Optim Eng (2008) 9: 257–271
DOI 10.1007/s11081-007-9033-z

Optimal oblivious routing under linear and ellipsoidal uncertainty

Pietro Belotti · Mustafa Ç. Pınar

Received: 23 October 2007 / Accepted: 20 November 2007 / Published online: 1 December 2007

Abstract In telecommunication networks, a common measure is the maximum congestion (i.e., utilization) on edge capacity. As traffic demands are often known with a degree of uncertainty, network management techniques must take into account traffic variability. The *oblivious performance* of a routing is a measure of how congested the network may get, in the worst case, for one of a set of possible traffic demands.

We present two models to compute, in polynomial time, the optimal oblivious routing: a linear model to deal with demands bounded by box constraints, and a second-order conic program to deal with ellipsoidal uncertainty, i.e., when a mean-variance description of the traffic demand is given. A comparison between the optimal oblivious routing and the well-known OSPF routing technique on a set of real-world networks shows that, for different levels of uncertainty, optimal oblivious routing has a substantially better performance than OSPF routing.

Keywords Traffic engineering · Oblivious routing · Linear programming · Second order cone programming

1 Introduction

Telecommunication networks are an important infrastructure of today's economy; the cost of connecting a community through wired or wireless technologies is paid off

Research partially supported by Bilateral Grant MISAG-CNR-1, jointly from the Scientific and Technological Research Council of Turkey and the Consiglio Nazionale delle Ricerche, Italy.

P. Belotti (✉)
Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA
e-mail: belotti@andrew.cmu.edu

M.Ç. Pınar
Department of Industrial Engineering, Bilkent University, Ankara, Turkey
e-mail: mustafap@bilkent.edu.tr

by the benefits offered by rapid data transfer. However, the variety of technologies available and the complexity of such invasive structure pose difficult problems to operators, both in the design and management of networks.

We focus on a class of problems where the link capacity is known in advance, and a set of origin-destination requests of flow, called *traffic demand* (or simply *demand*), is given. Then one faces the problem of *routing* these requests of flow such that the network capacity is not overloaded. Techniques of *traffic engineering* allow for a routing that does not affect the network performance (delay, data loss), thus guaranteeing a certain quality of service.

A common measure of the network usage is the maximum link *congestion* for a given routing, i.e., the maximum percentage of used link capacity. The more congested a network, the more prone to instability it is in the event of a change in the traffic requests. It is desirable then to devise a routing that gives the minimum congestion for a given demand. As shown in the next section, this amounts to solving a linear programming (LP) problem.

However, the traffic demand d is seldom known with accuracy, e.g. for difficulties in measurement or because d varies in time, and a set $\mathcal{D}$ of possible demands is considered. It is useful then to compute the *oblivious performance ratio* of a routing, i.e., the maximum ratio, for all $d \in \mathcal{D}$, between the congestion of the routing and the minimum congestion that can be attained for d.

We present two models that obtain, in polynomial time, the optimal oblivious routing assuming that either the demand has lower and upper bounds, or is described by mean-covariance information. These uncertainty models are motivated by the presence in the literature of techniques to estimate d with some accuracy, expressed by box constraints or mean-variance information (Tebaldi and West 1998; Vardi 1996; Zhang et al. 2003).

Often, real-world data networks follow the Open, Shortest Path First (OSPF) policy: routes are chosen as shortest paths between origin and destination, where arc weights are chosen depending on network parameters such as edge capacity. As arc weights are the only degree of freedom to play with, routing optimization consists in finding the value of weights so as to minimize some network performance measure (Ericsson et al. 2002; Fortz and Thorup 2000; Lin and Wang 1993). Other routing techniques such as *Multi-Protocol Label Switching* (MPLS) do not constrain route length, thus allowing for the implementation of any routing. As shown in (Fortz and Thorup 2000), this greater flexibility pays off in terms of network performance. The second contribution of this work is a comparison between the performance of a finely tuned routing and that of OSPF routing as commonly implemented in today's networks, that shows that the oblivious performance ratio of OSPF routing can be greatly improved.

In the next section we present the routing problem and some additional notation. The concept of oblivious routing is introduced in Sect. 3, and the two models are presented in Sects. 4 and 5. We report some tests on real-world networks in Sect. 6 and give some conclusions in Sect. 7.

2 Routing and congestion in telecommunication networks

Consider a network topology defined by an undirected graph $G = (V, E)$ whose edges $e \in E$ are assigned a capacity c_e. An edge e may also be denoted by the set $\{h, k\}$ of its endnodes. Associated with E is the set of directed arcs A containing all pairs (h, k) and (k, h) such that $\{h, k\} \in E$. The *neighborhood* of node h, i.e., the set of nodes adjacent to h, is defined as $N(h) = \{k \in V : \{h, k\} \in E\}$.

An *origin-destination pair* (o-d pair) is an oriented pair (i, j) of nodes in V requesting an amount of flow d_{ij} to be sent from i to j. Let D be the set of all o-d pairs. A *traffic demand* $d = (d_{ij})$ is a vector of requests between all $(i, j) \in D$.

The fraction of demand (i, j) flowing on edge $\{h, k\}$ in the direction $h \to k$ is denoted as f^{ij}_{hk}; for the sake of readability, we denote with f_{hk} the vector (f^{ij}_{hk}), $(i, j) \in D$, and with f_e, where $e = \{h, k\} \in E$, the vector $f_{hk} + f_{kh}$. Finally, vector f denotes the vector (f^{ij}_{hk}), $(i, j) \in D$, $(h, k) \in A$.

We denote the total flow on edge e as $\text{FLOW}(e, f, d) = \sum_{(i,j)\in D} d_{ij}(f^{ij}_{hk} + f^{ij}_{kh}) = d^T(f_{hk} + f_{kh}) = d^T f_e$. A *routing* of a demand vector d is the set of all f^{ij}_{hk} for each o-d pair $(i, j) \in D$ and arc $(h, k) \in A$ satisfying flow conservation:

$$\sum_{k\in N(h)} (f^{ij}_{hk} - f^{ij}_{kh}) = \begin{cases} 1 & \text{if } h = i, \\ -1 & \text{if } h = j, \\ 0 & \text{otherwise} \end{cases} \qquad \forall h \in V, (i, j) \in D.$$

A routing is *feasible* for a demand vector d if the network capacity can support it, i.e., $\text{FLOW}(e, f, d) \le c_e$ for all $e \in E$; a demand d is *feasible* when there exists a feasible f for d. The *congestion* of a network is the maximum fraction of capacity used on the graph edges:

$$\text{CONG}(f, d) = \max_{e\in E} (\text{FLOW}(e, f, d)/c_e).$$

Let us denote as $\mathcal{F}$ the set of all feasible routings. If the demand d is known a priori, then the routing with the minimum congestion ratio, $\text{OPT}(d) = \min_{f\in\mathcal{F}} \text{CONG}(f, d)$, is computed by solving the following linear problem:

$$\text{OPT}(d) = \min\ z \tag{1}$$

$$\text{s.t. } z \ge \sum_{(i,j)\in D} (g^{ij}_{hk} + g^{ij}_{kh})/c_e \quad \forall e = \{h, k\} \in E, \tag{2}$$

$$\sum_{k\in N(h)} (g^{ij}_{hk} - g^{ij}_{kh}) = \begin{cases} d_{ij} & \text{if } h = i \\ -d_{ij} & \text{if } h = j \\ 0 & \text{otherw.} \end{cases} \qquad \forall h \in V, (i, j) \in D, \tag{3}$$

$$\sum_{(i,j)\in D} (g^{ij}_{hk} + g^{ij}_{kh}) \le c_e \quad \forall e = \{h, k\} \in E, \tag{4}$$

$$g \ge \mathbf{0}, \tag{5}$$

where, unlike variables f, variables g represent a flow rather than a fraction of demand. Although constraint (4) is not necessary here due to the minimization of the bottleneck variable z, it is used in the following sections. Notice that, in view of constraint (4), it follows that $z \leq 1$.

3 Demand uncertainty and oblivious routing

From now on, we assume that the traffic demand d is not known a priori but can be any of a set $\mathcal{D}$ (non-empty and bounded) of traffic demands. Not surprisingly, the problem gets more difficult as *robust* network management is needed. The problem of routing a set of demands under uncertainty has received much attention recently. Li et al. (2004) deal with the *multicast* case, where a demand has one source but multiple destinations. Lin and Wang (1993) present a Lagrangian Relaxation-based algorithm for the problem where demands are routed on single paths, while Roughan et al. (2003) propose a simulation approach to solve both the estimation and the routing problem.

Given a routing f and a set $\mathcal{D}$ of possible traffic demands, the congestion ratio of f can be defined as the worst-case congestion ratio within $\mathcal{D}$, i.e., $\max_{d\in\mathcal{D}} \text{CONG}(f, d)$. However, a realistic measure should be independent from $\mathcal{D}$ and take into account only demands that can be routed. Let us denote $H(\mathcal{D})$ as the set of feasible demands. Hence the LP problem to compute $\text{OPT}(d)$ is feasible for all $d \in H(\mathcal{D})$. The *oblivious performance ratio* is the worst-case ratio, over all demands in $H(\mathcal{D})$, of $\text{CONG}(f, d)$ to the minimum congestion for d, $\text{OPT}(d)$:

$$\text{OPR}(f, \mathcal{D}) = \max_{d\in H(\mathcal{D})} \frac{\text{CONG}(f, d)}{\text{OPT}(d)}$$

is a measure of the redundancy of f with respect to the demand uncertainty $\mathcal{D}$. In this work, we consider two uncertainty representations, i.e., two specific subsets of all feasible demands, and propose a polynomial-time method for the *optimal oblivious routing* problem, which consists in finding, for a given set of o-d pairs, the routing with the optimal oblivious performance ratio:

$$\text{OOPR}(\mathcal{D}) = \min_{f\in\mathcal{F}} \max_{d\in H(\mathcal{D})} \frac{\max_{e\in E} \text{FLOW}(e, f, d)/c_e}{\text{OPT}(d)}. \tag{6}$$

This problem has been studied previously by Azar et al. (2003) and Applegate and Cohen (2003). Both works focus on a very general setting: the set $\mathcal{D}$ contains all demands admitting feasible routing on G. Azar et al. (2003) present an LP model with a class of constraints whose cardinality is exponential, and is thus dealt with as a family of valid inequalities whose separation has polynomial complexity. The authors then describe a cutting plane procedure that finds the oblivious routing in polynomial time. Applegate and Cohen (2003) propose a polynomial size LP model where the separation problem of the former work is solved implicitly. Assuming that $\mathcal{D}$ contains all feasible demands implies that, for the worst-case demand $\bar{d} = \text{argmax}_{d\in H(\mathcal{D})} \frac{\text{CONG}(f,d)}{\text{OPT}(d)}$, the capacity of at least one edge is totally used, i.e., $\text{OPT}(\bar{d}) = 1$. Thus, the optimal oblivious performance ratio (6) reduces to

$$\min_{f\in\mathcal{F}} \max_{d\in H(\mathcal{D})} \max_{e\in E} (\text{FLOW}(e, f, d)/c_e);$$

then, by swapping the two *max* operators and by strong duality, a linear programming model is obtained.

The assumption $\text{OPT}(d) = 1$ is no longer valid if the set of demands is further limited, e.g., by box constraints or within a mean-variance region. In fact, suppose that $\mathcal{D}$ is the set of demands admitting a routing in G and such that $d_{ij} \le \alpha \frac{\min_{e\in E} c_e}{|D|}$ for all $(i, j) \in D$, with $\alpha < 1$. For all demands $d \in \mathcal{D}$, there is a routing f such that even if all demands were routed on edge $\bar{e}$ with minimum capacity, $\text{FLOW}(\bar{e}, f, d) \le \alpha c_{\bar{e}}$ and hence $\text{OPT}(d) \le \alpha < 1$ for all $d \in \mathcal{D}$.

We observe that $\text{OPT}(d)$ does not depend on e, hence:

$$\text{OOPR}(\mathcal{D}) = \min_{f\in\mathcal{F}} \max_{e\in E} \max_{d\in H(\mathcal{D})} \frac{\text{FLOW}(e, f, d)/c_e}{\text{OPT}(d)}$$

(notice that we have swapped the *max* operators). The model is as follows:

$$\begin{aligned} &\min\ r \\ &\text{s.t.}\ \ r \ge \max_{d\in H(\mathcal{D})} \frac{\text{FLOW}(e, f, d)/c_e}{\text{OPT}(d)} \quad \forall e \in E, \\ &\qquad f \ \text{ is a routing.} \end{aligned} \tag{7}$$

Notice that constraint (7) can be written as

$$\max_{d\in H(\mathcal{D})} (\text{FLOW}(e, f, d) - r c_e \text{OPT}(d)) \le 0 \quad \forall e \in E. \tag{8}$$

4 A model with lower and upper bounds on demands

Suppose that vectors $a = (a_{ij})$ and $b = (b_{ij})$, $(i, j) \in D$, are given, and that $\mathcal{D}$ is the set of all feasible demands d such that $a \le d \le b$. For each edge $e \in E$, the left-hand side of (8) is the solution to an optimization problem over variables d supposing that f and r are fixed. We impose that d is a feasible demand by introducing auxiliary flow variables g. Let us write flow conservation constraints (3) in matricial form $A_1 g = d$ and $A_2 g = \mathbf{0}$; analogously, we use $Bg \le c$ instead of (4). As $\text{FLOW}(e, f, d) = d^T f_e$, the left-hand side of (8) is the following LP problem in variables g, d and ω, while f_e and r are taken as parameters:

$$\max\ (d^T f_e - r c_e \omega) \tag{9}$$

$$\text{s.t.}\ \ (\pi_e) \quad A_1 g = d, \tag{10}$$

$$(\sigma_e) \quad A_2 g = \mathbf{0}, \tag{11}$$

$$(\eta_e) \quad Bg \le c\,\omega, \tag{12}$$

$$(\chi_e) \quad \omega \le 1, \tag{13}$$

$$(\lambda_e) \quad -d \le -a, \tag{14}$$

$$(\mu_e) \quad d \le b, \tag{15}$$

$$(g, d, \omega) \ge \mathbf{0}. \tag{16}$$

We assume feasibility here, i.e., there exists at least one (g, d, ω) such that (10–16) hold. This is a rather general assumption, since infeasibility of (10–16) would imply that no demand in $\mathcal{D}$ admits a routing. Boundedness of $\mathcal{D}$ also implies that this LP is bounded. Furthermore, in any optimal solution to the above problem, we have $\omega = \text{OPT}(d)$: as any $d \in H(\mathcal{D})$ admits a routing g and ω has a negative coefficient in the objective function, any optimal value of ω is also optimal to (1–5). Constraint $\omega \leq 1$ requires that the network capacity support flow g, thus excluding infeasible demands of $\mathcal{D}$. To prove that the left-hand side of (8) equals an optimal solution of (9–16), we observe that:

$$\begin{aligned}
&\max_{d\in H(\mathcal{D})}\{d^T f_e - rc_e \text{OPT}(d)\} \\
&\quad = \max_{d\in H(\mathcal{D})}\left\{d^T f_e - rc_e \min_{g\in\mathcal{F},\omega\geq 0}\{\omega : (10), (11), (12)\}\right\} \\
&\quad = \max_{d\in H(\mathcal{D})}\left\{d^T f_e + \max_{g\in\mathcal{F},\omega\geq 0}\{-rc_e\omega : (10), (11), (12)\}\right\} \\
&\quad = \max_{d\in H(\mathcal{D}),g\in\mathcal{F},\omega\geq 0}\{d^T f_e - rc_e\omega : (10), (11), (12)\},
\end{aligned}$$

and that $d \in H(\mathcal{D})$ corresponds to constraints (13–15).

On the left of each constraint in (10–16) we give the corresponding dual variables. The dual is the following minimization problem:

$$\begin{aligned}
\min\ & \chi_e - a\lambda_e + b\mu_e \\
\text{s.t.}\ & \pi_e^T A_1 + \sigma_e^T A_2 + \eta_e^T B \geq 0, && (17) \\
& -\pi_e - \lambda_e + \mu_e \geq f_e, && (18) \\
& -c\eta_e + \chi_e \geq -rc_e, && (19) \\
& (\chi_e, \eta_e, \lambda_e, \mu_e) \geq \mathbf{0}. && (20)
\end{aligned}$$

Therefore, for each edge $e \in E$ we solve the dual of a maximization problem that gives the left-hand side of (8). The result below gives an LP model, which we call MB, to compute in polynomial time $\text{OOPR}(\mathcal{D})$, where $\mathcal{D}$ is the set of demands d with box constraints $a \leq d \leq b$.

Proposition 1 *The optimal oblivious routing with box constraints on d is obtained by solving the following linear problem*:

$$\begin{aligned}
(\text{MB}) \quad & \min r \\
& A_1 f = \mathbf{1}, \\
& A_2 f = \mathbf{0}, \\
& \chi_e - a\lambda_e + b\mu_e \leq 0 \quad \forall e \in E, \\
& \pi_e^T A_1 + \sigma_e^T A_2 + \eta_e^T B \geq \mathbf{0} \quad \forall e \in E, && (21) \\
& -\pi_e - \lambda_e + \mu_e \geq f_e \quad \forall e \in E, && (22) \\
& -c\eta_e + \chi_e \geq -rc_e \quad \forall e \in E, && (23) \\
& (r, f, \chi, \eta, \lambda, \mu) \geq \mathbf{0}. && (24)
\end{aligned}$$

Proof Consider the problem (P1):

$$\text{(P1)} \quad \min_{r,f}\left\{r : A_1 f = \mathbf{1}, A_2 f = \mathbf{0}, \max_{(d,g,\omega)\in X}(d^T f_e - rc_e\omega) \leq 0\ \forall e \in E\right\}$$

where $X = \{(d, g, \omega) \geq 0 : A_1 g = d, A_2 g = 0, Bg \leq c\omega, a \leq d \leq b, \omega \leq 1\}$. Due to strong duality, this problem is equivalent to (D1):

$$\text{(D1)} \quad \min_{r,f}\left\{r : A_1 f = \mathbf{1}, A_2 f = \mathbf{0}, \right.$$
$$\left. \min_{(\pi_e,\sigma_e,\eta_e,\chi_e,\lambda_e,\mu_e)\in Y_r(e)}(\chi_e - a\lambda_e + b\mu_e) \leq 0\ \forall e \in E\right\},$$

where $Y_r(e) = \{(\pi_e, \sigma_e, \eta_e, \chi_e, \lambda_e, \mu_e) : (21), (22), (23), (24)\}$. We can remove the *min* in problem (D1) as a result of the following observation. Consider the problem without the *min* operator, (D2):

$$\begin{aligned} \text{(D2)} \quad \min_{r,f}\{r : {} & A_1 f = \mathbf{1}, A_2 f = \mathbf{0}, \\ & (\chi_e - a\lambda_e + b\mu_e) \leq 0\ \forall e \in E, \\ & (\pi_e, \sigma_e, \eta_e, \chi_e, \lambda_e, \mu_e) \in Y_r(e)\ \forall e \in E\}. \end{aligned}$$

Obviously, any feasible solution to (D1) leads to a feasible point in (D2). Take now the feasible set of (D1) which is non-empty if and only if the feasible set of (D2) is non-empty. Let us fix an arbitrary $e \in E$. For fixed r and f feasible for (D2), let $Y^*_{r,f,e} = \{(\pi_e, \sigma_e, \eta_e, \chi_e, \lambda_e, \mu_e) \in Y_r(e) : (\chi_e - a\lambda_e + b\mu_e) \leq 0\}$. Then

$$\min_{(\pi_e,\sigma_e,\eta_e,\chi_e,\lambda_e,\mu_e)\in Y^*_{r,f,e}}(\chi_e - a\lambda_e + b\mu_e) \leq 0.$$

We have thus constructed a feasible solution to (D1). Hence, (D1) and (D2) must be equivalent, and (P1) is equivalent to (D2). □

5 Modeling ellipsoidal uncertainty

Suppose that d has a probability distribution (not necessarily known) with mean $\bar{d}$ and covariance matrix P. In this case, we can represent the uncertainty set conveniently as an ellipsoid, i.e., $d \in \{\bar{d} + Pu : \|u\|_2 \leq \epsilon\}$, with P positive (semi)definite as in Ben Tal and Nemirovski (1999) and $\epsilon \geq 0$ and finite. Hence, $\mathcal{D}$ is bounded and nonempty. Parameter ϵ determines the subset of demands that must be taken into account in the optimization – the greater ϵ, the more conservative the model. If we replace variable d with $\bar{d} + Pu$, the left-hand side of (8) is an optimum of the following maximization problem similar to (9–16) (analogously, we assume feasibility here):

$$\begin{aligned} & \max\ (\bar{d} + Pu)^T f_e - rc_e\omega \\ & \text{s.t.}\ (\pi_e) \quad A_1 g = d, \\ & \qquad (\sigma_e) \quad A_2 g = 0, \end{aligned}$$

$$\begin{aligned} (\eta_e) \quad & Bg \le c\omega, \\ & \omega \le 1, \qquad d = \bar{d} + Pu, \qquad \|u\|_2 \le \epsilon, \\ & (g, d, \omega) \ge \mathbf{0}, \end{aligned}$$

where we have associated multipliers π_e, σ_e, and η_e to the first three constraints. Consider its Lagrangian dual problem:

$$\begin{aligned} \min_{\pi_e,\sigma_e,\eta_e \ge 0} \max_{(u,\omega,g)\in S} & \{ f_e^T(\bar{d} + Pu) - rc_e\omega + \pi_e^T(\bar{d} + Pu - A_1 g) \\ & + \sigma_e^T(-A_2 g) + \eta_e^T(c\omega - Bg)\}, \end{aligned} \tag{25}$$

where $S = \{(u, \omega, g) : \|u\|_2 \le \epsilon, 0 \le \omega \le 1, g \ge 0\}$. The inner problem is decomposed in the sum of $f_e^T \bar{d} + \pi_e^T \bar{d} = (f_e + \pi_e)^T \bar{d}$ plus the following three maximization problems:

(a) $\max_{u:\|u\|_2 \le \epsilon} (f_e + \pi_e)^T Pu = \epsilon\|P(f_e + \pi_e)\|_2$;
(b) $\max_{0\le\omega\le 1}(\eta_e^T c - rc_e)\omega = \max(0, \eta_e^T c - rc_e)$ (we denote it as $[\eta_e^T c - rc_e]_+$);
(c) $\max_{g\ge 0}(-\pi_e^T A_1 - \sigma_e^T A_2 - \eta_e^T B)g$: this problem has null solution if and only if $\pi_e^T A_1 + \sigma_e^T A_2 + \eta_e^T B \ge 0$.

Thus, the problem (25) has the solution $(f_e + \pi_e)^T \bar{d} + \epsilon\|P(f_e + \pi_e)\|_2 + [\eta_e^T c - rc_e]_+$ if and only if the condition $\pi_e^T A_1 + \sigma_e^T A_2 + \eta_e^T B \ge 0$ holds.

Proposition 2 *The model MS for optimal oblivious routing with uncertainty expressed by mean-variance is the following*:

$$\begin{aligned} (\text{MS}) \quad & \min r \\ & A_1 f = \mathbf{1}, \\ & A_2 f = \mathbf{0}, \\ (f_e - \pi_e)^T \bar{d} + \epsilon\|P(f_e - \pi_e)\|_2 + \xi_e \le 0 \quad & \forall e \in E, \qquad (26) \\ \xi_e \ge \eta_e^T c - rc_e \quad & \forall e \in E, \qquad (27) \\ \pi_e^T A_1 + \sigma_e^T A_2 + \eta_e^T B \ge \mathbf{0} \quad & \forall e \in E, \qquad (28) \\ & (r, f, \eta, \xi) \ge \mathbf{0}. \end{aligned}$$

The proof is similar to the proof of Proposition 1 and, thus, is omitted. This model has the non-linear but convex, second-order cone constraint (26) and hence is solvable in polynomial time through interior point SOCP solvers. Notice that constraints (27) and (28) are equivalent to (19) and (17).

6 Computational results

We have adopted a test bed of four network instances available from the Rocketfuel project (Springs et al. 2004), providing data for the topology (V and E), link counts, and OSPF weights w_e of several real-world networks. We have also used an example instance (*Nsf*) from a work by Mitra and Ramakrishnan (1999), with demand and capacity data. We assume that weights follow Cisco's policy: a link e is assigned an OSPF weight w_e equal to the inverse of its capacity. Hence, we simply assume $c_e = 1/w_e$.

Traffic demand data, regarded as proprietary information by Internet Service Providers (ISPs), is rarely disclosed. Therefore, we have created the traffic demand under the *gravity* model: the demand is assumed proportional to a *repulsion* and an *attraction* parameter, R_i and A_i, associated with each node i, which in turn are proportional to the number of data packets exiting and entering node i, respectively. In order to test our model on a reasonable demand, we use a scalar factor β such that the demands $d_{ij} = \beta R_i A_j$ are feasible. Assume $\beta = \gamma \max\{u : (3), (4), (5), d_{ij} = u R_i A_j\}$ for a given $\gamma \in [0, 1]$. Hence, $(\beta R_i A_j)$ is a feasible traffic demand with congestion at most equal to γ.

We have tested our instances using different values of the uncertainty parameters. The scalar γ has been assigned values in the set $\{0.75, 0.95, 0.99\}$, so as to give $\bar{d}$ increasingly critical values. In model MB, the lower and upper bounds are $a = \bar{d}/p$ and $b = \bar{d}p$, where p has been assigned values in the set $\{1.2, 2, 5, 20\}$. In model MS, the covariance matrix P is a positive semidefinite, randomly generated matrix. The parameter ϵ is set to $\zeta \|\bar{d}\|_2$, where $\zeta \in \{0.01, 0.05, 0.1, 0.2, 0.4, 0.8\}$. It is worth noticing that large values of p and ζ correspond to a greater degree of uncertainty, hence we expect the oblivious performance ratio to grow as p and ζ grow.

The size of all instances tested could be reduced by neglecting those nodes in V with degree one, as the routing to and from such nodes is trivial. More precisely, if the removal of an edge $e \in E$ divides the graph in two components G_1 and G_2 such that G_1 is a tree, then $G_1 = (V_1, E_1)$ can be shrunk into a supernode i whose demands d_{ij} (resp. d_{ji}) for all nodes $j \notin V_1$ are given by $\sum_{h \in V_1} d_{hj}$ (resp. $\sum_{h \in V_1} d_{jh}$), while all *internal* demands d_{hk} with both h and k in V_1 can be ignored as they are routed within G_1. A polynomial procedure to reduce the graph consists in repeatedly shrinking all nodes i with only one neighbor j to j itself, until no such nodes i are found. It is worth noting that the flow on edge $e = \{i, j\}$ is fixed, hence $\text{FLOW}(e, f, d)/c_e \leq \text{OPT}(d)$; as $\text{OOPR}(\mathcal{D}) \geq 1 \geq \frac{\text{FLOW}(e,f,d)/c_e}{\text{OPT}(d)}$, edge e can be ignored. As appears from columns 2 to 5 in Tables 1 and 2, this reduces the size of almost all instances, which could be solved to optimality in reasonable time.

We have tested the MB model on a Sun Fire 240 workstation equipped with a 1.66 GHz Sparc64 processor and 4 GB RAM; the MS model instead has been tested on a computer with a 1.5 MHz Pentium processor and 512 MB of RAM memory. Both models for optimal oblivious routing have been coded in AMPL (Fourer et al. 1990); model MB has been solved by the linear programming solvers of CPLEX 9.0 (Ilog Inc 2003) (we have chosen to use the barrier instead of the simplex method due to a substantial improvement in solution time), whereas the SOCP model MS has been solved through the interior point method in the MOSEK 3.1 software package (Andersen and Andersen 2000). The source and the data files used in our tests are available from the ftp page:

ftp://ftp.elet.polimi.it/users/Pietro.Belotti/oblivious

Table 1 shows the results obtained with the MB model. Columns 4 and 5 give the size of the instance after the reduction above described, then for each value of γ and p we report the optimal oblivious performance ratio "oopr" and the performance ratio "ospf" obtained by OSPF routing. We obtain "ospf" by simply computing, for each pair (s,t), the shortest path from s to t according to the OSPF weights $w_e = 1/c_e$, and then fixing the flow variables f in MB accordingly. The last column reports the computing time required, on average, to solve the LP problem associated with MB.

Analogously, Table 2 shows the results obtained with the MS model. For each value of γ and ζ we report the optimal oblivious performance ratio "oopr" and the performance ratio "ospf" obtained by OSPF routing, which is obtained similarly as for MB. Due to its size, instance *Sprintlink* of model MS could not fit into the RAM memory and hence has not been solved.

It is apparent from the tables that, in all cases, OSPF routing has an oblivious performance ratio that is much worse than the optimal oblivious routing, computed through our models. With low degrees of uncertainty, while OSPF routing has a sensible performance loss (from 40% for *Sprintlink* network, under the MB model, to 151% for *Abovenet*, under the MS model), the optimal oblivious performance ratio is one in most cases, as is expected since, for $p \to 1$ or $\zeta \to 0$, the optimal routing is the one obtained with model (1–5). Nevertheless, $\text{OPR}(\mathcal{D}) = 1$ even for larger p and ζ, i.e., even greater degrees of uncertainty do not affect the routing performance.

As p and ζ get large values, OSPF routing has a performance ratio of up to 12 as in instance *Sprintlink*, whereas the best oblivious routing does not worsen significantly, as the performance loss is not greater than 99%, indicating high robustness of the optimal oblivious routing; notice that $p = 20$ and $\zeta = 0.8$ give a large percentage of the feasible demands. We also observe that the performance ratio has almost no dependence on γ, which can be explained by the fact that γ specifies how critical the demand is w.r.t. the network capacity, but it does not drive the level of uncertainty, which is specified by p and ζ.

We have depicted the dramatic gain in performance in Fig. 1 for network *Nsf*, under the MS model, for $\gamma = 0.95$ and for ζ varying in the interval $[10^{-5}, 5]$. It is worth emphasizing that the optimal performance ratio is 1 for small and medium values of ζ, whereas the OSPF routing has a performance ratio of 1.6, i.e., a loss of 60%, even for very low degrees of uncertainty. For higher degrees of uncertainty, the OSPF routing attains a performance ratio of 4 while the optimal oblivious ratio stabilizes at 1.818. This shows that a finely tuned routing, at least for low degrees of uncertainty, may have a performance ratio which is the best possible, and does not increase significantly even with high uncertainty.

The optimization time is reasonably short for all instances except *Sprintlink*, that has a size almost double as that of the remaining ones and has required greater memory and processor resources. For larger networks it could be necessary to study an alternative approach, e.g., a column generation technique based on a path formulation of the problem.

Table 1 Oblivious performance ratio for OSPF and the optimal oblivious routing with box uncertainty

Name	Orig.		Redu.		p	$\gamma = 0.75$		$\gamma = 0.90$		$\gamma = 0.99$		t_{avg} [s]
	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$		oopr	ospf	oopr	ospf	oopr	ospf	
Telstra (AU)	44	46	7	9	1.2	1.0000	1.9852	1.0000	1.9739	1.0000	1.9694	0.01
					2.0	1.0000	2.0303	1.0000	2.0167	1.0000	2.0140	
					5.0	1.2343	2.0607	1.2343	2.0553	1.2343	2.0542	
					20.0	1.2827	2.0759	1.2827	2.0746	1.2827	2.0743	
VNSL (India)	9	11	8	10	1.2	1.0000	2.1163	1.0000	2.1142	1.0000	2.1135	1.05
					2.0	1.0000	2.1231	1.0000	2.1210	1.0000	2.1206	
					5.0	1.0000	2.1279	1.0000	2.1270	1.0000	2.1269	
					20.0	1.0000	2.1303	1.0000	2.1301	1.0000	2.1300	
Nsf (US)	8	10	8	10	1.2	1.1824	2.0926	1.1824	2.0926	1.1822	2.0926	0.90
					2.0	1.6069	3.0118	1.6069	3.0025	1.6069	2.9605	
					5.0	1.8057	3.6850	1.8040	3.6010	1.8036	3.5842	
					20.0	1.8155	3.9213	1.8148	3.9002	1.8147	3.8961	
Abovenet (US)	19	34	15	30	1.2	1.0000	2.3789	1.0000	2.3507	1.0000	2.3241	229.41
					2.0	1.0000	3.5284	1.0000	3.5284	1.0000	3.5284	
					5.0	1.0000	3.9172	1.0000	3.9124	1.0000	3.9114	
					20.0	1.3440	7.0130	1.3440	7.0130	1.3440	7.0130	
Sprintlink (US)	52	85	33	65	1.2	1.0000	1.4346	1.0000	1.4219	1.0000	1.4010	51495.02
					2.0	1.0000	1.6827	1.0000	1.6201	1.0000	1.6076	
					5.0	1.3703	7.7900	1.3703	7.7900	1.3703	7.7900	
					20.0	1.9901	12.8507	1.9901	12.8507	1.9901	12.8507	

Table 2 Oblivious performance ratio for OSPF and the optimal oblivious routing under ellipsoidal uncertainty

Name	Orig.		Redu.			$\gamma = 0.75$		$\gamma = 0.90$		$\gamma = 0.99$		
	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$	ζ	oopr	ospf	oopr	ospf	oopr	ospf	t_{avg} [s]
Telstra (AU)	44	46	7	9	0.01	1.0000	2.0810	1.0000	2.0810	1.0000	2.0810	0.06
					0.05	1.0000	2.0810	1.0000	2.0810	1.0000	2.0810	
					0.10	1.0609	2.0810	1.0346	2.0810	1.0278	2.0810	
					0.20	1.1016	2.0810	1.0737	2.0810	1.0680	2.0810	
					0.40	1.1603	2.0810	1.1445	2.0810	1.1413	2.0810	
					0.80	1.2827	2.0810	1.2827	2.0810	1.2827	2.0810	
VNSL (India)	9	11	8	10	0.01	1.0000	2.1311	1.0000	2.1311	1.0000	2.1311	1.84
					0.05	1.0000	2.1311	1.0000	2.1311	1.0000	2.1311	
					0.10	1.0000	2.1311	1.0000	2.1311	1.0000	2.1311	
					0.20	1.0000	2.1311	1.0000	2.1311	1.0000	2.1311	
					0.40	1.0000	2.1311	1.0000	2.1311	1.0000	2.1311	
					0.80	1.0041	2.1311	1.0041	2.1311	1.0041	2.1311	
Nsf (US)	8	10	8	10	0.01	1.0835	2.0147	1.0802	2.0147	1.0792	2.0147	3.28
					0.05	1.2781	2.4483	1.2733	2.4483	1.2716	2.4483	
					0.10	1.4060	2.7815	1.4029	3.7815	1.4019	3.7815	
					0.20	1.5365	3.1327	1.5354	3.1327	1.5351	3.1327	
					0.40	1.6461	3.4622	1.6461	3.4622	1.6461	3.4622	
					0.80	1.7415	3.6602	1.7415	3.6602	1.7415	3.6602	

Table 2 (*continued*)

Name	Orig.		Redu.			$\gamma = 0.75$		$\gamma = 0.90$		$\gamma = 0.99$		
	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$	ζ	oopr	ospf	oopr	ospf	oopr	ospf	t_{avg} [s]
Abovenet (US)	19	34	15	30	0.01	1.0000	2.5162	1.0000	2.5162	1.0000	2.5162	584.17
					0.05	1.0000	2.8322	1.0000	2.8322	1.0000	2.8322	
					0.10	1.0000	2.9830	1.0000	2.9829	1.0000	2.9829	
					0.20	1.0003	3.9292	1.0000	3.9292	1.0003	3.9293	
					0.40	1.0004	3.9353	1.0002	3.9353	1.0007	3.9353	
					0.80	1.1995	6.2078	1.1993	6.2069	1.1995	6.2066	

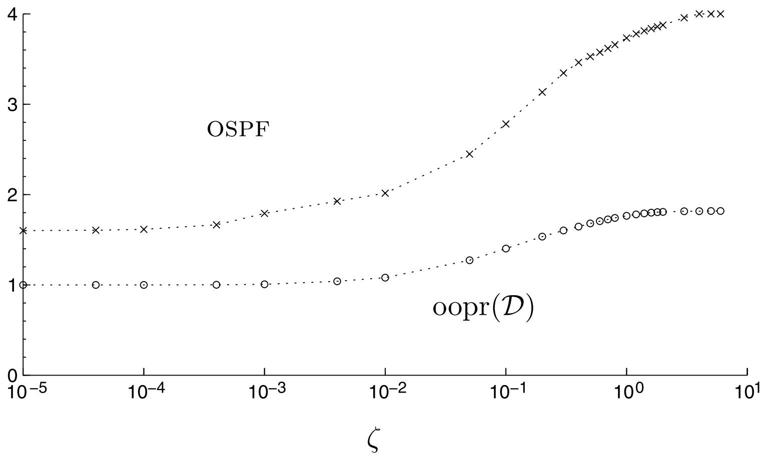

Fig. 1 Comparison of the OSPF and the optimal oblivious performance ratio for network *Nsf* with $\gamma = 0.95$ and different values of ζ

7 Concluding remarks

We have proposed two models to obtain a routing with optimal oblivious performance with respect to two models of demand uncertainty. The first is a linear programming model that deals with demands whose uncertainty is modeled by box constraints. In order to deal with ellipsoidal uncertainty, we have proposed a second-order cone programming model to obtain the optimal routing given a mean-covariance representation of the traffic demand. This proves that the problem of finding optimal oblivious routing with box or ellipsoidal uncertainty can be solved in polynomial time.

From a more practical viewpoint, we compare the optimal oblivious routing with the more common OSPF routing technique, where edge weights are fixed according to a simple rule. We have observed that an optimized routing has a much better performance ratio, and a good level of robustness even with high uncertainty. It remains to be investigated whether a better choice of OSPF weights can improve the performance observed in our tests.

References

Andersen ED, Andersen KD (2000) The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In: Frenk H, Roos K, Terlaky T, Zhang S (eds) High performance optimization. Kluwer Academic, Dordrecht, pp 197–232. Code available from http://www.mosek.com/

Applegate D, Cohen E (2003) Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In: Proceedings of *SIGCOMM '03*, Karlsruhe, Germany, pp 313–324

Azar Y, Cohen E, Fiat A, Kaplan H, Räcke H (2003) Optimal oblivious routing in polynomial time. In: Proceedings of STOC '03, San Diego, California, pp 383–388

Ben Tal A, Nemirovski A (1999) Robust solution of uncertain linear programs. Oper Res Lett 25(1):1–13

Ericsson M, Resende MGC, Pardalos PM (2002) A genetic algorithm for the weight setting problem in ospf routing. J Comb Optim 6(3):229–333

Fortz B, Thorup M (2000) Internet traffic engineering by optimizing OSPF weights. In: Proceedings of IEEE INFOCOM, pp 519–528

Fourer R, Gay DM, Kernighan BW (1990) A modeling language for mathematical programming. Manag Sci 36:519–554. See also http://www.ampl.com

Ilog Inc (2003) CPLEX 9.0 users's manual

Li Z, Li B, Jiang D, Lau LC (2004) On achieving optimal end-to-end throughput in data networks: theoretical and empirical studies. ECE Technical Report, University of Toronto, May 2004

Lin FYS, Wang JL (1993) Minimax open shortest path first routing in networks supporting the SMDS service. In: Proceedings of IEEE ICC, vol 2, pp 666–670

Mitra D, Ramakrishnan KG (1999) A case study of multiservice, multipriority traffic engineering design for data networks. In: Proceedings of IEEE Globecomm '99, Rio de Janeiro, Brasil, pp 1087–1093

Roughan M, Thorup M, Zhang Y (2003) Traffic engineering with estimated traffic matrices. In: Proceedings of IMC '03, Miami Beach, FL, USA, pp 248–258

Springs N, Mahajan R, Wetherall D (2004) Measuring ISP topologies with rocketfuel. IEEE/ACM Trans Netw 12(1):2–16

Tebaldi C, West M (1998) Bayesian inference on network traffic using link count data. J Am Stat Assoc 93(442):557–576

Vardi Y (1996) Network Tomography: estimating source-destination traffic intensities from link data. J Am Stat Assoc 91(433):365–377

Zhang Y, Roughan M, Lund C, Donoho D (2003) An information-theoretic approach to traffic estimation. In: Proceedings of SIGCOMM '03, Karlsruhe, Germany, pp 301–312