

Optimizing OSPF/IS-IS Weights in a Changing World

Bernard Fortz, *Associate Member, IEEE* and Mikkel Thorup

Abstract—A system of techniques is presented for optimizing open shortest path first (OSPF) or intermediate system–intermediate system (IS-IS) weights for intradomain routing in a changing world, the goal being to avoid overloaded links. We address predicted periodic changes in traffic as well as problems arising from link failures and emerging hot spots.

Index Terms—Combinatorial optimization, intermediate system–intermediate system (IS-IS), local search, open shortest path first (OSPF), shortest path first, traffic engineering, traffic management.

I. INTRODUCTION

THE OPTIMIZATION goals of traffic engineering are to enhance the performance of Internet protocol (IP) traffic while utilizing network resources economically. In this paper, our focus is on *optimizing open shortest path first (OSPF) or intermediate system–intermediate system (IS-IS) traffic routing so as to make more efficient use of network resources in scenarios of changes to traffic and networks.*

Shortest path first (SPF) protocols such as OSPF [23] or IS-IS [8] are the most commonly used intradomain Internet routing protocols today. Traffic is routed along shortest paths to the destination. The weights of the links, and thereby the shortest path routes, can be changed by the network operator. A simple default weight setting suggested by Cisco [11] is to make the weight of a link inversely proportional to its capacity. As an alternative to OSPF/IS-IS, the more flexible multiprotocol label switching (MPLS) protocol has been suggested [5], [25]. MPLS is not yet widely deployed, but in principle, it would allow arbitrary routing in networks.

Our general objective in this paper is to route demands through an OSPF/IS-IS based network so as to avoid congestion in terms of link loads exceeding capacities with resulting packet loss and back off in TCP.

In the context of a *fixed* network with a fixed known demand matrix this problem has already been addressed experimentally in [20] with real and synthetic data, showing that we can find weight settings supporting 50%–110% more demands

than Cisco's defaults inverse-capacity-weights, and get within a few percent of the best possible with general routing, including MPLS. Similar positive findings have been reported in [6], [14], [21], and [24]. Here, the demand matrix could be based on concrete measurements, as described in [17] (see also [10] and [12]), but could also be based on concrete service level agreements (SLAs).

However, as stipulated in [3], *demand matrices and networks change*. The obvious idea for dealing with change is to just reset the weights using the above mentioned techniques. However, as we shall see shortly, there are several reasons why one should *avoid weight changes* as much as possible. Our target in this paper is to match the above mentioned improvements for fixed networks and demand matrices in scenarios of changing networks and demand matrices, changing as few weights as possible.

Why Weight Changes are Bad: There are multiple reasons why weight changes are to be avoided as much as possible. First we note that even a single weight change is disruptive for a network. The weight change has to be flooded in the network. As the routers learn about the change, they recompute their shortest paths to update their routing tables, and it may take seconds before all routers agree on the new shortest paths. Meanwhile, packets may arrive out of order, degrading the performance of TCP. Obviously, the more weight changes we try to flood simultaneously, the more chaos we introduce in the network with packets being sent back and forth between routers with very diverging pictures of the weights and shortest paths in the network.

A quite different problem with many weight changes is that a human network operator oversees the configuration of weights and is making sure that the resulting routing is well behaved, making sure that a diverse set of requirements is satisfied. These requirements may be very specific to the network, and may never be explicitly formulated until the network operator sees a concrete problem. In contrast, we are developing a generic tool for reducing congestion in arbitrary networks according to a specific mathematical formulation of congestion. Indeed, congestion is important, but our tool can only act as a decision support system over which the human operator has full control. If the tool suggests a few congestion reducing weight changes, the operator has a chance of checking that these are consistent with other requirements, but if we suggest too many changes, he is less likely to be able to endorse them.

In fact, our tool allows the operator to limit the range of different weights. If he does not approve of suggested weight changes, he can limit the ranges, thus, forcing the tool to come up with new suggestions. This kind of interactive play is only

Manuscript received March 8, 2001; revised December 20, 2001. The work of B. Fortz was supported by the AT&T Research Prize 1997.

B. Fortz is with the Institut d'Administration et de Gestion, Université Catholique de Louvain, Louvain-la-Neuve, Belgium (e-mail: fortz@poms.ucl.ac.be).

M. Thorup is with AT&T Labs-Research, Shannon Laboratory, Florham Park, NJ 07932 USA (e-mail: mthorup@research.att.com).

Publisher Item Identifier S 0733-8716(02)03143-8.

meaningful if the number of weight changes is so limited that a human can comprehend them. We note that limiting the range of weights is technically trivial, and it will not be discussed further in this paper.

Good Performance With Few Weight Changes: Motivated by the problems in changing weights, our first contribution is a technique for optimizing weights, changing as few weights as possible.

The technique was applied to a real weight setting of the AT&T IP backbone, and it was found that increasing a single weight from 1024 to 1025 reduced the max-utilization by 8%. Here, the *utilization* of a link is the ratio of its load over its capacity, so utilization above 100% means an overloaded link. The max-utilization is the maximum utilization over all links in the network. Thus, from the perspective of max-utilization, the weight change was worth an 8% increase in link capacities. Checking the impact of a single weight change is relatively easy, and executing it is much cheaper than buying and installing new links with higher capacity.

Our technique was also tried on the experimental networks from [20], including a proposed AT&T IP backbone as well as the synthetic two-level graphs of Zegura *et al.* [9], [29], [30]. Here it was found that if we started with Cisco's suggested default link weights inversely proportional to link capacities, using at most ten weight changes, we could support around 50% more demands, achieving about 2/3 of the gains reported in [20] with a complete change of all weights.

Of particular interest, we considered the problem of reestablishing performance after a local change such as a link failure or a new hot spot. That is, our starting point was a set of weights optimized for a network with a given demand matrix. We then simulated all possible single link failures and hot spots, and discovered that even in the worst cases, three weight changes sufficed to get within a few percent of the best possible with a complete dynamic change of all weights.

We note that improved routing may cause TCP to send packets more aggressively, thus, changing the demand matrix. Allowing TCP to send more traffic thanks to less loaded links and packet loss is in itself considered positive, and here we ignore the more complicated interaction. We note, however, that we in response to demand changes from TCP could apply the weight changes iteratively.

Periodic Daily Changes: Even on a daily basis there are large structural differences between day and evening traffic. However, this pattern is relatively similar for different days, that is, the traffic undergoes quite predicted periodic changes on a daily basis [17], [18].

The most obvious idea for dealing with daily changes would be to use the above techniques to adapt the weights to the changes. However, even a few daily weight changes are best avoided.

We propose here to seek *one* weight setting that is good for *all* the typical periodic changes happening during a day. To deal with periodic changes, we suggest optimizing the weight setting based on a few representative demand matrices, such as one day and one evening matrix. Our basic idea is that if one set of weights works well for some representative demand matrices, it works as well for all demand matrices dominated (\leq in each

entry) by convex combinations of these representative demand matrices, thus, allowing us to cover an infinite space of demand matrices rather than just a few singular points. This combined with general robustness to noise gives us a chance of addressing all the periodic changes happening throughout a day.

For our test instances, we succeeded in finding single weight settings for pairs of structurally different demand matrices that simultaneously for each individual demand matrix gave routing close to the optimum for general routing (including MPLS) for that demand matrix. That is, our fixed weight setting is competitive with ideal dynamic routing with unlimited power to adapt independently to each demand matrix.

This finding is somewhat surprising in that it is easy to construct pairs of demand matrices so that however we fix the routing, even with MPLS, it is going to be bad for at least one of the demand matrices.

Our successful optimization over multiple demand matrices is also interesting for *differentiated services* with quality-of-service (QoS) constraints. Here, we want to find one weight setting that on the one hand satisfies certain performance guarantees for high priority customers, and on the other hand gives good best effort service for the general traffic. Having one demand matrix for the high priority customers and another for the regular traffic, we can apply the above mentioned techniques.

Summary of Contribution: This is the first paper addressing the problem of finding good weight settings for OSPF/IS-IS without the freedom to change all weights. We believe that respecting the problem in changing weights takes us from a nice classical style combinatorial optimization problem to something of true practical relevance.

Our technique for tuning a weight setting with as few weight changes as possible is a novel general adaptation of the classical local search approach of combinatorial optimization [1], and we expect it to find many other applications in areas where parameter changes are undesirable. From a programming perspective, the adaptation had the advantage that we could largely reuse our previous local search code for a good setting of all weights [20].

Concerning periodic changes, our main new idea is *not* to do the obvious thing of precomputing weight changes for a transition between, say, day and evening, but to strive for one weight setting good for the whole day.

The value of our techniques is demonstrated experimentally on synthetic and real networks, showing that we can improve significantly over standard default weight settings, and in fact get close to the best dynamic routing with the more general MPLS scheme. As pointed out in [4], [19], and [20], it is trivial to construct networks and demand matrices for which the best OSPF/IS-IS routing is worse than the best MPLS routing. Also, in this paper, we present a negative example with two demand matrices that do not have a good common routing, even with MPLS. Thus, our positive results cannot be achieved in general, but they indicate that the negative examples are unlikely to dominate in real networks.

In combination, our techniques provide an efficient tool for OSPF/IS-IS traffic management in a changing world. It is now an integrated part of AT&T's NetScope/Bravo traffic engineering tool for IP networks [16]. Note that our work may also be applicable to MPLS traffic engineering when label-switched paths

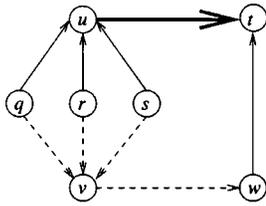


Fig. 1. The need for global weights optimization.

Link	Default unit weights		Increasing weight of overloaded link		Optimal global single change	
	Weight	Load	Weight	Load	Weight	Load
(q, u)	1	1	1	0.5	1	1
(r, u)	1	1	1	0.5	1	1
(s, u)	1	1	1	0.5	3	0
(u, t)	1	3	2	1.5	1	2
(q, v)	1	0	1	0.5	1	0
(r, v)	1	0	1	0.5	1	0
(s, v)	1	0	1	0.5	1	1
(v, w)	1	0	1	1.5	1	1
(w, t)	1	1	1	2.5	1	2

or tunnels become logical links with weights like the physical links in OSPF/IS-IS.

Contents: First, in Section II, we define our exact model and objectives. Next, in Section III, we present our technique for optimizing with few weight changes. In Section IV, this technique is applied in connection with link failures and hot spots. Then, in Section V, we describe and evaluate our approach with multiple demand matrices. Finally, we have some concluding remarks in Section VI.

II. MODELS AND EXPERIMENTAL TEST-BED

This section presents the basic framework from [20] used to evaluate weight settings for a given network and demand matrix. Whereas [20] kept the network and demand matrix fixed, we are going to show how to deal with changes in the subsequent sections. The original presentation of the framework from [20] was rather short, and here we provide much more motivating discussions and background.

A. The General Routing Problem

Optimizing the use of existing network resources can be seen as a general routing problem defined as follows. We are given directed graph $G = (N, A)$, $A \subseteq N \times N$, whose nodes and arcs represent routers and the links between them, arc capacities $c := (c_a)_{a \in A}$, and a demand matrix D that, for each pair (s, t) of nodes, tells us how much traffic flow we need to send from s to t . We refer to s and t as the source and the destination of the demand. Many of the entries of D may be zero, and in particular, $D[s, t]$ should be zero if there is no path from s to t in G . A routing solution specifies for each source–destination pair how the demanded traffic should flow in the network. The load ℓ_a on an arc a is then the total traffic flow through the arc, including the contributions from each source–destination pair.

In the left side of Fig. 1 is a concrete example of a network in which each of the nodes q , r , s , and w demands a flow of 1 to t . In the displayed routing solution, the flows just follow shortest paths to the destination, and thereby a load of three is accumulated on link (u, t) , whereas, all other links have loads zero or one. The table on the right will be discussed in the next subsection.

For real instances of the problem, additional complicating constraints such as nodes forbidden for transit traffic or point-to-multipoint demands arise [17]. These kind of constraints can be integrated by modifying the graph including artificial links, but

these constraints do not affect the methods and results presented here and are left out for the sake of clarity.

So far, we have been rather vague about our objective of “avoiding overloaded arcs,” and we will now define some more exact objectives. Recall that the utilization of an arc a is the load divided by the capacity, i.e., ℓ_a/c_a , and a link is overloaded if the utilization exceeds 100%. The max-utilization is the maximum utilization over all links. Thus, if all links in Fig. 1 had capacity 2, the max-utilization would be $3/2 = 150\%$.

Minimizing the max-utilization as in [21] is a natural and intuitive objective for routing. In our example, we can reduce the max-utilization to one by sending the traffic from s via v and w to t instead of via u . We note that there may be some links for which we are particularly concerned about high utilization, but we can just view such links as having a reduced capacity. We will consider the max-utilization in this paper, but it suffers from allowing a single bottle-neck, e.g., an ingress link from another domain over which we have no control, to dominate the whole picture. Also, it does not penalize using very long detours. To get a measurement considering the whole network, we consider cost functions of the form

$$\Phi = \sum_{a \in A} \phi(\ell_a, c_a)$$

summing a cost $\phi(\ell_a, c_a)$ from each arc a depending on the relation between the load ℓ_a and the capacity c_a . In [24], they use a queuing theory style link cost function $\phi(\ell_a, c_a) = \ell_a/(c_a - \ell_a)$. With this function, it is more expensive to send flow along arcs whose loads approach capacity, which is what we want. However, the function does not deal with overloaded links, and in reality, a single overloaded link does not take down a whole network. To overcome this problem, we resorted to a piece-wise linear approximation of $\ell_a/(c_a - \ell_a)$, defined with heavy penalty for overloaded links. More precisely, for some fixed capacity c_a , we define $\phi(x, c_a)$ as the continuous function with $\phi(0, c_a) = 0$ and derivative in the load ℓ_a of

$$\phi'(\ell_a, c_a) = \begin{cases} 1 & \text{for } 0 \leq x/c_a < 1/3 \\ 3 & \text{for } 1/3 \leq x/c_a < 2/3 \\ 10 & \text{for } 2/3 \leq x/c_a < 9/10 \\ 70 & \text{for } 9/10 \leq x/c_a < 1 \\ 500 & \text{for } 1 \leq x/c_a < 11/10 \\ 5000 & \text{for } 11/10 \leq x/c_a < \infty \end{cases} \quad (1)$$

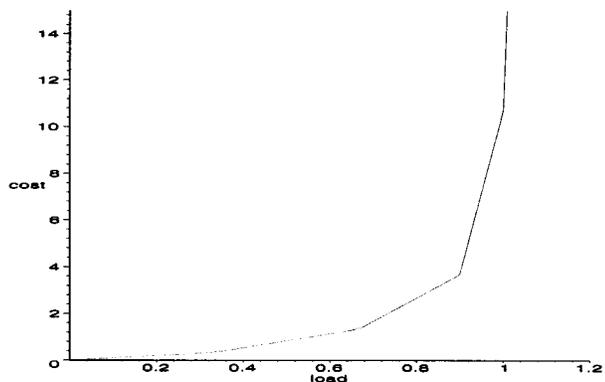


Fig. 2. Arc cost $\phi(\ell_a, 1)$ as a function of load ℓ_a with capacity $c_a = 1$.

The arc cost function ϕ is illustrated in Fig. 2. Generally, it is cheap to send flow over an arc with a small utilization ℓ_a/c_a . The cost increases progressively as the utilization approaches 100%, and explodes when we go above 110%.

Because of the explosive increase in cost as loads exceed capacities, our objective typically implies that we keep the max-utilization below 1, or at least below 1.1, if at all possible.

The objective function was chosen on the basis of discussions on costs with people close to the AT&T IP backbone. The exact coefficients are not important. We tried many variations and found that this did not change the quality of our results. Also, it had no substantial impact to use a smoother objective function with smaller segments. More importantly, the routing solutions found were very robust to changes in the objective function. In particular, when optimizing routings for Φ , our solutions tended to also do very well with respect to max-utilization.

The piece-wise linearity of our cost function has the advantage that using a linear programming (LP) solver, we can find the optimal solution to the general routing problem with no limitations to how we can distribute the flow between the paths. We can then compare ourselves against this unrealistic ideal to see how competitive we are with *any* other approach, including MPLS.

A problem in the current formulation of Φ is that it does not provide a universal measure of congestion. With the max-utilization, it is clear for any routing in any network that we have a problem if it exceeds 100%, and we would like a similar universal cutoff for our summed link costs. To achieve this, we use a normalized cost function¹

$$\Phi^* = \Phi/\Psi$$

where Ψ is the cost we would have had if all flow was sent along hop-count shortest paths and the capacities matched the loads. When capacities match loads, we pay $\phi(1, 1)$ per unit load over a link, so if $\Delta(s, t)$ is the hop-count distance between s and t , $\Psi = \sum_{(s,t) \in V^2} (D[s, t] \cdot \Delta(s, t) \cdot \phi(1, 1))$. Note that for a given network and demand matrix, the division by Ψ does not affect which routings are considered good. Summing up, $\Phi^* \geq 1$ implies that we are performing as badly as if all flows were along hop-count shortest paths with loads matching the capacities. The same cost can, of course, also stem from some

¹The normalization from [20] was defined differently so that it was $\phi(1, 1) = 10 (2/3)$ times bigger.

loads going above capacity and others going below, or by flows following longer detours via less utilized arcs. Nevertheless, it is natural to say that a routing *congests* a network if $\Phi^* > 1$.

B. OSPF/IS-IS Routing

As mentioned earlier, this paper focuses on routing with OSPF [23] and IS-IS [8], which are the most commonly used intradomain Internet routing protocols today. The network operator assigns a weight to each link, and shortest paths from each router to each destination are computed using these weights as lengths of the links. In each router, the next link on all shortest paths to all possible destinations is stored in a table, and a flow arriving at the router is sent to its destination by splitting the flow between the links that are on the shortest paths to the destination. The details of the splitting/tie-breaking depends on the configuration of the router. In this paper we assume a hash-based splitting, as used in the AT&T IP network: if a router has multiple outgoing links on shortest paths to a destination, packets are assigned an outgoing link based on a hash function of some information in their header. The hashing ascertains that packets from the same flow follow the same path, which is important for the packets to arrive in the order they were transmitted. The hash-based splitting generally results in a roughly even split, and for simplicity, in our analysis and simulations, we assume that the split is exactly even.

In [19] and [20], it was found that Cisco's [11] suggested default of setting weights inversely proportional to the capacity was as good or better than setting all weights to one, or using physical distance. This justifies that inverse capacity is the only default weight setting that we compare ourselves against in the current paper.

C. The Need for Global Optimization

A naive approach to the problem of trying to improve a given routing by modifying only one or few weights would be to just increase the weight of overloaded links, thus, avoiding the need for sophisticated tools as presented in this paper. (Un)fortunately, such a strategy is often insufficient to provide the best results. To illustrate this, reconsider the network of Fig. 1 assuming that all capacities are two. Suppose that we want to send one unit of flow from q , r , s and w to t . The table in Fig. 1 presents the loads resulting from three different weight settings. For units weights, we get a maximal load of three on link (u, t) . If we wish to decrease this maximal load by increasing the weight of this heavily loaded link, the best solution is to set the weight to two. We then get that the flows coming from q , r and s are split along the two possible paths, and the maximal load becomes 2.5 on link (u, t) . If we instead look for the best change over all links, we easily see that setting the weight of link (s, u) to three changes the routing for the flow from s to t , but not those from q and r to t , and the maximal load in the network is thereby decreased to two, matching the capacity.

Note, that in this optimal solution, the shortest paths are unique, hence, that we made no use of the even splitting feature. Rather it just singles out part of the traffic from s and redirects it. A similar pattern was found in many of the optimized weight settings from [20]. A somewhat similar finding is reported in

[26], where the authors, in a slightly different context, show that the best routing strategy usually is obtained with a limited number of different paths.

D. Test Instances

Our basic experimental networks and demand matrices are the same as in [19] and [20]. We have a proposed AT&T IP backbone with 90 nodes, 274 arcs and projected demands. Also, we have a synthetic two-level networks produced using the generator GT-ITM [29], based on a model of Calvert, Bhattacharjee, Daor, and Zegura [9], [30]. This model places nodes in a unit square, thus, getting a distance $\delta(x, y)$ between each pair of nodes. These distances lead to random distribution of two-level graphs, with arcs divided in two classes: *local access* arcs and *long distance* arcs. Arc capacities were set equal to 200 for local access arcs and to 1000 for long distance arcs. Many networks were generated and tested, also including Waxman graphs [28] and random graphs, but for space reasons, we only present experiments over one two-level graph with 50 nodes and 148 arcs. Results for other networks are consistent with those reproduced here and follow the same patterns.

The above synthetic network model does not include a model for the demands. Inspired by classical entropy models for urban traffic [27], we decided to model the demands as follows. For each node x , we pick two random numbers $o_x, d \in [0, 1]$. Further, for each pair (x, y) of nodes we pick a random number $c_{(x,y)} \in [0, 1]$. Now, if the Euclidean distance (L_2) between x and y is $\delta(x, y)$, the demand between x and y is

$$\alpha o_x d_y c_{(x,y)} e^{-\delta(x,y)/2\Delta}, \quad (2)$$

Here, α is a parameter and Δ is the largest Euclidean distance between any pair of nodes. Above, the o_x and d_x model that different nodes can be more or less active senders and receivers, thus, modeling hot spots on the net. Because we are multiplying three random variables, we have a quite large variation in the demands. The distance factor $e^{-\delta(x,y)/2\Delta}$ implies that we have relatively more demand between close pairs of nodes, yet the distance on its own never has an impact bigger than a factor $\sqrt{e} = 1.648 \dots$. In our experiments, we also tried not using the distance factor $e^{-\delta(x,y)/2\Delta}$, and the results were essentially unchanged. This model—without the distance factor—has also been used in [13] for Internet traffic. They also consider voice and transaction data but with a large distance factor.

In our experiments, we generated one demand matrix for each network, then we scaled it (by multiplying each entry by a constant) at different levels to obtain different total demands. All our results are reported for these increasing total demands, which allows to measure, for a given routing, at which level of total demand congestion occurs for the given topology and demand pattern.

Our technique for few weight changes was also tested on the real AT&T IP backbone with its real weight setting and measured demands [17].

Discussion of Test Instances: We will now briefly discuss our choice of synthetic data. On the positive side, for the network model itself, it has already been established that it gives good approximations to real networks in many aspects [9], [30].

Further, the similarity in results found in [20] between the proposed AT&T IP backbone and our synthetic networks indicate that the model is reasonably faithful for our purpose.

On the negative side, it has recently been found that the degree distribution for IP networks has heavy tails [15] (we note that [15] actually considers the Internet as a whole, whereas, we are only considering a single domain), and similar findings have been found for the demands [17], that is, we should have some nodes with very high degrees, and some very large demands. Our synthetic models do not provide such heavy tails.

Generating networks and demand matrices with heavy tails is in itself easy [2]. For the networks, one could even use some of the official networks available at [7]. However, in real networks, the heavy tails of the networks and demand matrices are highly dependent. Generally, heavy tails means structure. High degree nodes are placed at strategic positions relative to the structure of the expected demand matrix. Indeed, we have found that AT&T's network performs very poorly on a random demand matrix. A random heavy tail demand matrix placing a large demand in, say, Alaska, would be even worse. Thus, networks and demand matrices with heavy tails have to be tailored to each other to be realistic. So far, there is no established model for such tailored combination, and tailoring the input data ourselves would make our results less convincing.

The advantage of Zegura's networks is that they are constructed in a generic way not favoring any particular demands, and for that reason they work reasonably well for generic demand matrices such as ours.

Thus, we argue there is no obvious better choice of a model, given that, to claim good optimization results, we cannot convincingly construct a model by hand. However, for the sake of optimization work like ours, we see it as an important challenge to develop a single heavy tail model integrating both IP networks and demand matrices where the IP network is reasonably tuned for the demand matrix. In fact, as argued above, we think that such an integrated approach may give a better understanding of the heavy tails found in IP network, viewing the height degree nodes as strategically positioned hops.

III. FEW WEIGHT CHANGES

In this section, we consider the problem of optimizing, making as few weight changes as possible. We apply our technique for few changes in two scenarios. One is when the weight setting has not previously been optimized, e.g., if the network is so far just run with inverse capacity weights following Cisco's recommendation [11]. The other one, described in Section IV, is in response to problematic changes in demands or topology like critical link failures or emerging hot spots.

A. Technique for Few Changes

We want to make as few weight changes as possible. The local search from [19] and [20] that we used for multiple demand matrices works with a single solution that is iteratively improved by small changes in the current solution. It typically performs a lot of iterations (5000 in our practical experiments) and, therefore, produces a solution completely different from the starting one. This can be seen as a depth-first search in the solution space, but

```

Number of iterations set to 3.
Will profile top 4 utilized links.
Demand scaling factor set to 2.916967.
Data file: experiments/hier100.dat

No input weights. Using InvCap as default.
Scaling factor : 16113.563
InvCap : 8.873 (1.210)
Utilization profile:
69:2.420e+02/2.000e+02= 1.210
170:1.639e+02/2.000e+02= 0.820
176:1.584e+02/2.000e+02= 0.792
89:1.564e+02/2.000e+02= 0.782
Iteration 1 : 1.323 (0.820)
w[69] : 5 -> 11
Utilization profile:
170:1.639e+02/2.000e+02= 0.820
69:1.623e+02/2.000e+02= 0.811
89:1.564e+02/2.000e+02= 0.782
81:1.533e+02/2.000e+02= 0.766

Iteration 2 : 1.308 (0.811)
w[69] : 5 -> 11
w[170] : 5 -> 10
Utilization profile:
69:1.623e+02/2.000e+02= 0.811
89:1.564e+02/2.000e+02= 0.782
81:1.533e+02/2.000e+02= 0.766
231:1.499e+02/2.000e+02= 0.750
Iteration 3 : 1.294 (0.782)
w[100] : 5 -> 1
w[69] : 5 -> 12
w[170] : 5 -> 9
Utilization profile:
89:1.564e+02/2.000e+02= 0.782
81:1.533e+02/2.000e+02= 0.766
180:1.466e+02/2.000e+02= 0.733
63:1.456e+02/2.000e+02= 0.728
Adaptive CPU time : 221
Best found : 1.294 (0.782)
    
```

Fig. 3. Optimizing with few changes on two-level graph with 100 nodes and 280 arcs.

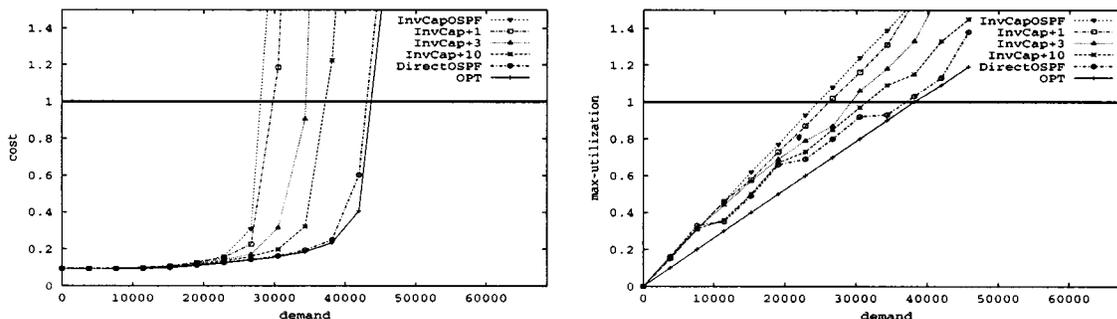


Fig. 4. Few changes on AT&Ts proposed backbone.

since we want to make as few changes as possible, our approach should rather be a breadth first search.

Our heuristic for improving an input weight setting w_0 with as few weight changes as possible works as follows: first we consider about 1000 single weight changes to w_0 , corresponding to about five weight changes for each arc in our largest networks. As in [19] and [20], we limit the number of weight changes considered by applying random sampling to the neighborhood structure, as exploring the full neighborhood is too time consuming. Instead of selecting only the best weight change as in [19] and [20], we keep the 100 best weight changes in a family F of “best weight settings.” The process is iterated with F instead of w_0 : we consider 1000 single weight changes for each weight setting in F and a new F is selected containing the 100 best of the old weight settings in F and the about 100 000 new weight settings considered. After i iterations, including the start from w_0 , the family consists of weight settings with up to i weight changes from w_0 . The size of F corresponds to the breadth of our search. All the above numbers are just parameters that experimentally were found to give a good compromise between quality of solution and time. For the largest networks considered, 10 iterations took about 1 1/2 hours on a single processor of a 194 MHz CPU SGI Challenge XL. If running time is an issue, one can, for example, reduce the family size or the number of iteration, or even better, buy a better faster computer or parallelize the code.

This technique for few changes has, to our knowledge, not been used before. Its main interest is that it provides a general framework for optimizing with few changes that can be easily adapted for other applications. It has the advantage that if a local search such as our previous heuristic from [20] is available, the

main ingredients such as the changes applied to one solution to get a neighbor of it, or the procedures to evaluate a new solution, can be reused in this new framework, saving a lot of implementation work.

B. Experiments With Few Changes

As our base experiment with few changes, we took the inverse-capacity weight setting (InvCapOSPF) as a starting point. A concrete output for running the code for 3-iterations is presented in Fig. 3.

It tells the user how much improvement can be obtained with one, two, and three changes and which changes to make. Further, it shows more precisely how the loads on the most utilized links change as more changes are applied.

The results of applying our technique are presented in Figs. 4 and 5, with InvCap+ k denoting that we applied k changes. DirectOSPF represents a weight setting, involving all weights, optimized with our local search from [20], and OPT is the solution of the general routing problem, including the possibilities with MPLS. Note that the variations in the right side of Figs. 4 and 5 are due to the fact that the optimization is done using our metric and that max-utilization is computed afterwards.

First, as in [19], and [20], we note that all curves start off pretty flat, and then, quite suddenly, start increasing rapidly. This behavior follows our cost function that explodes when the load of a link reaches its capacity [cf. (1), Fig. 2].

The most interesting comparison between the different schemes is the amount of demand they can cope with before the network suddenly gets congested in the sense that its normalized cost exceeds one. For the real network, one change

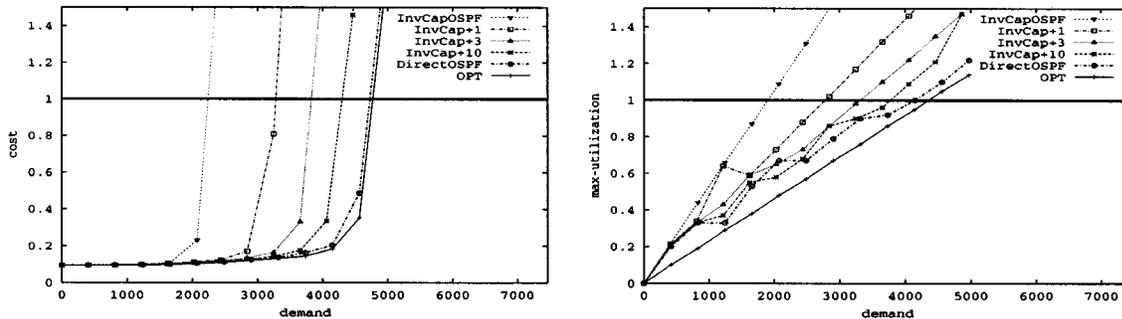


Fig. 5. Few changes on two-level graph.

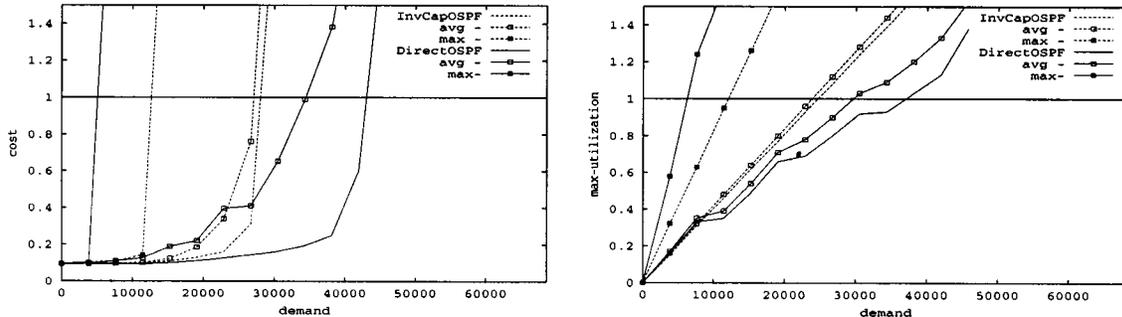


Fig. 6. Link failure on AT&T's proposed backbone.

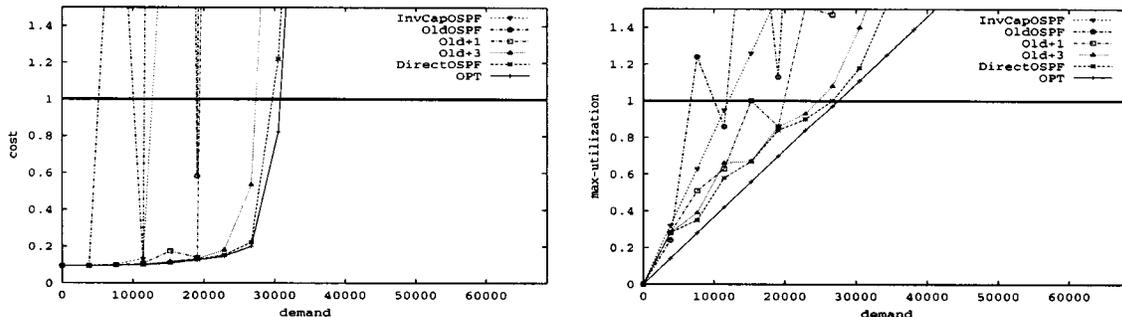


Fig. 7. Resurrecting performance after worst possible link failure in AT&T's proposed backbone.

gives 10%, three changes gives 25% while ten changes gives 35%. On the synthetic network, one change gives 50%, three changes give 75% while ten changes give 90%. Thus, quite substantial improvements can be obtained with few changes.

The code was also run with on the real AT&T IP backbone with its real weight setting and measured demands. It was found that changing the weight of a single link from 1024 to 1025 improved the max-utilization by 8%. Further details on this are, however, proprietary.

IV. LINK FAILURES AND HOT SPOTS

In this section, we consider the problem of link failures and emerging hot spots. We show that they are not normally problematic. However, there are a few critical cases, but for these, we reestablish good performance with at most three weight changes.

A. Link Failures

For AT&T's proposed IP backbone, we tried all possible link failures, and computed the routing both with InvCapOSPF, and

with a weight setting optimized from before the link failure (DirectOSPF). The results are depicted in Fig. 6, giving the values, first without link failures, then as average over all link failures, and finally with worst-case link failures. Here, by worst case, we mean the link whose deletion decreases the objective function the most for the given weight setting. This link is, thus, found independently for each weight setting considered.

For average link failures, we see that DirectOSPF performs 25% better than InvCapOSPF, which is still pretty good. However, for worst-case link failures for the proposed AT&T IP backbone in Fig. 6, we see that our optimized weights do 65% worse than InvCapOSPF. Shortly, we will turn this defeat into victory.

If we remove the worst-case link that caused the biggest problem for weight settings found by our DirectOSPF for the proposed AT&T IP backbone and reoptimize with as few weight changes as possible, as described in the last section, we obtain results depicted in Fig. 7. Here, OldOSPF denotes our optimized weight setting from before the link failure, and Old+ k denotes the result of optimizing this weight setting with

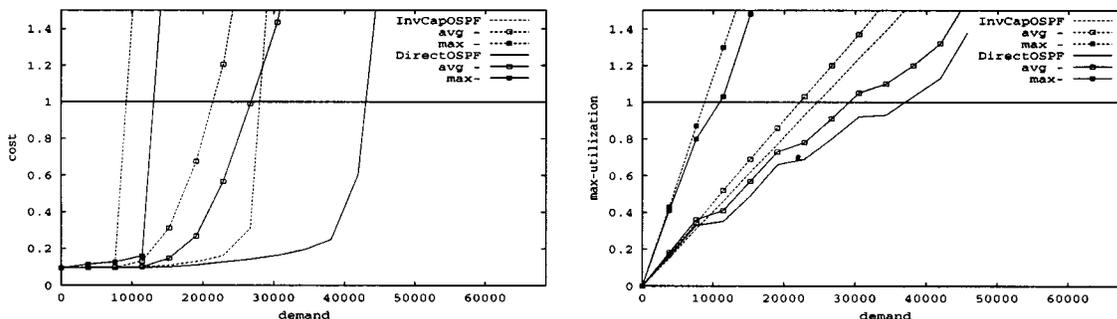


Fig. 8. Hot spot on AT&T's proposed backbone.

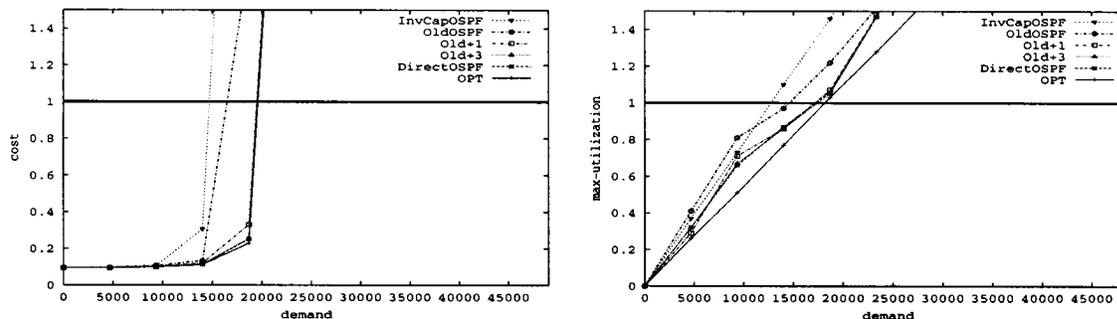


Fig. 9. Resurrecting performance the worst hot spot in AT&T's proposed backbone.

k changes. InvCapOSPF, DirectOSPF, and OPT are all based directly on the network without the bad link.

First, observe that OldOSPF comes down and performs as well or better than InvCapOSPF for demands scaled around 11 000 and 19 000. The reason for this is that for the original weight settings generated for these two demand scalings, it was actually a different link that was bad. However, for all the other demand levels, it is the worst possible link we have deleted. This essentially means that our weight setting had some different choices in what links to make critical, and the choice happened to come out differently for two of the scalings. The jumps nearly disappear after just one weight change, which makes us beat InvCap everywhere.

What we now see is that after just one weight change, we do 60% better than InvCap and with three changes we do 160% better, getting within 10% of OPT.

B. New Hot Spots

Our next experiments concern the developments of new hot spots. The essential experiment is to take one router and multiply all incoming demands by a factor of three. This corresponds to multiplying one d_x by three in (2), which again amounts to multiplying all entries in a column of the demand matrix by three. This is likely to give more structural difference than just noise, turning links toward x into bottlenecks.

The results of the experiments are depicted in Fig. 8. As for link failures, the average and the worst performance are obtained over hot-spotting all nodes, one at the time. We do this with the weight settings both from InvCapOSPF and DirectOSPF. We see that DirectOSPF performs 20% better than InvCapOSPF for average hot spots, and 30% better for worst-case hot spots. Without hot spots, DirectOSPF performed 50% better than InvCapOSPF. Thus, we do lose some of our

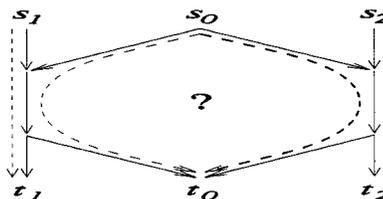


Fig. 10. Bad choice.

advantage over InvCapOSPF, but our optimized weight setting still provides a clear advantage over InvCapOSPF.

As for link failures, we tried to take the worst possible hot spot (in the sense that it deteriorates the most the objective function), and reestablish good performance. The result is depicted in Fig. 9. The nice thing is that we only need a single weight change to get back very close to optimum, which is even better than for the worst possible link failure.

V. MULTIPLE DEMAND MATRICES

Our motivation for working with multiple demand matrices is the general experience from AT&T that traffic follows quite regular periods with a peak in the day and in the evening. The network operators do not want to change weights on a regular basis so we want just one weight setting which is good for the whole period. We then collect a peak demand matrix for the day and one for the evening. A weight setting performing good on both performs good on all convex combinations and, hence, it has a good chance of performing well for the whole period.

A. An Impossible Example

In Fig. 10, it is illustrated that if we have to fix one set of routes for two different demand matrices, this may force us to

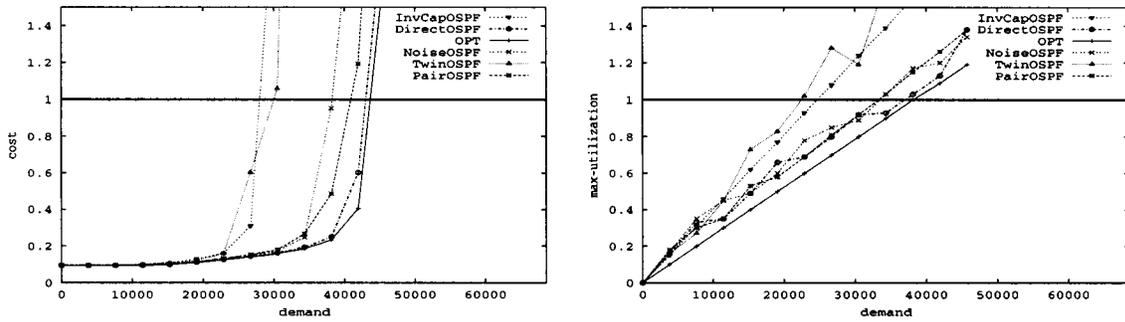


Fig. 11. Simultaneous optimization on AT&Ts proposed backbone and scaled projected demands (D_1).

increase the max-utilization by 50% for one of them. In the example, one demand matrix wants one unit of demand from s_0 to t_0 and from s_1 to t_1 while the other wants one unit of demand from s_0 to t_0 and from s_2 to t_2 . However, any flow from s_0 to t_0 has to share a link with either flow from s_1 to t_1 , or from s_2 to t_2 . In each situation independently, we can get max-utilization 1, but if we have to fix the routing from s_0 to t_0 , the best we can do is to split the s_0 – t_0 flow evenly, getting a max-utilization of $3/2$ in both cases. As for the negative examples in [20], this worst-case example is far from real networks, and we will see that in practice, real networks perform much better.

B. Optimizing for Multiple Demand Matrices

Given a network $G = (N, A, c)$ with several demand matrices D_1, \dots, D_k , we want to find a single weight setting $w := (w_a)_{a \in A}$ which works well for all of them. In general, we will use $\ell_a(G, D, w)$ to denote the load on link a with network G , demand matrix D , and weight setting w . Similarly for our cost function, we have $\Phi(G, D, w) = \sum_a \Phi_a(\ell_a(G, D, w))$ with Φ_a as defined in (1).

Now consider a demand matrix D dominated by a convex combination of D_1, \dots, D_k , that is $D \leq \alpha_1 D_1 + \dots + \alpha_k D_k$ where $\alpha_1 + \dots + \alpha_k = 1$. Here everything is understood to be entry-wise, so for all x, y , $D[x, y] \leq \alpha_1 D_1[x, y] + \dots + \alpha_k D_k[x, y]$.

Since the routing for each source–destination pair is fixed by the weight setting w , for each arc $a \in A$, $\ell_a(G, D, w) \leq \alpha_1 \ell_a(G, D_1, w) + \dots + \alpha_k \ell_a(G, D_k, w)$. In particular, it follows that the max-utilization for D is no worse than the worst max-utilization for the D_i . Further, since each arc cost function Φ_a is convex, $\Phi_a(\ell_a(G, D, w)) \leq \alpha_1 \Phi_a(\ell_a(G, D_1, w)) + \dots + \alpha_k \Phi_a(\ell_a(G, D_k, w))$ and, hence, $\Phi(G, D, w) \leq \alpha_1 \Phi(G, D_1, w) + \dots + \alpha_k \Phi(G, D_k, w)$. Thus, our weight setting w does no worse for D than for the worst of the D_i , neither with respect to our cost function Φ , nor with respect to max-utilization. Note that the same observation holds true with MPLS, as long as the routing for each source–destination pair is fixed.

From [19], we know that it is NP -hard even to approximate a good weight setting for a single fixed demand matrix, that is, unless $NP = P$, there cannot be any general method providing guaranteed good results. However, an efficient local search heuristic for the problem is suggested in [19] and [20] that for a given network G and demand matrix D looks for a weight setting w that minimizes $\Phi(G, D, w)$.

To optimize simultaneously for several demand matrices D_1, \dots, D_k , we simply modify the local search heuristic to minimize

$$\Phi(G, D_1, \dots, D_k, w) = \sum_{i \leq k} \Phi(G, D_i, w). \quad (3)$$

As in our original motivation for defining Φ , this has the effect of penalizing highly loaded links, this time, for all the demand matrices instead of just one. Our negative theoretical example shows that we cannot in general hope for good results, but we can still hope to do well in practice.

C. Experiments for Multiple Demand Matrices

In [19] and [20], there was only one demand matrix D_1 for each network. Here, we generated a second independent matrix D_2 , using the same distribution, and scaled to have the same total demand as D_1 . We call this demand matrix a twin of D_1 . The difference between D_1 and D_2 models structural differences in traffic, with a complete change in who are the big senders and receivers, thus modeling, e.g., the difference between day and evening traffic. We also derived a noisy version \tilde{D}_1 of D_1 that will be discussed later.

We compared OPT, InvCapOSPF, and DirectOSPF (obtained with our heuristic from [20]) against OSPF/IS-IS routings of D_1 using weight settings optimized relative to D_2 (TwinOSPF) and for D_1 and D_2 simultaneously (PairOSPF) as in (3). A symmetric set of experiments were performed with the roles of D_1 and D_2 interchanged. A good performance of PairOSPF on both D_1 and D_2 shows that we have successfully found one weight setting that compromised neither demand matrix.

The results of our experiments are presented in Figs. 11–14 with different scalings of the demand matrices.

In the experiments, we see that DirectOSPF allows us to cope with 50%–110% more demand than the InvCapOSPF. Also, DirectOSPF is less than 2% from being able to cope with the same demands as the optimal general routing OPT.

If we look at TwinOSPF, we see that it is often at least as bad, or worse than InvCapOSPF. Nevertheless, PairOSPF does very well, never getting more than 10% worse than DirectOSPF and OPT. In particular, PairOSPF is always doing simultaneously well for both demand matrices it optimizes over.

Put in terms of our day–evening example, this means that if we just optimize relative to a peak demand matrix from the day, we cannot expect to gain anything for the evening. However,

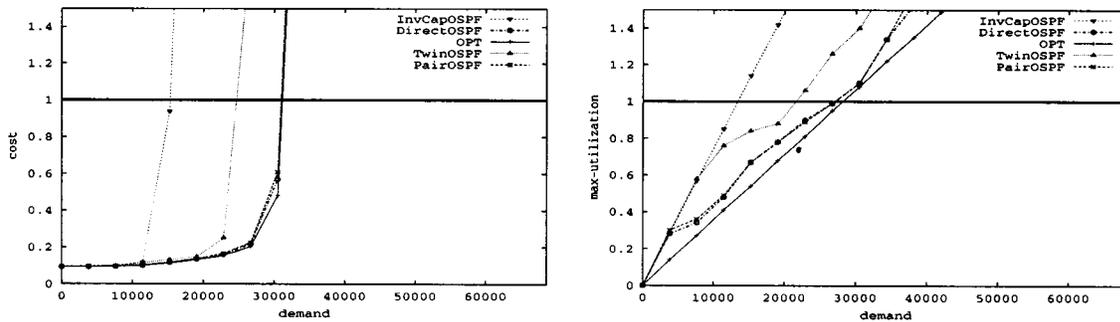


Fig. 12. Simultaneous optimization on AT&T's proposed backbone and synthetic twin demands (D_2).

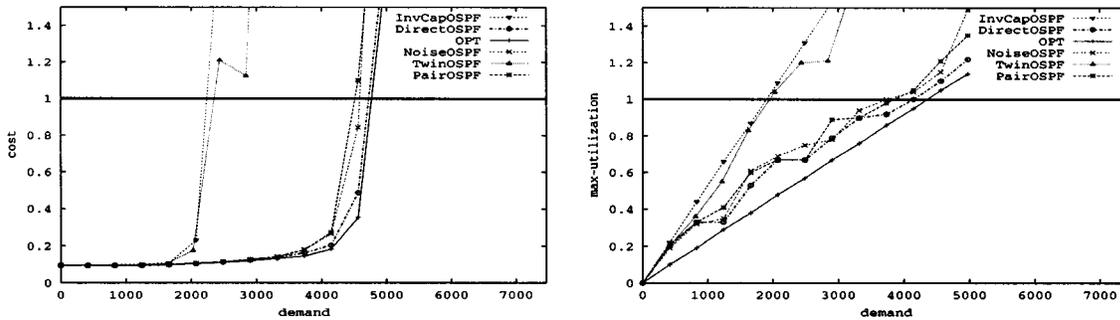


Fig. 13. Simultaneous optimization on two-level graph (D_1).

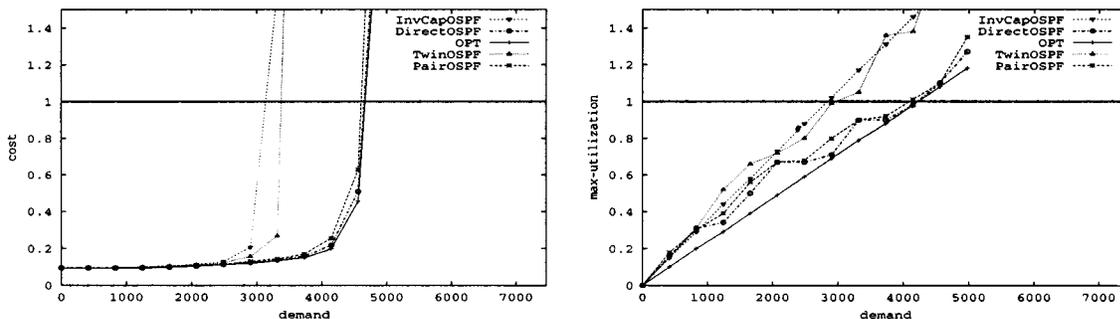


Fig. 14. Simultaneous optimization on two-level graph and twin demands (D_2).

if we do simultaneous optimization for a day and evening demand matrix, we can hope to do well on both. This combined with our robustness to noise (see Section V-D below) and our automatic benefit for all convex combinations gives us a good chance of dealing well for typical daily periodic changes with a single weight setting.

One may wonder why things work so well in spite of simple negative examples like the one presented in Fig. 10. A simple explanation comes from the fact that even if a bad example occurs as a subgraph of the network, real networks are usually large enough that we can steer the traffic around such a small bottleneck even with limited control over the routing. Indeed, in our experiments it usually turned out that only a small number of links were overloaded, so our bottlenecks were local, and making just a few good weight changes allowed us to smooth out these local discrepancies.

D. Robustness to Noise

We derived a noisy version \tilde{D}_1 of D_1 by multiplying each entry by a random number between zero and two. On the average

this changes each entry by 50% without changing the expected total demand. The discrepancy between D_1 and \tilde{D}_1 could represent problems of getting exact measures [17], or general fluctuations in traffic.

The results of our experiments with noise are also presented in Figs. 11 and 13, where we compared DirectOSPF against routings of D_1 using weight settings optimized relative to \tilde{D}_1 (NoiseOSPF). NoiseOSPF is doing quite well, gaining a minimum of 40% over InvCapOSPF. Thus, our weight settings are pretty robust to noise, and this implies that we only need a rough estimate of the demand matrix.

Note that the results obtained with D_2 with weights optimized for D_1 (TwinOSPF) are much worse than those with \tilde{D}_1 (NoiseOSPF), meaning weight settings are more sensitive to structural changes than to noise.

E. Differentiated Service

We now outline how our positive experience with multiple demand matrices could also be applied with differentiated service. Suppose we have two classes of customers: gold customers that

were promised a guaranteed bandwidth, and normal customers. We want to ensure that gold customers packets are routed in the first 60% of the capacity of each link. Suppose we can find a weight setting such that we remain below 60% of the capacity if we only send gold customers demands. Then, we can ensure with this weight setting that gold customers packets will be sent below 60% of the capacity by giving them the priority over normal customers using weighted fair queuing. Let D_1 be the demand matrix describing the bandwidth promised to gold customers. We need to route D_1 with max-utilization $3/5$. Recall that our objective function Φ is designed to keep max-utilization below one, which is also evident in our experiments with DirectOSPF. To satisfy our gold customers, keeping their max-utilization below $3/5$, we can, therefore, optimize with respect to $\Phi(G, (10/6 D_1), w)$.

However, this approach does not take normal customers into account, and could lead to really bad situations for them. To remedy this, let D_2 be a demand matrix estimating the total traffic, including the gold customers. We want to provide good best-effort service with respect to D_1 , while keeping the objective of getting the gold customers max-utilization below $3/5$. This problem has been addressed for MPLS [22] where it is just a multicommodity flow problem.

This leads us to a combined objective function of the form

$$\alpha\Phi(G, (10/6 D_1), w) + \Phi(G, D_2, w)$$

where α is a parameter that we can optimize. If α is sufficiently high, we only care about the first term, and then we are pretty sure to stay below 60% capacity for gold customers, if at all possible. However, as soon as we have satisfied the gold customers, we better start worrying about best effort service for everybody as in the second term, so the natural optimization is to use the smallest α for which the gold customers get satisfied. This smallest α can be found by a binary search.

The interesting thing here is that our successful experiments with multiple demand matrices indicate that we can satisfy both gold customers and normal customers with OSPF/IS-IS without compromising either.

VI. CONCLUSION

We have presented a system of algorithms for efficient OSPF traffic management in a changing world. As described in Section V, we can optimize efficiently over a few peak demand matrices, say representing day and evening US traffic and international traffic in the night, and thereby produce a weight setting covering all demand matrices dominated by convex combinations of these peaks. The noise tolerance from Section V implies that we do not need to find the absolute peaks, as long as we get within a reasonable neighborhood of them. Further, the results from Section IV indicate that our weight setting is good for new hot spots, and most link failures, so generally, we do not expect to have to change the weight setting.

However, by simulation, we may discover that a few critical link failures can cause problems. For these few links we precompute a few weight changes to be applied in case they fail. Sim-

ilar proactive strategies can be applied against other predicted changes.

In case the whole structure of the demand matrix evolves to a degree that our weight setting is no longer satisfactory, we expect to reestablish good performance with just a few changes. Here the code from Section III can be used to generally monitor the network, keeping the network operator informed whether he can improve network performance with a few changes. The current running times of one–two hours may be considered too slow for such on-line traffic engineering. However, as is typical for local search, we can always make a less exhaustive search to terminate earlier. Within 10–15 minutes, one often get within 5% of the presented results. Also, there are much faster machines on the market. Anyhow, our system is expected to be most interesting to network operators in cases where real problems are experienced in the network, and where the alternative is to buy new hardware. In such a case, one–two hours of CPU time is cheap.

Besides the above concrete results, we believe that we have set a quite general framework for efficient traffic management, and many of the objectives could be reused within MPLS routing. Particularly this covers the convexity idea.

In combination, our techniques provide an efficient system for OSPF/IS-IS traffic management in a changing world. It is now an integrated part of AT&Ts NetScope/Bravo traffic engineering tool for IP networks [16].

ACKNOWLEDGMENT

The authors would like to thank J. Rexford for some very useful and patient comments to different drafts of this paper.

REFERENCES

- [1] E. H. L. Aarts and J. K. Lenstra, Eds., "Local search in combinatorial optimization," in *Discrete Mathematics and Optimization*. Chichester, England: Wiley, 1997.
- [2] W. Aiello, F. Chung, and L. Lu, "A random model for massive graphs," in *Proc. 32nd ACM Symp. Theory Computing (STOC)*, 2000, pp. 171–179.
- [3] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao. (2000) A framework for internet traffic engineering. Network Working Group, Internet Draft (work in progress). [Online]. Available: <http://search.ietf.org/internet-drafts/draft-ietf-tewg-framework-02.txt>.
- [4] D. O. Awduche, "MPLS and traffic engineering in IP networks," *IEEE Communications Magazine*, vol. 37, pp. 42–47, Dec. 1999.
- [5] D. O. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus. (1999, Sept.) Requirements for traffic engineering over MPLS. Network Working Group, Request for Comments. [Online]. Available: <http://search.ietf.org/rfc/rfc2702.txt>.
- [6] A. Bley, M. Grötchel, and R. Wessälly, "Design of broadband virtual private networks: Model and heuristics for the B-WiN," in *Proc. DIMPCS Workshop on Robust Communication Networks and Survivability AMS-DIMACS Series 53*, 1998, pp. 1–16.
- [7] Caida. [Online]. Available: <http://www.caida.org/tools/visualization/mapnet/Data/>.
- [8] R. Callon. (1990, Dec.) Use of OSI IS-IS for routing in TCP/IP and dual environments. Network Working Group, Request for Comments: Rep. 1195. [Online]. Available: <http://search.ietf.org/rfc/rfc1195.txt>.
- [9] K. Calvert, M. Doar, and E. W. Zegura, "Modeling Internet topology," *IEEE Commun. Mag.*, vol. 35, pp. 160–163, 1997.
- [10] J. Cao, D. Davis, S. V. Wiel, and B. Yu, "Time-varying network tomography: Router link data," *J. Amer. Statist. Assoc.*, vol. 95, pp. 1063–1075, 2000.
- [11] Cisco. (1997) Configuring OSPF. [Online]. Available: Documentation at http://www.cisco.com/uni-verc/cc/td/doc/product/software/ios113ed/113ed_cr/np1_c/1cospf.htm.

- [12] N. G. Duffield and M. Grossglauser, "Trajectory sampling for direct traffic observation," *ACH/IEEE Trans. Networking*, vol. 9, pp. 280–292.
- [13] A. Dwivedi and R. E. Wagner, "Traffic model for USA long-distance optimal network," in *Proc. Optical Fiber Communication Conf. (OFC)*, vol. TuK1, 2000, pp. 156–158.
- [14] M. Ericsson, M. G. C. Resende, and P. M. Pardalos, "A genetic algorithm for the weight setting problem in OSPF routing," *J. Combinatorial Optimization*, 2002, to be published.
- [15] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proc. ACM SIGCOMM'99: Conf. Applications, Technologies, Architectures, and Protocols for Computer Communications*, 1999, pp. 251–262.
- [16] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "NetScope: Traffic engineering for IP networks," *IEEE Network Mag. (Special Issue on Internet Traffic Eng.)*, vol. 14, no. 2, pp. 11–19, Mar./Apr. 2000.
- [17] 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 Trans. Networking*, pp. 265–279, 2001.
- [18] A. Feldmann, J. Rexford, and R. Caceres, "Efficient policies for carrying web traffic over flow-switched networks," *IEEE/ACM Trans. Networking*, vol. 6, pp. 673–685, 1998.
- [19] B. Fortz and M. Thorup. (2000) Increasing Internet capacity using local search. Université Libre de Bruxelles, Tech. Rep. IS-MG 2000/21. [Online]. Available: http://smg.ulb.ac.be/Preprints/Fortz00_21.html.
- [20] —, "Internet traffic engineering by optimizing OSPF weights," in *Proc. 19th IEEE Conf. Computer Communications (INFOCOM)*, 2000, pp. 519–528.
- [21] F. Y. S. Lin and J. L. Wang, "Minimax open shortest path first routing algorithms in networks supporting the SMDS services," in *Proc. IEEE Int. Conf. Communications (ICC)*, vol. 2, 1993, pp. 666–670.
- [22] D. Mitra and K. G. Ramakrishnan, "A case study of multiservice, multipriority traffic engineering design for data networks," in *Proc. IEEE GLOBECOM*, 1999, pp. 1077–1083.
- [23] J. T. Moy. (1991, July) OSPF version 2. Network Working Group, Request for Comments: 1247. [Online]. Available: <http://search.ietf.org/rfc/rfc1247.txt>.
- [24] K. G. Ramakrishnan and M. A. Rodrigues, "Optimal routing in shortest-path data networks," *Lucent Bell Labs Technical Journal*, vol. 6, no. 1, 2001.
- [25] E. C. Rosen, A. Viswanathan, and R. Callon. (2001) Multiprotocol label switching architecture. Network Working Group, Request for Comments. [Online]. Available: <http://search.ietf.org/rfc/rfc3031.txt>.
- [26] A. Sridharan, S. Bhattacharyya, C. Diot, R. Guerin, and J. Jetcheva, "On the impact of aggregation on the performance of traffic aware routing," [Online]. Available: <http://www.seas.upenn.edu:8080/guerin/>, 2000, submitted for publication.
- [27] P. L. Toint, "Transportation modeling and operations research: A fruitful connection," in *Operations Research and Decision Aid Methodologies in Traffic and Transportation Management*. ser. NATO ASI Series: Ser. F, Computer and Systems Sciences, M. Labbé, G. Laporte, K. Tanczos, and P. L. Toint, Eds. New York: Springer-Verlag, 1998, vol. 166, pp. 1–27.
- [28] B. M. Waxman, "Routing of multipoint connections," *IEEE J. Select. Areas Commun. (Special Issue on Broadband Packet Communications)*, vol. 6, pp. 1617–1622, 1988.
- [29] E. W. Zegura. (1996) GT-ITM: Georgia Tech internetwork topology models (software). [Online]. Available: <http://www.cc.gatech.edu/fac/Ellen.Zegura/gt-itm/gt-itm.tar.gz>.
- [30] E. W. Zegura, K. L. Calvert, and S. Bhattacharjee, "How to model an internetwork," in *Proc. 15th IEEE Conf. Computer Communications (INFOCOM)*, 1996, pp. 594–602.



Bernard Fortz (M'99–A'00) received the Ph.D. degree in operations research from the Université Libre de Bruxelles, Belgium.

He is a Professor of Operations Research at the Institut d'Administration et de Gestion of the Université Catholique de Louvain, Belgium. His main research interests are combinatorial optimization, network design problems, and the optimization of Internet resources using efficient routing protocols.

Dr. Fortz was awarded the AT&T research prize, in 1997.



Mikkel Thorup received the Ph.D. degree from Oxford University, Oxford, U.K., in 1993.

From 1993 to 1998, he was with the University of Copenhagen, Denmark. In the summer of 1997, he was a Distinguished Visiting Professor at the Max-Planck-Institut für Informatik, Germany. He is currently a Principal Technical Staff Member at AT&T Labs-Research, Florham Park, NJ, where he has been since 1998. His main interest is algorithms and data structures, and he serves on the Editorial Board of *Journal of Algorithms*.