Chance Constrained Network Design

Fausto Pascali* Maria Grazia Scutellà* Antonio Frangioni *

*Department of Computer Science, University of Pisa
Largo B. Pontecorvo 3, Pisa, Italy

Abstract

When designing or upgrading a communication network, operators are faced with a major issue, as uncertainty on communication demands makes it difficult to correctly provision the network capacity. When a probability on traffic matrices is given, finding the cheapest capacity allocation that guarantees, within a prescribed level of confidence, that each arc can support the traffic demands peaks turns out to be, in general, a difficult non convex optimization problem belonging to the class of chance constrained problems. Drawing from some very recent results in the literature we highlight the relationships between chance constrained network design problems and robust network optimization. We then compare several different ways to build uncertainty sets upon deviation measures, comprised the recently proposed backward and forward deviation measures that capture possible asymmetries of the traffic demands distribution. We report results of a computational study aimed at comparing the performance of different models when built upon the same set of historical traffic matrices.

Keywords: *chance constraint, robust optimization, network design, deviation measures, routing.*

1 Introduction

When designing or upgrading a communication network, operators are faced with a major issue, as uncertainty on communication demands makes it difficult to correctly provision the network capacity. In fact, providing large capacity, while making the network resilient to unexpected peaks of demand, may be very costly, and therefore render the network operations uneconomical. It is therefore necessary to carefully balance the network failure probability due to high demand, on one side, and the capacity provision cost on the other. When a probability on traffic matrices is given, finding the cheapest capacity allocation that guarantees, within a prescribed level of confidence, that each arc can support the traffic demand peaks turns out to be, in general, a difficult non convex optimization problem, belonging to the class of chance constrained problems. An alternative approach is to properly formulate an uncertainty set to which all demand matrices supposedly belong, and to find a minimum cost capacity allocation that supports all the demands in such a set; this is the so-called Robust Network Design problem, which, depending on the structure of the uncertainty set, may be polynomially solvable [4, 5, 14]. Drawing from some very recent results in the literature, we highlight the relationships between the two approaches, describing several different ways in which uncertainty sets can be constructed that "well approximate" the chance constraints while preserving the computational tractability of the model [9]. Since each uncertainty set approximates in a different way the actual non convex set inherent to the "true" chance constraints, we report results of a computational study aimed at finding what approximation provides better results, in the sense of producing the capacity allocation that sits within an allotted monetary budget and that minimizes the actual ex-post network failure probability. We are not aware of previous results along these lines except those in [3]; however, there the capacities of the arcs are given, and the focus is on minimizing a network congestion measure, i.e., the oblivious performance ratio.

2 The chance constrained network design problem

Consider a communication network represented by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V}$ is the set of nodes, with $n = |\mathcal{V}|$, and $\mathcal{A}$ is the set of arcs, with $m = |\mathcal{A}|$. We are given a set $\mathcal{K} \subseteq \mathcal{V} \times \mathcal{V}$ of ordered pairs which represents users that wish to communicate, with $k = |\mathcal{K}|$. We are also given a nonnegative cost c_{ij}, $\forall\ (i,j) \in \mathcal{A}$, which is the cost to reserve (or allocate) a unit of capacity on the arc (i,j) of the network. Our task is to provide the network with a capacity allocation and to decide how to accommodate the commodities. Many applications require that the commodities are routed along the same paths independently of the traffic demands. To model this assumption, we denote by y_{ij}^{st} the fraction of commodity $(s,t) \in \mathcal{K}$ routed along the arc $(i,j) \in \mathcal{A}$. Moreover, let $\delta^-(i)$ ($\delta^+(i)$) be the set of incoming (outgoing) arcs of node i, $\forall\ i \in \mathcal{V}$. Formally, we define an oblivious routing as follows:

Definition 1 *Given a directed network $G = (\mathcal{V}, \mathcal{A})$ and a set of ordered pairs $\mathcal{K} \subseteq \mathcal{V} \times \mathcal{V}$, we define an* ***oblivious routing*** *as a function $y : \mathcal{A} \times \mathcal{K} \to [0,1]$, such that:*

$$\sum_{(i,j)\in\delta^-(i)} y_{ij}^{st} - \sum_{(j,i)\in\delta^+(i)} y_{ji}^{st} = \phi_i^{st} \quad = \begin{cases} -1 & \text{if } i = s, \\ \ \ 1 & \text{if } i = t, \\ \ \ 0 & \text{otherwise.} \end{cases} \qquad i \in \mathcal{V},\ (s,t) \in \mathcal{K}; \tag{1}$$

We denote by $\mathcal{Y}$ the set of all oblivious routings so that $y \in \mathcal{Y}$ means, in a compact way, that y is an oblivious routing.

If the amount of future connection request d_{st} from node s to node t is given, $\forall (s,t) \in \mathcal{K}$, then we face the following deterministic problem:

$$\min\ cx \tag{2}$$
$$y_{ij}d - x_{ij} \le 0 \qquad (i,j) \in \mathcal{A}; \tag{3}$$
$$(y,x) \in \mathcal{Y} \times \mathcal{R}_+^m$$

where y_{ij} is the vector with components $\{y_{ij}^{st}\}$. In order to address uncertainty, let us model the traffic demand vector d as a multivariate random variable $\tilde{d}$ with a probabilistic density function $\mathbb{P} : \mathbb{R}_+^k \to [0,1]$, $(d \mapsto \mathbb{P}(d))$. Let $\epsilon \in [0,1]$ be the confidence level required for the each arc of the network. Then, we consider the following individual chance constraint problem [6]:

$$\min\ cx \tag{4}$$
$$\mathbb{P}(y_{ij}\tilde{d} - x_{ij} \le 0) \ge \epsilon \qquad (i,j) \in \mathcal{A}, \tag{5}$$
$$(y,x) \in \mathcal{Y} \times \mathcal{R}_+^m.$$

This is, in general, a difficult non convex optimization problem, due to the chance constraints (5). However, we can build alternative tractable models based on a set $\mathcal{S}$ of historical traffic matrices. Let us denote the generic approximation model as follows:

$$c(\epsilon) = \min\ cx \tag{6}$$
$$H_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S}) \le 0 \qquad (i,j) \in \mathcal{A}; \tag{7}$$
$$(y,x) \in \mathcal{Y} \times \mathcal{R}_+^m.$$

where $H_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S})$ denotes a multivalued function: $H_{ij} : (\mathcal{R} \times \mathcal{R}^k \times \mathcal{R} \times \mathcal{R}^{|\mathcal{S}|}) \to \mathcal{R}^q$ (where q is a integer number depending on the model). To preserve the computational tractability, we shall assume that H is a convex function.

The most simple approach is to specify a single representative matrix [8], which somehow "synthesizes" the future possible traffic demands; in this case model (6)-(7) reduces to the classical linear programming model (2)-(3). Different choices of d into constraints (3) result in different models. When no statistical information is known (apart from a set of previous realizations), one may set $d = \epsilon d_r$ where d_r is a traffic matrix randomly chosen in $\mathcal{S}$; we refer to this strategy as the *single random model* (*BSO*). By using instead the first moment statistical information we may set $d = \epsilon \overline{d}$, where $\overline{d}_{st} = \sum_{h \in \mathcal{S}} d^h_{st}/|\mathcal{S}|$ is the mean vector of the sample set $\mathcal{S}$; we denote this as the *single average model* (*BSA*). Alternatively, we can use the vector of demand peaks by setting $d = \epsilon d^{max}$, where $d^{max}_{st} = \max_{h \in \mathcal{S}} d^h_{st}$; we denote this as the *single peaks model* (*BSP*).

A different approach, still preserving the linearity of the model, is based on the assumption that all the traffic demands in $\mathcal{S}$ will realize non-simultaneously, as in [7]. In this case $q = |\mathcal{S}|$, and for all $(i,j) \in \mathcal{A}$ $H_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S})$ is a vector with q components $H^h_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S})$, where:

$$H^h_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S}) = -\, x_{ij} + \sum_{(s,t) \in \mathcal{K}} \epsilon d^h_{st} y^{st}_{ij} \qquad \forall h \in \mathcal{S}. \tag{8}$$

We denote this as the *multiple model* (*BMU*). It is important to stress the conceptual difference between the parameter ϵ in constraints (5), which is a pure probability value, and the parameter ϵ in the above models, which tunes the expected value of future demands in order to fit with a certain available budget.

Finally, we discuss how to get a convex tractable approximation of the chance constraints (5) following the results of [9]. Let $z(\tilde{d}) = (y_{ij}\tilde{d} - x_{ij})$. In [9] it is proven, in a more general setting, that

$$\min_{\tau} \left\{ \tau + \frac{1}{\epsilon} \mathbb{E}[z(\tilde{d}) - \tau]^+ \right\} \leq 0 \Rightarrow \mathbb{P}(z(\tilde{d}) \leq 0) \geq \epsilon. \tag{9}$$

The left hand side inequality of (9) corresponds to constrain the ϵ-Conditional Value at Risk of the random variable $z(\tilde{d})$ to be non positive [2, 10, 11]. In order to avoid expensive numerical calculation related to the expected value function, in [14] it is proposed to bound $\mathbb{E}[\cdot]^+$ with a suitable function which preserves the computational tractability. The choice of the bounding function depends on the assumption we can make about the probability distribution of the random parameters. In order to capture possible asymmetries of the random variable distribution, some deviation measures are introduced in [12], which are called *Forward and Backward deviation measures.* The authors define the set of the values associated with *forward deviations* of $\tilde{d}$ as follows:

$$\mathcal{P}\left(\tilde{d}\right) = \left\{ \alpha \;:\; \alpha \geq 0,\ \mathbb{E}\big[e^{\frac{\theta(\tilde{d} - E[\tilde{d}])}{\alpha}}\big] \leq e^{\frac{\theta^2}{2}}\ \forall \theta \geq 0 \right\}.$$

Likewise, for *backward deviations*, they define the set

$$\mathcal{Q}\left(\tilde{d}\right) = \left\{ \alpha \;:\; \alpha \geq 0,\ \mathbb{E}\big[e^{-\frac{\theta(\tilde{d} - \mathbb{E}[\tilde{d}])}{\alpha}}\big] \leq e^{\frac{\theta^2}{2}}\ \forall \theta \geq 0 \right\}.$$

These deviation measures can be used in order to bound the ϵ-Conditional Value at Risk of the random variable $z(\tilde{d})$ [13]. Specializing the result in [13] to our application, and after some algebraic calculation, we can derive our last model (BDE), where:

$$H_{ij}(x_{ij}, y_{ij}, \epsilon, \mathcal{S}) = -\, x_{ij} + y_{ij}\overline{d} + \sqrt{-2\ln(\epsilon) \sum_{(s,t) \in \mathcal{K}} (y^{st}_{ij} p^{st})^2}. \tag{10}$$

Notice that (6) when equipped with (10) becomes a second order cone program;note also that these problems are nowadays routinely solvable by standard approaches. The computation of forward and backward deviation measures of $\tilde{d}$ out of a sample set $\mathcal{S}$ is discussed in [14].

Aim of this paper is to compare the ex-post performance of all the previous models in a *budget perspective*. That is, for $l \in \{BSO,\ BSA,\ BSP,\ BMU,\ BDE\}$ we consider the problem

$$\max_{\epsilon} \ \left\{ \ \epsilon \ : \ C^l(\epsilon) \leq B \ \right\}, \tag{11}$$

where $C^l(\epsilon)$ denotes the minimum cost provided by model l for a fixed value of ϵ, while B denotes a given budget. In general we are not able to solve (11) in one shot; however, since we can compute $C^l(\epsilon)$ in polynomial time, we solve (11) through a binary search on ϵ. Our testing methodology consists in extracting a *sample set* from a given set of realizations of the demand matrices, solving (11) on these sets, and then comparing the ex-post resilience properties of the (y, x) solutions obtained by each model on appropriately defined *testing sets*.

3 Computational results

For our tests we used data of the *Abilene network*, which has 30 arcs, 12 nodes, and 132 commodities; six months of traffic matrices at the hourly discretization are provided at [1]. We set the arc cost proportionally to the physical distance of the endpoints. From this pool of $\mathcal{H}$ historical traffic matrices we extracted the sample sets $\mathcal{S}$ and a testing set $\mathcal{T}$. We built the sample sets $\mathcal{S}$ and $\mathcal{T}$ in 2 possible ways, by selecting

(A) one hour at random for each of the 7 days of a week;

(C) all 24 hours of one random day.

Each sample set $\mathcal{S}$, which collects the traffic matrices of 7 weeks, is equipped with a benchmark budget Br which corresponds to the optimal value of model (6)-(7) equipped with (8) with $\epsilon = 1$. Then, for each sample set we built and solved the five models for three different levels of budget: *low*, *nominal* and *high* level, corresponding to $B = p * Br$ with $p \in \{0.5, 1, 1.2\}$, respectively. Each pair (model, level of budget) produced a different solution, i.e., a capacity allocation x and a routing scheme y. We evaluated the empirical probability of each solution (y, x) according to

$$EP(x, y) = \frac{\sum_{h \in \mathcal{T}} \sum_{(i,j) \in A} \mathcal{I}[x_{ij} - \sum_{(s,t) \in K} d^h_{st} y_{ij}]}{|\mathcal{T}||A|} \tag{12}$$

on the testing set $\mathcal{T}$ which collects the traffic matrices of the remaining 14 weeks. In (12), the operator $\mathcal{I}[\cdot]$ is the indicator function: for each traffic matrix $d \in \mathcal{T}$, it counts how many times the capacity x_{ij} of the arc (i, j) is able to support the flow $\sum_{(s,t) \in K} d^h_{st} y_{ij}$ routed along it. We ran all the models over 240 different sample sets of type A and 210 sample sets of type C, and for each model we report the mean empirical probability with respect to the produced solutions. We do not report computational time, since such a comparison is out of the scope of this paper; we just mention that the commercial solver `Cplex 11` required on average less then one second to solve models (BSA), (BSP), (BSO), around 4 seconds for (BMU), and around 90 seconds for (BDE).

Since we do not have any a priori information about the characteristics of the historical traffic demands, we arbitrarily selected the two different types of sample sets. As a first step, we compared the results produced by the two types of sampling. The mean values are reported in Figure 3: it is evident that the sample set of type A always leads to solutions with a better performance, particularly for the nominal level of the budget. This implies that the considered traffic matrices are better described by a hourly fluctuation, rather than daily fluctuation.

Figure 2 reports the results of the sample sets of type A by varying the level of budget. The model which provided a network with a higher mean empirical success probability is (BSA), while the model which performed worst is (BDE). Models (BMU) and (BSP) have essentially the some behavior, with (BSP) performing just a little bit better (about 0.01 %) than (BMU). The behavior of model (BSO) is interesting since its performance falls as the available budget increases.

Looking at Figure 3, which reports the results on the sample sets of type C, we observe that the model (BSA) still performs better, while the performance of model (BDE) considerably increases as the level of the budget does. This seems to indicate that when we have some information about the traffic demands, a model using just first moment statistical information may be able to attain good performances, while

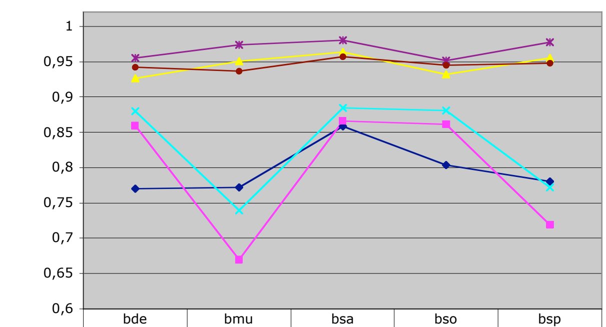

	bde	bmu	bsa	bso	bsp
0.5A	0,769912368	0,772053571	0,858731812	0,803354828	0,780253347
0.5c	0,859060847	0,669417989	0,866061508	0,861180556	0,719231151
1A	0,926390542	0,950650751	0,963388889	0,932141204	0,955352183
1C	0,879545304	0,739323743	0,884196429	0,880472884	0,771635251
1.2A	0,954834656	0,973862906	0,98005291	0,951063195	0,977375992
1.2C	0,941641865	0,936337632	0,95667328	0,944611442	0,947941468

Figure 1: *Empirical Probability, compare sample sets of type A and C, independently of the models*

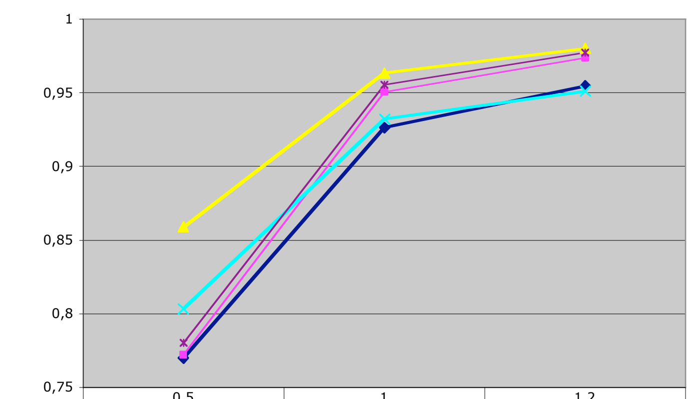

	0,5	1	1,2
BDE	0,769912368	0,926390542	0,954834656
BMU	0,772053571	0,950650751	0,973862906
BSA	0,858731812	0,963388889	0,98005291
BSO	0,803354828	0,932141204	0,951063195
BSP	0,780253347	0,955352183	0,977375992

Figure 2: *Empirical Probability, sample sets of type A*

when the knowledge on traffic demands reduces, then the use of higher order statistical information becomes more important.

Finally, in Figure 4 we study the behavior of a fixed sample set (of type A) on a wide range of possible budgets; we also show the relation between the empirical values of probability and the theoretical one that we can infer from the solution of model (10). As we expected, as the budget increases both the empirical and the theoretical probability increase. The gap between the empirical and theoretical values tends to reduce for higher budgets, i.e., for higher level of probability. However, for smaller values the

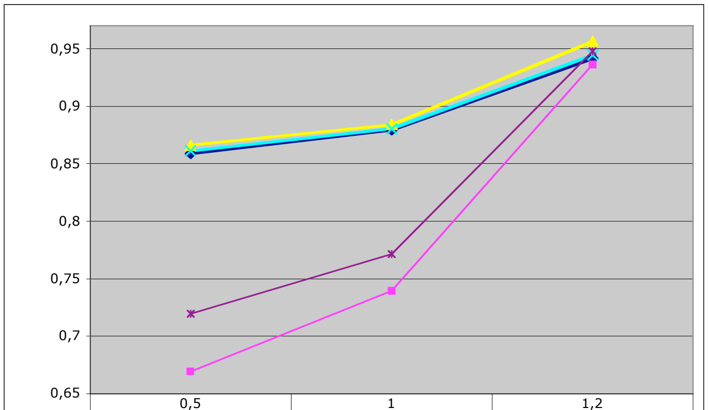

	0,5	1	1,2
BDE	0,859060847	0,879545304	0,941641865
BMU	0,669417989	0,739323743	0,936337632
BSA	0,866061508	0,884196429	0,95667328
BSO	0,861180556	0,880472884	0,944611442
BSP	0,719231151	0,771635251	0,947941468

Figure 3: *Empirical Probability, sample sets of type C*

differences among the models become more significant, in particular for the sophisticated model (BDE).

4 Conclusion

Our preliminary results show that the different models attain significantly different ex-post results when the sampling set and the budget level change. Therefore, it is of relevant interest to devise guidelines which allow modelers and decision makers to properly choose, among the available ones, the uncertainty model which is best suited to the specific situation they face.

References

[1] The Abilene network. `http://www.cs.utexas.edu/~yzhang/research/AbileneTM/`.

[2] P. Artzener, F. Delbean, J. M. Eber and D. Health, *Coherent measures of risk* , Mathematical Finance **9** (1999) pp. 203–228.

[3] P. Belotti and M. Ç. Pinar, *Optimal oblivious routing under linear and ellipsoidal uncertainty*, Optim. Eng. **9** (2008) 257–271.

[4] D. Bertsimas and M. Sim, *Robust discrete optimization and network flows* , Math. Prog. **98** (2003) pp. 49–71.

[5] A. Ben-Tal and A. Nemirovski, *Robust solutions of linear programming problems contaminated with uncertain data* , Math. Prog. **88** (2000) pp. 411–424.

[6] A. Charnes and W. W. Cooper, *Chance constrained programming* , Man. Science **6** (1959) pp. 73–79.

[7] M. Labbé, R. Séguin, P. Soriano, and C. Wynants, *Network synthesis with non- simultaneous multycommodity flow requirements: Bounds and heuristics* Tech. report, ULB, (1999).

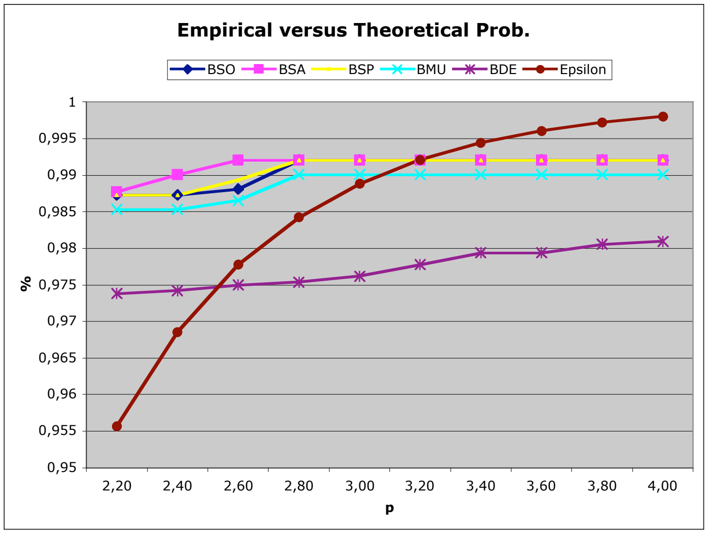

Figure 4: *Empirical Probability, sample sets of type A, different levels of budget*

[8] M. Minoux *Network synthesis and optimum network design problems: Models, solution methods and applications*, Networks **19** (1986) pp. 313–360.

[9] A. Nemirovski and A. Shapiro, *Convex approximations of chance constrained programs* , SIAM J. on Opt. **17** (2006) pp. 969–996.

[10] R. T. Rockafellar and S. Uryasev, *Optimization of conditional value-at-risk* , J. of Risk **2** (2000) pp. 21–41.

[11] R. T. Rockafellar, S. Uryasev and M Zabarankin, *Generalized deviations in risk analysis* , Finance and Stochastics **10** (2006) pp. 51–74 .

[12] M. Sim and W. Chen, *Goal driven optimization* , accepted in Operations Research (2006).

[13] M. Sim, X. Chen, and P. Sun. *A robust optimization perspective of stochastic programming* , Operations Research **55** (2007) pp. 1058–1107.

[14] M. Sim, C. Piaw Teo, W. Chen and S. Jie, *From cvar to uncertainty set: Implications in joint chance constrained optimization* , submitted to Operations Research (2007).