}_d^s be the set of shortest paths based on the current solution of the ($\mathbf{LP}(\phi_i)$) problem. We find the first demand d which has multiple shortest paths. If the demand has more than two paths,

Algorithm 2 *Link Identifying Procedure*

```

while ( $d \in \mathcal{D}$ ) do
  if ( $\mathcal{P}_d^s > 1$ ) then
     $Q_1$  and  $Q_2$  be two min-hop paths in  $\mathcal{P}_d^s$ 
    define set  $\mathcal{S}_1 = Q_1 \setminus \{Q_1 \cap Q_2\}$ 
    define set  $\mathcal{S}_2 = Q_2 \setminus \{Q_1 \cap Q_2\}$ 
    define set  $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$ 
    while ( $n \in \mathcal{S}$ ) do
      if ( $n \in \mathcal{S}_1$ ) then
        find  $m \in \mathcal{S}_2$  such that  $\pi_n = \pi_m$ 
        and  $u_n \neq u_m$ .
        return ( $m, n$ )
      else
        find  $m \in \mathcal{S}_1$  such that  $\pi_n = \pi_m$ 
        and  $u_n \neq u_m$ .
        return ( $m, n$ )
  return -1

```

we pick the paths with minimum number of hops, say Q_1 and Q_2 . For these two paths, we define the set \mathcal{S}_1 as the set of links present in Q_1 but are not common to both Q_1 and Q_2 . Similarly, we construct \mathcal{S}_2 . We next find a link pair $m \in \mathcal{S}_1$ and $n \in \mathcal{S}_2$ (or viceversa), such that dual solution, $\pi_m = \pi_n$ and utilization, $u_m \neq u_n$.

4.2 Computing ϕ_{i+1}

We construct the new link cost function (ϕ_{i+1}) based on the current cost function (ϕ_i) and value of u_m and u_n . We explain the technique using Figure 1(b). Observe that the two links have utilization corresponding to the same segment (since $\pi_m = \pi_n$) but have different utilizations. Without loss of generality, we assume that $u_m < u_n$. The shaded area shows the acceptable region (except the border points) from which point mn can be picked and segment $(u_1, \xi_1)-(u_2, \xi_2)$ in ϕ_i is replaced by two segments, $(u_1, \xi_1)-(u_{mn}, \xi_{mn})$ and $(u_{mn}, \xi_{mn})-(u_2, \xi_2)$ in ϕ_{i+1} . Other segments of ϕ_i are copied as such to ϕ_{i+1} .

The shaded area is constructed based on the following three restrictions on point mn :

- (1) u_{mn} should be bounded on both sides by u_m and u_n .
- (2) gradient of segment $(u_0, \xi_0)-(u_1, \xi_1)$ should be strictly less than the gradient of segment $(u_1, \xi_1)-(u_{mn}, \xi_{mn})$, which is strictly less than the gradient of segment $(u_1, \xi_1)-(u_2, \xi_2)$.
- (3) gradient of segment $(u_1, \xi_1)-(u_2, \xi_2)$ should be strictly less than the gradient of segment $(u_{mn}, \xi_{mn})-(u_2, \xi_2)$, which is strictly less than the gradient of segment $(u_2, \xi_2)-(u_3, \xi_3)$.

First condition is to ensure that the two links m and n have utilizations corresponding to different segments in cost function ϕ_{i+1} . Second and third conditions ensure that the link cost function is still convex in nature. Observe that while constructing ϕ_{i+1} from ϕ_i , the link cost function is only minimally altered. The points links $(u_0, \xi_0), (u_1, \xi_1), \dots$ of ϕ_i are not changed, rather new points like (u_{mn}, ξ_{mn}) are added to ϕ_{i+1} without disturbing the fundamental nature of the link cost function (we discuss more on this in results section).

We would also like to point out that the approach used to choose the links and the procedure of splitting the cost function are not the only ways, rather it is one of the ways. Other modifications and alterations are being investigated. But, they do not change the overall behavior of the heuristic.

5 Numerical Results

We have implemented our approach using C^{++} and CPLEX callable libraries. In our study, we have considered several performance metrics:

- ▷ Fraction of Demands (FD) with multiple shortest paths.
- ▷ Maximum Link Utilization (ML).
- ▷ Network-wide Fortz-Thorup (NFT) [2] metric: it captures the total scaled cost incurred by current allocation. The scaled (normalized) cost $(\sum_{\ell \in \mathcal{L}} \phi(y_\ell; c_\ell) / \varphi)$ is the ratio of total cost of current allocation $(\sum_{\ell \in \mathcal{L}} \phi(y_\ell; c_\ell))$ for the given capacitated network as compared to the cost in case the network was uncapacitated (φ). Observe that for an uncapacitated network with convex link cost function, cost is minimal when flows are allocated to hop count based shortest paths.
- ▷ (Δf) : the difference in the value of objective function (f) for ECMP flow allocation and optimal routing as the percentage of the value of objective function for optimal routing. Δf captures the suboptimality introduced by incorporating ECMP principle and integer weights in a solution derived by LP solution.

5.1 Experimental Networks (EN)

For our first set of studies, we have considered four experimental networks as shown in Figures 2-5. For our study, we have assumed that there are traffic demand volume between all pairs of nodes; we consider both uniform and random demand volume. For each demand, we find the shortest paths based on hop count and then assign all the demand volume on one of the shortest paths. Upon such allocation, we find the flow on each link and determine the capacity of each link so that all the links are at 40% utilization. We use this utilization value since to our knowledge, most real-world networks try to maintain utilization around a similar value.




Fig. 2. EN-I




Fig. 3. EN-II




Fig. 4. EN-III




Fig. 5. EN-IV




Fig. 6. FD and Δf value vs. iterations of MLP for ENs




Fig. 7. Functions ϕ and ϕ_f for ENs

First, we analyze the behavior of **MLP** and study the convergence properties of the heuristic; see Figure 6 where we plot the value of FD and Δf with increasing number of iterations for experimental networks. The results depict two important points. The first one is that only a few iterations are required; for instance, minimum FD values were reached at 4, 3, 6 and 3 iterations for EN-I, EN-II, EN-III and EN-IV, respectively. The second point is that decreasing value of FD leads to smaller value of Δf .

Next, we compare the constructed cost function ϕ_f from the **MLP** approach with the initial cost function ϕ for the experimental networks. For the assumed demands and initial capacities, only the second segment of the cost function is modified and all the other segments of ϕ remain unchanged. Hence, we only show second segment in Figure 7. Observe that for all the experimental networks, **MLP** makes a very minor modification in the cost function ϕ . Moreover, the modified cost function ϕ_f forms a lower envelope of the cost function ϕ and yet, piece-wise linear and convex in nature.

Finally, we present results for ENs in Table 1 for **LP**(ϕ) and in Table 2 for **MLP**. Since OSPF requires integer weights, we have also considered scaling the weight for **MLP** by multiplying by 1000 and rounding the result to the closest integer; note that if the link utilization is less than 90% (see (1)), then weight values stay within 10,000; such values are well within the limits of weights allowed by OSPF. We refer to the newly computed integer weight system derived from **LP**(ϕ) as $w_{LP(\phi)}$ and that from **MLP** as w_{MLP} , and corresponding allocation as **EA**($w_{LP(\phi)}$) and **EA**(w_{MLP}), respectively. We can see from Table 1 and Table 2 our approach obtains single shortest paths for almost all pairs without increasing maximum link utilization and significantly affecting Δf .

We have also considered random networks with 10 and 20 nodes by varying the number of links in the networks. Results are reported in Table 3 and Table 4. We can see that our approach is very effective in attaining single shortest paths for most pairs without drastically affecting Δf .

Table 1
Results for $\mathbf{LP}(\phi)$ for Experimental Networks

ENs	FD	$\mathbf{LP}(\phi)$		$\mathbf{EA}(\mathbf{w}_{LP(\phi)})$		Δf
		ML	NFT	ML	NFT	
EN I	0.33	0.47	1.33	0.72	1.45	0.094
EN II	0.20	0.57	1.33	0.77	1.62	0.217
EN III	0.30	0.67	1.33	1.12	-	-
EN IV	0.38	0.67	1.33	1.41	-	-

Table 2
Results for \mathbf{MLP} for Experimental Networks

ENs	FD	\mathbf{MLP}		$\mathbf{EA}(\mathbf{w}_{MLP})$		Δf
		ML	NFT	ML	NFT	
EN I	0.09	0.40	1.33	0.48	1.33	0.006
EN II	0.00	0.40	1.33	0.40	1.33	0.000
EN III	0.03	0.40	1.33	0.47	1.33	0.000
EN IV	0.00	0.40	1.33	0.40	1.33	0.000

Table 3
Results for 10-Node Random Networks

RNs	$\mathbf{LP}(\phi)$		\mathbf{MLP}		
	$ \mathcal{L} $	FD	Δf	# Iterations	FD
17	0.311	43.779	7	0.000	0.000
24	0.422	71.491	5	0.000	0.000
30	0.333	3.070	2	0.000	0.000
35	0.222	0.390	5	0.000	0.000
39	0.133	0.183	6	0.000	0.000

Table 4
Results for 20-Node Random Networks

RNs	$\mathbf{LP}(\phi)$		\mathbf{MLP}		
	$ \mathcal{L} $	FD	Δf	# Iterations	FD
37	0.394	65.652	3	0.000	0.000
70	0.500	59.019	6	0.150	0.050
99	0.463	156.474	10	0.047	0.052
124	0.347	23.220	10	0.005	0.002
145	0.237	0.595	3	0.005	0.006

6 Summary

In this work, we show the relation between the slopes of the FT function and the link weight for OSPF traffic engineering through linear programming duality. We have developed a heuristic that updates the FT function in such a way that the modified function is still piece-wise linear and is a convex envelope of the FT function; furthermore, the modified function leads to generating link weights from the dual solution which is better suited for generating single-shortest paths for most pairs in a network. We present numerical examples to show that our approach is very effective in obtaining single shortest paths for most pairs without drastic increase in the optimal cost or maximum link utilization.

References

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," *Proc. IEEE Conference on Computer Communications (INFOCOM'2000)*, pp. 519–528, 2000.
- [3] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 756–767, 2002.
- [4] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*, Morgan Kaufmann Publishers, 2004.
- [5] M. Pióro, A. Szentési, J. Harmatos, A. Jüttner, P. Gajowniczek, and S. Kozdrowski, "On OSPF related network optimization problems," *Performance Evaluation*, vol. 48, pp. 201–223, 2002. (A preliminary version of this paper appeared in the Proc. IFIP ATM IP 2000, Ilkley, England, July 2000).
- [6] M. Thorup, M. Roughan, "Avoiding Ties in Shortest Path First Routing," Technical Report, AT&T Labs-Research, 2003.
- [7] Y. Wang, Z. Wang, and L. Zhang, "Internet traffic engineering without full mesh overlaying," *Proc. IEEE INFOCOM'2001*, pp. 565–571, New York, USA, 2001.
- [8] A. Sridharan, R. Guerin, C. Diot, "Achieving Near-Optimal Traffic Engineering Solutions for current OSPF/IS-IS Networks," *Proceedings of INFOCOM'2003*, San Francisco, 2003.