2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

A Chernoff Bound Approximation for Risk-Averse Integer Programming

Randy Cogill Zhou Zhou
rcogill@virginia.edu zz5x@virginia.edu
Dept. of Systems and Information Engineering
University of Virginia

***Abstract*— In this paper we consider optimization problems with a stochastic performance measure, where the goal of the problem is to find a solution that minimizes the probability that this performance measure exceeds a given threshold. It is known that this and related problems are computationally intractable, so we consider an approach that seeks to minimize an upper bound on the probability of exceeding the given threshold. From this approach, we obtain a suboptimal solution, together with a guaranteed upper bound on the achieved exceedance probability.**

First, we present an algorithm that minimizes a Chernoff bound by solving a binary integer program. For problems with totally unimodular constraint sets, this Chernoff bound can be minimized by solving a linear program. This formulation is shown to recover several known results for the cases of Gaussian and stochastically dominant costs. We then briefly consider these problems in a closed loop setting, where solutions can be refined as the values of uncertain quantities in the model are revealed. We propose an open-loop feedback control algorithm where a binary integer program (or possibly linear program) is solved in each time step given the current state of the system.

I. INTRODUCTION

In this paper we consider optimization problems with a stochastic performance measure. Specifically, we consider problems with a risk-based objective, where the goal of the problem is to find a solution that minimizes the probability that the stochastic performance measure exceeds a given threshold.

An example of an application where this problem arises is in travel route planning, where the objective is to travel between two locations within a specified amount of time [12]. Here, the ideal route is one which minimizes the probability of exceeding the given threshold on travel time. This is particularly relevant in emergency response settings, where the common quality of service measures are specified in terms of the fraction of calls that must be attended to within a given response time (*i.e.*, 90% of calls must be responded to within 5 minutes).

It is known that problems related to the ones considered in this paper are computationally intractable, primarily due to the fact that even evaluating the objective can be as computationally intensive as enumerating the solutions to an NP-hard optimization problem. So, we consider an approach that seeks to minimize an *upper bound* on the probability of exceeding the given threshold. From this approach, we obtain a suboptimal solution, together with a guaranteed upper bound on the exceedance probability achieved by this solution. The approach we consider seeks to minimize a Chernoff bound on the exceedance probability. Due to the fact that the problems we consider have binary variables, minimization of this Chernoff bound can be performed exactly by solving a binary integer program. For problems with totally unimodular constraint sets, minimization of the Chernoff bound can be performed exactly by solving a linear program. Such problems include classical combinatorial optimization problems such as the shortest path problem, the weighted matching problem, and the minimum spanning tree problem. Hence, stochastic versions of these problems can benefit directly from the algorithms presented in this paper.

We also briefly discuss the risk-averse optimization problem in a closed-loop setting. In the closed loop setting, solutions can be refined as the values of uncertain quantities in the model are revealed. We consider an open loop feedback control algorithms, where a binary integer program (or possibly linear program) is applied in each time period. This allows us to apply the Chernoff bound technique in the closed-loop setting, where decisions are made in each time period given the current state of the system.

II. PROBLEM FORMULATION

In this paper we consider optimization problems where the objective is to minimize the probability that a stochastic cost exceeds a given threshold. Specifically, the problem considered in this paper is the following:

$$\begin{aligned} \text{minimize:} \quad & \mathbf{P}\big(C^T x \geq L\big) \\ \text{subject to:} \quad & Ax \leq b \\ & x \in \{0,1\}^n \end{aligned} \tag{1}$$

In this problem the constraint matrix A and vector b are deterministic, and the vector C is stochastic. Specifically, we consider a model where the n-dimensional random vector C has statistically independent components, and the moment generating function exists for each component. So, there is no uncertainty in the feasible set of this problem. All uncertainty in the model is focused on $C^T x$, the performance measure for a given feasible solution. This is precisely the structure of the uncertainty found in a shortest path problem with stochastic edge costs, one of the main problems motivating this work.

III. Prior Work

The problems considered in this paper are most closely related to chance constrained optimization problems, first introduced by Charnes, Cooper, and Symonds in [9], Miller and Wagner [14], and Prekopa [24]. Chance constrained optimization problems are problems with probabilistic constraints of the form

$$\mathbf{P}(C^T x \geq L) \leq \alpha,$$

possibly in addition to deterministic constraints. This provides a generalization of robust optimization, since the special case where $\alpha = 0$ requires the constraint $C^T x < L$ to hold for all realizations of C. For more details, we refer to [2].

The problems we consider here are binary integer programs, which are NP-hard in general. However, chance constraints add another level of complexity to these problems. For example, the feasible region defined by a chance constraint may not be convex. Also, depending on the nature of the random variables, simply testing a chance constraint for a given solution might be computationally challenging. However difficult to tackle in general, there are elegant results that apply when the feasible region defined by the chance constraint is proved convex. The mathematical tools to prove convexity of distributions can be found in [26] and the literature therein. Based on these results, an exact algorithm called cone generation and its application are given, which is similar to the well-known column generation method [1], see [10] and [11] for details. The major computational difficulties of column generation also apply to cone generation: the worst-case running time of this approach is still exponential.

Aside from exact approaches, there are four representative approximation methods for chance constrained problems: Normal approximation, K-σ approximation, Scenario approximation and Bernstein approximation. All of them share the same spirit of replacing the chance constraints with deterministic constraints. K-σ approximation requires the assumption that the distributions or random quantities are symmetric (see [3], [2] and [15]). It was implied in [20] that if the random variables are simultaneously perfectly correlated, the K-σ approximation reduces to a linear approximation.

The Scenario approximation approach is based on Monte Carlo sampling. This approach generates numerous realizations for the random quantities, and solves a deterministic problem containing these deterministic constraints. Calafiore and Campi showed in [7], [8] that as the sample size N increases, the probability that the chance constraint is violated decreases rapidly. Normal approximation is a natural way to approximate the chance constraints and is already used in engineering problems (e.g. [21]). However, it should be noted that normal approximation does not provide a guarantee of feasibility of the approximation, especially when the number of random variables is small. We again refer to [26] for the details of Normal approximation. In Bernstein approximation, Bernstein inequalities are applied to the chance constraints. Bernstein inequalities were used by Pinter to approximate chance constraints in the example given in [22] and later in [4] to robustly solve NETLIB problems. This approximation framework was later refined and generalized in [16]. Bernstein approximation is conservative, unlike Normal approximation, and there is no restriction of the distribution. Also, it is much less demanding than Scenario approximation in computational intensity. The approach taken in our paper is similar to Bernstein approximation, however we are able to exploit the fact that our variables are binary to minimize a Chernoff bound exactly.

Also, the related problem of constructing a risk-averse path in a graph with stochastic weights has been studied by numerous researchers in a variety of fields. Approaches range from numerically computing the dynamic programming value function [12], [17], to providing exact algorithms for special classes of link travel time distributions [18], [19].

The results in the current paper are most closely related to those in [18], [19], [13]. In that series of papers, the authors present a broad range of theoretical and practical results for risk averse integer programs. In particular, the authors characterize the problems and solutions for families of cost distributions (most notably, Gaussian and stochastically dominant), and give algorithms with performance guarantees for a problem approximation based on the Chebyshev bound. These papers present many strong results, and we view the current paper as complementing that work in that we present a practical approximation that is not covered in those papers, and recovers the problem formulations presented in those papers for Gaussian and stochastically dominant costs.

The complexity of evaluating $\mathbf{P}(C^T x \geq L)$

Here we discuss the complexity of evaluating the objective $\mathbf{P}(C^T x \geq L)$ for generally distributed C_i. It turns out that simply *evaluating* this objective is as hard as counting the number of feasible solutions to the knapsack problem, an NP-hard optimization problem. This complexity result is the main justification for the approximation method we use here, where $\mathbf{P}(C^T x \geq L)$ is replaced by a tractable upper bound. The complexity of this problem is discussed in [13], but we present an analysis of the complexity here to keep our treatment self-contained.

The knapsack problem is

$$\begin{aligned} \text{maximize:} \quad & \textstyle\sum_{i=1}^n r_i z_i \\ \text{subject to:} \quad & \textstyle\sum_{i=1}^n a_i z_i \leq W \\ & z \in \{0,1\}^n \end{aligned}$$

where $r_1, \ldots, r_n$ and $a_1, \ldots, a_n$ are nonnegative integers. Here we'll show that any algorithm capable of evaluating $P(C^T x \geq L)$ for arbitrarily distributed C_i can also be used to count all feasible solutions to any instance of the knapsack problem.

Suppose $C_1, \ldots, C_n$ are independent and have

$$C_i = \begin{cases} 0 & \text{with probability } \frac{1}{2} \\ a_i & \text{with probability } \frac{1}{2} \end{cases}$$

Also, without loss of generality, suppose $x = \mathbf{1}$. Then

$$\begin{aligned}\mathbf{P}(C^T x \geq W+1) &= 1 - \mathbf{P}(C^T x < W+1) \\ &= 1 - \mathbf{P}(C^T x \leq W) \\ &= 1 - \frac{K}{2^n}\end{aligned}$$

where K is the number of $z \in \{0,1\}^n$ such that

$$\sum_{i=1}^{n} a_i z_i \leq W$$

Hence, any algorithm capable of evaluating $\mathbf{P}(C^T x > W)$ can also be used to compute the number of feasible solutions to an instance of the knapsack problem.

The problem of counting knapsack solutions is in a complexity class called #P-complete. If a polynomial-time algorithm exists for any #P-complete problem, then all problems in NP can be solved in polynomial time. Hence, #P-complete problems are generally considered to computationally intractable.

Although the problem of evaluating $\mathbf{P}(C^T x \geq L)$ is generally hard, of course there are classes of distributions for which this objective can be evaluated and optimized. Although several of these cases have been discussed in previous literature, we will show in later sections how the known problem formulations for these cases can also be obtained from our Chernoff bound approach.

IV. Chernoff bound minimization

We will begin this section by presenting a theorem that characterizes a class of upper bounds on the solutions of (1). This theorem leads directly to an algorithm for minimizing an upper bound on the objective value of (1).

Theorem 1: For a given $t > 0$, let $u(t)$ be the minimum achievable objective value of the binary integer program

$$\begin{array}{rl}\text{minimize:} & \sum_{i=1}^{n} x_i \ln\big(\mathbf{E}\big[e^{tC_i}\big]\big) \\ \text{subject to:} & Ax \leq b \\ & x \in \{0,1\}^n\end{array} \tag{2}$$

A solution x^* achieving this minimum satisfies

$$\mathbf{P}\big(C^T x^* \geq L\big) \leq e^{u(t)-tL}$$

Proof: By the Chernoff bound, the inequality

$$\mathbf{P}\big(C^T x \geq L\big) \leq \mathbf{E}[e^{t(C^T x - L)}]$$

is satisfied for all $t > 0$. Also, since the components of C are independent,

$$\mathbf{E}[e^{t(C^T x - L)}] = e^{-tL} \prod_{i=1}^{n} \mathbf{E}[e^{tC_i x_i}] \tag{3}$$

Since $\ln(z)$ is monotonically increasing and $\mathbf{E}[e^{t(C^T x - L)}]$ is positive, minimizing the right-hand side of (3) is equivalent to minimizing

$$\ln\left(e^{-tL} \prod_{i=1}^{n} \mathbf{E}[e^{tC_i x_i}]\right) = -tL + \sum_{i=1}^{n} \ln\left(\mathbf{E}[e^{tC_i x_i}]\right)$$

However, note that

$$\ln\left(\mathbf{E}[e^{tC_i x_i}]\right) = \begin{cases} 0 & \text{if } x_i = 0 \\ \ln\left(\mathbf{E}[e^{tC_i}]\right) & \text{if } x_i = 1 \end{cases}$$

Since $x \in \{0,1\}^n$, minimizing the right-hand side of (3) is equivalent to minimizing

$$\sum_{i=1}^{n} x_i \ln\left(\mathbf{E}[e^{tC_i}]\right) \tag{4}$$

subject to the constraints on x. Letting x^* be the feasible x that minimizes (4) and $u(t)$ be the associated minimum value, this gives the result

$$\begin{aligned}\mathbf{P}\big(C^T x^* \geq L\big) &\leq \mathbf{E}[e^{t(C^T x^* - L)}] \\ &= e^{u(t)-tL}\end{aligned}$$

■

For any given parameter $t > 0$, one can minimize an upper bound on the probability of violating the threshold L. This leads to the natural question of how the parameter t should be chosen. For certain special cases, such as Gaussian or Poisson random variables, one can write a closed-form expression for the optimal t in terms of a given x. For these special cases, this leads to algorithms that can minimize the upper bound produced by the Chernoff bound jointly in t and x.

Note that if the optimization variables were not binary, it would not be possible to convert the Chernoff bound directly into an equivalent linear objective. The fact that an upper bound on $\mathbf{P}\big(C^T x^* \geq L\big)$ can be minimized by solving a binary integer program is particularly good news for certain combinatorial problems, such as stochastic shortest path problems. Of particular interest are combinatorial optimization problems where, in their deterministic versions, are efficiently solvable. These problems have the special property that linear program relaxations of their integer programming formulations yield integer optimal solutions. One sufficient condition ensuring that this is true is that a problem has an integer constraint vector b and a *totally unimodular* constraint matrix A [6], [27]. The matrix A is said to be totally unimodular if every square submatrix has determinant 1, 0, or -1. The following theorem shows that locally optimal solutions to our Chernoff bound approximation have integer solutions when relaxations of the underlying deterministic optimization problems have integer optimal solutions.

Lemma 1: Consider the optimization problem in the variables t and x:

$$\begin{array}{rl}\text{minimize:} & \sum_{i=1}^{n} x_i \ln\big(\mathbf{E}\big[e^{tC_i}\big]\big) - tL \\ \text{subject to:} & Ax \leq b \\ & x \in [0,1]^n \\ & t \geq 0\end{array} \tag{5}$$

If A is totally unimodular and b has integer components, the objective value achieved by any local optimum can be achieved by a solution with $x \in \{0,1\}^n$.

Proof: Suppose (t, x) is a local optimum for (5). By holding t fixed and optimizing over x we cannot obtain any further improvement in the objective value, but obtain solution that is no worse than (t, x). With fixed t, this problem is a linear program with a unimodular constraint set and integer b. So, there is an optimal solution $x^* \in \{0,1\}^n$ to this linear program. The solution (t, x^*) obtains the same objective value as (t, x) and has $x^* \in \{0,1\}^n$. ■

We are motivated by the pursuit of algorithms that can be applied to uncertainties with the most general distributions possible. In fact, for routing problems on road networks we model the random travel times on each link as a sum of independent exponential random variables with possibly different parameters. This provides a family of random variables that generalizes the gamma random variable, which is commonly used to model travel times in transportation models [23].

To accommodate the widest class of distributions possible, we use a heuristic that alternates between minimizing t and minimizing x. For given x, the parameter t can be optimized using binary search due to the convexity of

$$e^{-tL} \prod_{i=k}^{K} \mathbf{E}[e^{tC_i^T x_i}]$$

Then for fixed t, x can be optimized by solving (5). We point out that, although the following algorithm globally optimizes t for given x and globally optimizes x for given t, it does not necessarily produce a solution that is jointly optimal in x and t.

The procedure just described can be summarized by the following algorithm:

Algorithm 1:

- To initialize, compute the feasible x that minimizes $\sum_{i=1}^n x_i \mathbf{E}[C_i]$.
- For this x, compute the t that minimizes (3).
- In each subsequent iteration, alternate between solving (5) with the previously computed t, and finding the t that minimizes (3) for the previously computed x.
- Stop when the same x has been computed in two subsequent iterations.

ALGORITHMS FOR SPECIAL DISTRIBUTIONS

Here we show how the Chernoff bound approach can be refined when we restrict our attention to special cases of distributions on C. We will consider the cases of Gaussian costs and stochastically dominant costs. Both cases have been studied previously in [19], and here we recover the known problem formulations using the Chernoff bound approach.

Gaussian costs

When C_i is Gaussian with mean μ_i and variance σ_i^2, its cumulant generating function is given by

$$\ln\left(\mathbf{E}[e^{tC_i}]\right) = \mu_i t + \frac{1}{2}\sigma_i^2 t^2$$

So,

$$\begin{aligned} &\sum_{i=1}^n x_i \ln\left(\mathbf{E}[e^{tC_i}]\right) - tL \\ &\quad = \quad t\left(\sum_{i=1}^n \mu_i x_i - L\right) + \frac{1}{2}t^2\left(\sum_{i=1}^n \sigma_i^2 x_i\right) \end{aligned}$$

We restrict our analysis to values of L such that there exists at least one solution x with

$$\sum_{i=1}^n \mu_i x_i \le L \tag{6}$$

Generally, the Chernoff bound is only less than 1 for such solutions. For a given x satisfying (6), the t that minimizes the expression above is

$$t = \frac{L - \sum_{i=1}^n \mu_i x_i}{\sum_{i=1}^n \sigma_i^2 x_i} \tag{7}$$

Restricting to solutions satisfying (6) and substituting (7) in the objective, we obtain the quasiconcave minimization problem

$$\begin{aligned} \text{minimize:} \quad & -\tfrac{1}{2}\tfrac{(L-\sum_{i=1}^n \mu_i x_i)^2}{\sum_{i=1}^n \sigma_i^2 x_i} \\ \text{subject to:} \quad & Ax \le b \\ & \textstyle\sum_{i=1}^n \mu_i x_i \le L \\ & x \in [0,1]^n \end{aligned}$$

From Lemma 1, we also see that all locally optimal solutions to this problem can be chosen to have $x \in \{0,1\}^n$ when A is totally unimodular and b is integer. Also, we note that locally optimal solutions to this problem can be computed numerically by simply using the algorithm in Section V.

The optimal solution to this problem is known to minimize $\mathbf{P}(C^T x \ge L)$ exactly, and is shown in [19]. In that paper, an entirely different approach is used to arrive at this optimization problem and to show that optimal solutions can be chosen to binary. That paper further gives an upper bound on the number of local optimal contained in this problem.

Stochastic dominance

For random variables C_i and C_j, C_j is said to *stochastically dominate* C_i if

$$\mathbf{P}(C_j \ge L) \ge \mathbf{P}(C_i \ge L)$$

for all L. Here we will consider the case where C is distributed such that for any $i \ne j$, either C_i stochastically dominates C_j or C_j stochastically dominates C_i. One example is the case where the components of C are all exponentially distributed.

Without loss of generality, assume that C_j stochastically dominates C_i. As a consequence, it can be shown that [25]

$$\ln\left(\mathbf{E}[e^{tC_i}]\right) \leq \ln\left(\mathbf{E}[e^{tC_j}]\right)$$

for all t. Suppose that a solution x^* minimizes

$$\sum_{i=1}^{n} x_i \ln\left(\mathbf{E}[e^{tC_i}]\right) \tag{8}$$

for a given value of t. Under the stochastic dominance assumption, this solution minimizes (8) for all $t > 0$.

We can use additional structure of the cumulant generating to show that (8) can be minimized by selecting the solution with the smallest expected cost.

Lemma 2: Suppose C is distributed such that for any $i \neq j$, either C_i stochastically dominates C_j or C_j stochastically dominates C_i. The objective (8) is minimized by minimizing $\mathbf{E}[C^T x]$.

Proof: We can express the cumulant generating function of each C_i as

$$\begin{aligned}\ln\left(\mathbf{E}[e^{tC_i}]\right) &= t\mathbf{E}[C_i] + \left(\ln\left(\mathbf{E}[e^{tC_i}]\right) - t\mathbf{E}[C_i]\right) \\ &= t\mathbf{E}[C_i] + r_i(t)\end{aligned}$$

For any fixed $t > 0$, the solution that minimizes (8) also minimizes

$$\frac{1}{t}\sum_{i=1}^{n} x_i \ln\left(\mathbf{E}[e^{tC_i}]\right) = \sum_{i=1}^{n} x_i \left(\mathbf{E}[C_i] + \frac{r_i(t)}{t}\right)$$

Note that

$$r_i(t) = \ln\left(\mathbf{E}[e^{tC_i}]\right) - t\mathbf{E}[C_i]$$

is continuous on the domain of the cumulant generating function. Also, by L'Hospital's rule,

$$\begin{aligned}\lim_{t\to 0} \frac{r_i(t)}{t} &= \lim_{t\to 0}\left(\sum_{j=1}^{\infty} \kappa_{ij} \frac{t^{j-1}}{(j-1)!} - \mathbf{E}[C_i]\right) \\ &= 0\end{aligned}$$

where κ_{ij} is used to denote the j-th cumulant of C_i, and $\kappa_{i1} = \mathbf{E}[C_i]$.

Since there are finitely many $x \in \{0,1\}^n$, for any $\epsilon > 0$ there exists t' sufficiently small such that

$$\frac{1}{t}\sum_{i=1}^{n} x_i r_i(t) \leq \epsilon$$

for all $x \in \{0,1\}^n$ and $t \leq t_1$. So, t can be chosen sufficiently small such that the minimizer of

$$\sum_{i=1}^{n} x_i \mathbf{E}[C_i] \tag{9}$$

is the minimizer of

$$\frac{1}{t}\sum_{i=1}^{n} x_i \ln\left(\mathbf{E}[e^{tC_i}]\right)$$

Hence, the solution that minimizes (8) for any $t > 0$ is the solution that minimizes (9). ■

By computing the solution that minimizes $\mathbf{E}[C^T x]$, then finding the t that provides the tightest Chernoff bound, we are able to find the t and x that globally minimize (5).

Suppose that in addition to stochastic dominance, the C_i have the further property of being *additive*. That is, if C_i and C_j are distributed according to some family of distributions, this family is additive if $C_i + C_j$ is distributed according to this family as well. One example of an additive, stochastically dominant distribution is the Poisson distribution. When the C_i come from an additive, stochastically dominant distribution, there exists a solution $x \in \{0,1\}^n$ such that

$$\mathbf{P}(C^T x \geq L) \leq \mathbf{P}(C^T z \geq L)$$

for all solutions $z \in \{0,1\}^n$ and all L. Such a solution necessarily has

$$\sum_{i=1}^{n} x_i \ln\left(\mathbf{E}[e^{tC_i}]\right) \leq \sum_{i=1}^{n} z_i \ln\left(\mathbf{E}[e^{tC_i}]\right)$$

for all solutions $z \in \{0,1\}^n$ and all $t > 0$. Hence, as also discussed in [19], the solution that minimizes $\mathbf{E}[C^T x]$ also minimizes $\mathbf{P}(C^T x \geq L)$ when the C_i are additive and stochastically dominant.

V. Open-loop feedback control for the shortest path problem

One of the primary applications of this work is in determining reliable routes in transportation networks. These problems have a dynamic component that we have not discussed in the models presented so far. The static solutions discussed in the previous section can be used as part of an *open-loop feedback control* (OLFC) [5] scheme that adapts routes as the values of uncertain quantities are revealed.

Suppose an instance of a risk-averse shortest path problem is specified by a directed graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, a stochastic edge cost vector C, a cost threshold L, a source vertex s, and a destination vertex d. The algorithms in the previous section can be used to compute a path that minimizes an upper bound on $\mathbf{P}(C^T x \geq L)$. Although one could simply follow this computed path, note that information might be revealed as edges in the graph are traversed. In the case of a transportation problem with travel time as the cost, the travel times across traversed edges are revealed as these edges are crossed. Unlike a deterministic shortest path problem, the feedback revealed when crossing an edge might cause a change the most preferable route to the destination.

To cast this problem in the framework of open-loop feedback control, we will introduce some notation. Let $v(k)$ denote the vertex occupied at time k (where $v(0) = s$) and let $e(k)$ denote the edge crossed between times k and $k+1$. Also, let $L(k)$ denote the remaining 'cost budget' at time k. That is, $L(0) = L$ and

$$L(k+1) = L(k) - c_{e(k)}$$

Finally, let $b(k)$ denote the constraint vector corresponding to a shortest path problem for computing a path between $v(k)$ and d.

At each time period $k \geq 0$, the open-loop feedback control algorithm computes a locally optimal solution to the optimization problem

$$\begin{array}{rl} \text{minimize:} & \sum_{i=1}^{n} x_i \ln\big(\mathbf{E}\big[e^{tC_i}\big]\big) - tL(k) \\ \text{subject to:} & Ax \leq b(k) \\ & x \in [0,1]^n \\ & t \geq 0 \end{array} \tag{10}$$

We refer to the OLFC algorithm as Algorithm 2, the steps of which are described below:

Algorithm 2:

- Initialize $v(0) = s$, $L(0) = L$, and $b(0) = b$.
- While $v(k) \neq s$,
  - Solve the problem (10) using Algorithm 1.
  - Let $v(k+1)$ be the vertex connected to $v(k)$ by edge $e(k)$
  - Let $L(k+1) = L(k) - c_{e(k)}$
  - Let $b(k+1)$ be the constraint vector for computing a path from $v(k+1)$ to d

Here we note that the route traversed by the OLFC policy is exactly the open-loop route when the components of C satisfy stochastic dominance.

VI. Conclusions

In this paper we considered optimization problems with a stochastic performance measure, where the goal of the problem is to find a solution that minimizes the probability that this performance measure exceeds a given threshold. We provided algorithms that minimize an upper bound on the exceedance probability. These algorithms obtain a suboptimal solution, together with a guaranteed upper bound on the exceedance probability achieved by this solution. The algorithms can be applied to a fairly general class of risk-averse combinatorial optimization problems. The primary requirements on the models are independence of the random variables appearing in the performance measure associated with a feasible solution, and existence of moment generating functions for each random variable.

Although the results are not presented here, the algorithms discussed in this paper have been tested on examples involving dynamic computation of paths in graphs. In particular, the algorithms have been applied to the problem of computing risk-averse routes in the road network of a city. This network used in our tests contains approximately 3000 edges and 2000 vertices. For this network, paths can be computed very quickly by implementing our algorithms in existing routing software.

References

[1] C. Barnhart, E.L. Johnson, G.L. Nemhauser, M.W.P. Savelsbergh, and P.H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46(3):316–329, 1998.

[2] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[3] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–14, 1999.

[4] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88(3):411–424, 2000.

[5] D. Bertsekas. Dynamic programming and suboptimal control: A survey from adp to mpc. *European Journal of Control*, 11:310–334, 2005.

[6] D. Bertsimas and J. Tsitsiklis. *Intro. to Linear Optimization*. Athena Scientific, 1997.

[7] G. Calafiore and MC Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

[8] GC Calafiore, MC Campi, and D.A. e Informatica. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

[9] A. Charnes, WW Cooper, and GH Symonds. Cost horizons and certainty equivalents: an approach to stochastic programming of heating oil. *Management Science*, pages 235–263, 1958.

[10] D. Dentcheva, A. Prékopa, and A. Ruszczynski. Concavity and efficient points of discrete distributions in probabilistic programming. *Mathematical Programming*, 89(1):55–77, 2000.

[11] D. Dentcheva, A. Prekopa, and A. Ruszczyski. Bounds for probabilistic integer programming problems. *Discrete Applied Mathematics*, 124(1-3):55–65, 2002.

[12] Y.Y. Fan, R.E. Kalaba, and J.E. Moore. Arriving on time. *Journal of Optimization Theory and Applications*, 127(3):497–513, 2005.

[13] D. Karger and E. Nikolova. Stochastic Combinatorial Optimization with Risk. *MIT-CSAIL-TR-2008-055*, 2008.

[14] B.L. Miller and H.M. Wagner. Chance constrained programming with joint constraints. *Operations Research*, 13(6):930–945, 1965.

[15] A. Nemirovski. On tractable approximations of randomly perturbed convex constraints. In *42nd IEEE Conference on Decision and Control, 2003. Proceedings*, volume 3, 2003.

[16] A. Nemirovski and A. Shapiro. Scenario approximations of chance constraints. *Probabilistic and randomized methods for design under uncertainty*, pages 3–47, 2006.

[17] Y.M. Nie and X. Wu. Shortest path problem considering on-time arrival probability. *Transportation Research Part B: Methodological*, 43(6):597–613, 2009.

[18] E. Nikolova, M. Brand, and D.R. Karger. Optimal route planning under uncertainty. In *Proceedings of International Conference on Automated Planning and Scheduling*, 2006.

[19] E. Nikolova, J. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. *Algorithms–ESA 2006*, pages 552–563, 2006.

[20] D.L. Olson and S.R. Swenseth. A linear approximation for chance-constrained programming. *Journal of the Operational Research Society*, 38(3):261–267, 1987.

[21] R.K. Patro and S. Bhatnagar. A probabilistic constrained nonlinear optimization framework to optimize RED parameters. *Performance Evaluation*, 66(2):81–104, 2009.

[22] J. Pinter. Deterministic approximations of probability inequalities. *Mathematical Methods of Operations Research*, 33(4):219–239, 1989.

[23] A. Polus. A study of travel time and reliability on arterial routes. *Transportation*, 8(2):141–151, 1979.

[24] A. Prékopa. On probabilistic constrained programming. In *Proceedings of the Princeton Symposium on Mathematical Programming*, pages 113–138. Citeseer, 1970.

[25] S.M. Ross. *Introduction to Stochastic Dynamic Programming*. Probability and Mathematical Statistics. Academic Press, San Diego, CA, 1983.

[26] A. Shapiro, D. Dentcheva, A. Ruszczyński, and A.P. Ruszczyński. *Lectures on stochastic programming: modeling and theory*. Society for Industrial Mathematics, 2009.

[27] L. Wolsey. *Integer Programming*. Wiley Interscience, 1998.