Computational complexity of stochastic programs

A. Shapiro
School of Industrial and Systems Engineering,
Georgia Institute of Technology,
Atlanta, Georgia 30332-0205, USA

East Coast Optimization Meeting 2019

Consider optimization problem

$$\min_{x \in X} \Big\{ f(x) = \mathbb{E}[F(x, \xi)] \Big\},$$

where $X \subset \mathbb{R}^n$, $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ and ξ is an m-dimensional random vector. In case of two-stage linear stochastic programming with recourse, $X = \{x \in \mathbb{R}^n_+ : Ax = b\}$ and $F(x, \xi)$ is the first stage cost $c^\top x$ plus the optimal value of the second stage problem

$$\min_{y \in \mathbb{R}^m} q^\top y \ \ subject\ to\ Tx + Wy = h,\ y \geq 0,$$

with ξ formed from random components of q, T, W, h.

For fixed $x \in X$ the expectation $\mathbb{E}[F(x, \xi)]$ is given by the integral

$$\mathbb{E}[F(x, \xi)] = \int F(x, z) dP(z),$$

where P is the probability distribution of ξ.

A standard approach to solving such stochastic programs is to discretize distribution P, i.e., to construct scenarios ξ_k, $k = 1, ..., K$, with assigned probabilities $p_k > 0$, and hence to approximate $\mathbb{E}[F(x, \xi)]$ by $\sum_{k=1}^{K} p_k F(x, \xi_k)$. In the two-stage linear case this leads to the linear program

$$\begin{array}{ll} \min\limits_{x, y_1, ..., y_K} & c^\top x + \sum_{k=1}^{K} p_k q_k^\top y_k \\ s.t. & T_k x + W_k y_k = h_k,\ k = 1, ..., K, \\ & Ax = b,\ x \geq 0,\ y_k \geq 0,\ k = 1, ..., K. \end{array}$$

In order to have an accurate approximation of the 'true' distribution P the number K of required scenarios typically growths exponentially with dimension m.

Computational complexity of solving two-stage linear stochastic programs (deterministic point of view): *the approximate solutions, with a sufficiently high accuracy, of linear two-stage stochastic programs with fixed recourse are #P-hard even if the random problem data is governed by independent uniform distributions* (Dyer and Stougie, 2006, Hanasusanto, Kuhn and Wiesemann, 2016).

Sample complexity of solving stochastic programs

Generate a sample ξ^j, $j = 1, ..., N$, of random vector ξ and approximate the expectation $\mathbb{E}[F(x, \xi)]$ by the respective sample average. This leads to the following so-called Sample Average Approximation (SAA) of the 'true' problem

$$\min_{x \in X} \left\{ \widehat{f}_N(x) = \frac{1}{N} \sum_{j=1}^{N} F(x, \xi^j) \right\}.$$

Slow convergence of the sample average $\widehat{f}_N(x)$ to the expectation $f(x)$. By the Central Limit Theorem, for fixed x the error

$$\widehat{f}_N(x) - f(x) = O_p(N^{-1/2}).$$

Let $\widehat{v}_N$ be the optimal value of the SAA problem and v^0 and $\mathcal{S}^0$ be the optimal value and set of optimal solutions of the true problem. Then under mild regularity conditions

$$\widehat{v}_N = \min_{x \in \mathcal{S}^0} \widehat{f}_N(x) + o_p(N^{-1/2}).$$

In particular, if $\mathcal{S}^0 = \{x^0\}$, then

$$N^{1/2}[\widehat{v}_N - v^0] \Rightarrow N(0, \sigma^2(x^0))$$

(Shapiro, 1991).

Large Deviations type bounds. Suppose that: $\varepsilon > \delta \geq 0$, the set X is of finite diameter D, there is a constant $\sigma > 0$ such that

$$M_{x',x}(t) \leq \exp\{\sigma^2 t^2/2\}, \quad t \in \mathbb{R}, \ x', x \in X,$$

where $M_{x',x}(t)$ is the moment generating function of the random variable $F(x',\xi) - F(x,\xi) - \mathbb{E}[F(x',\xi) - F(x,\xi)]$, there exists $\kappa(\xi)$ such that its moment generating function is finite valued in a neighborhood of zero and

$$\left|F(x',\xi) - F(x,\xi)\right| \leq \kappa(\xi)\|x' - x\|, \quad x', x \in X \text{ and } a.e.\ \xi.$$

Then for $L = \mathbb{E}[\kappa(\xi)]$ and sample size

$$N \geq \frac{8\sigma^2}{(\varepsilon - \delta)^2}\left[n \log\left(\frac{O(1)DL}{(\varepsilon - \delta)^2}\right) + \log\left(\frac{2}{\alpha}\right)\right],$$

we are guaranteed that $\Pr\left(\widehat{\mathcal{S}}_N^{\delta} \subset \mathcal{S}^{\varepsilon}\right) \geq 1 - \alpha$. Here $\widehat{\mathcal{S}}_N^{\delta}$ and $\mathcal{S}^{\varepsilon}$ are the sets of δ-optimal and ε-optimal solutions of the SAA and true problems respectively.

Stochastic Approximation (SA) approach. Suppose that the problem is convex, i.e., the feasible set X is convex and $F(\cdot, \xi)$ is convex for a.e. ξ. Classical SA algorithm

$$x_{j+1} = \Pi_X(x_j - \gamma_j G(x_j, \xi^j)),$$

where $G(x, \xi) \in \partial_x F(x, \xi)$ is a calculated (sub)gradient, Π_X is the orthogonal (Euclidean) projection onto X and $\gamma_j = \theta/j$. Theoretical bound (assuming $f(\cdot)$ is *strongly convex and differentiable*)

$$\mathbb{E}[f(x_j) - v^0] = O(j^{-1}),$$

for an *optimal* choice of constant θ (recall that v^0 is the optimal value of the true problem). This algorithm is very sensitive to choice of θ.

Robust SA approach (B. Polyak, 1990, Nemirovski). Constant step size variant: fixed in advance sample size (number of iterations) N and step size $\gamma_j \equiv \gamma$, $j = 1, ..., N$: $\tilde{x}_N = \frac{1}{N}\sum_{j=1}^N x_j$. Theoretical bound

$$\mathbb{E}[f(\tilde{x}_N) - v^0] \leq \frac{D_X^2}{2\gamma N} + \frac{\gamma M^2}{2},$$

where $D_X = \max_{x \in X} \|x - x_1\|_2$ and $M^2 = \max_{x \in X} \mathbb{E}\|G(x, \xi)\|_2^2$. For optimal (up to factor θ) $\gamma = \frac{\theta D_X}{M\sqrt{N}}$ we have

$$\mathbb{E}\left[f(\tilde{x}_N) - v^0\right] \leq \frac{D_X M}{2\theta\sqrt{N}} + \frac{\theta D_X M}{2\sqrt{N}} \leq \frac{\kappa D_X M}{\sqrt{N}},$$

where $\kappa = \max\{\theta, \theta^{-1}\}$. By Markov inequality it follows that

$$\Pr\left\{f(\tilde{x}_N) - v^0 > \varepsilon\right\} \leq \frac{\kappa D_X M}{\varepsilon\sqrt{N}},$$

and hence to the sample size estimate $N \geq \frac{\kappa^2 D_X^2 M^2}{\varepsilon^2 \alpha^2}$.

Multistage stochastic programming. Let ξ_t be a random (stochastic) process. Denote $\xi_{[t]} := (\xi_1, .., \xi_t)$ the history of the process ξ_t up to time t. The values of the decision vector x_t, chosen at stage t, may depend on the information $\xi_{[t]}$ available up to time t, but not on the future observations. The decision process has the form

$$\text{decision}(x_0) \rightsquigarrow \text{observation}(\xi_1) \rightsquigarrow \text{decision}(x_1) \rightsquigarrow \\ ... \rightsquigarrow \text{observation}(\xi_T) \rightsquigarrow \text{decision}(x_T).$$

Risk neutral T-stage stochastic programming problem:

$$\begin{array}{ll} \min\limits_{x_1, x_2(\cdot), ..., x_T(\cdot)} & \mathbb{E}\Big[F_1(x_1) + F_2(x_2(\xi_{[2]}), \xi_2) + \cdots + F_T\left(x_T(\xi_{[T]}), \xi_T\right)\Big] \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t),\ t = 2, \ldots, T. \end{array}$$

In linear case, $F_t(x_t, \xi_t) := c_t^\top x_t$ and

$$\mathcal{X}_t(x_{t-1}, \xi_t) := \{x_t : B_t x_{t-1} + A_t x_t = b_t,\ x_t \geq 0\},\ \ t = 2, ..., T.$$

Optimization is performed over feasible policies (also called decision rules). A policy is a sequence of (measurable) functions $x_t = x_t(\xi_{[t]})$, $t = 1, ..., T$. Each $x_t(\xi_{[t]})$ is a function of the data process up to time t, this ensures the *nonanticipative* property of a considered policy.

If the number of realizations (scenarios) of the process ξ_t is finite, then the above (linear) problem can be written as one large (linear) programming problem.

Dynamic programming equations. Going recursively backwards in time. At stage T consider

$$Q_T(x_{T-1}, \xi_T) := \inf_{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)} F_T(x_T, \xi_T).$$

At stages $t = T-1, ..., 2$, consider

$$Q_t(x_{t-1}, \xi_{[t]}) := \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} F_t(x_t, \xi_t) + \underbrace{\mathbb{E}\left[Q_{t+1}(x_t, \xi_{[t+1]}) \middle| \xi_{[t]}\right]}_{\mathcal{Q}_{t+1}(x_t, \xi_{[t]})}.$$

At the first stage solve:

$$\operatorname*{Min}_{x_1 \in \mathcal{X}_1} F_1(x_1) + \mathbb{E}[Q_2(x_1, \xi_1)].$$

If the random process is stagewise independent, i.e., ξ_{t+1} is independent of $\xi_{[t]}$, then $\mathcal{Q}_{t+1}(x_t) = \mathbb{E}[Q_{t+1}(x_t, \xi_{t+1})]$ does not depend on $\xi_{[t]}$.

For example, suppose that the problem is linear and only the right hand side vectors b_t are random and can be modeled as a (first order) autoregressive process

$$b_t = \mu + \Phi b_{t-1} + \varepsilon_t,$$

where μ and Φ are (deterministic) vector and regression matrix, respectively, and the error process ε_t, $t = 1, ..., T$, is stagewise independent. The corresponding feasibility constraints can be written in terms of x_t and b_t as

$$B_t x_{t-1} + A_t x_t \leq b_t, \;\; \Phi b_{t-1} - b_t + \mu + \varepsilon_t = 0.$$

That is, in terms of decision variables (x_t, b_t) this becomes a linear multistage stochastic programming problem governed by the stagewise independent random process $\varepsilon_1, ..., \varepsilon_T$.

Discretization by Monte Carlo sampling Independent of each other random samples $\xi_t^j = (c_t^j, B_t^j, A_t^j, b_t^j)$, $j = 1, ..., N_t$, of respective ξ_t, $t = 2, ..., T$, are generated and the corresponding scenario tree is constructed by connecting every ancestor node at stage $t-1$ with the same set of children nodes $\xi_t^1, ..., \xi_t^{N_t}$. In that way the stagewise independence is preserved in the generated scenario tree. We refer to the constructed problem as the Sample Average Approximation (SAA) problem.

The total number of scenarios of the SAA problem is given by the product $\mathcal{N} = \prod_{t=2}^{T} N_t$ and quickly becomes astronomically large with increase of the number of stages even for moderate values of sample sizes N_t.

For $T = 3$, under certain regularity conditions, for $\varepsilon > 0$ and $\alpha \in (0, 1)$, and the sample sizes N_1 and N_2 satisfying

$$O(1)\left[\left(\frac{D_1 L_1}{\varepsilon}\right)^{n_1} \exp\left\{-\frac{O(1)N_1\varepsilon^2}{\sigma_1^2}\right\} + \left(\frac{D_2 L_2}{\varepsilon}\right)^{n_2} \exp\left\{-\frac{O(1)N_2\varepsilon^2}{\sigma_2^2}\right\}\right] \leq \alpha,$$

we have that any first-stage $\varepsilon/2$-optimal solution of the SAA problem is an ε-optimal first-stage solution of the true problem with probability at least $1 - \alpha$.

In particular, suppose that $N_1 = N_2$ and take $L := \max\{L_1, L_2\}$, $D := \max\{D_1, D_2\}$, $\sigma^2 := \max\{\sigma_1^2, \sigma_2^2\}$ and $n := \max\{n_1, n_2\}$. Then the required sample size $N_1 = N_2$:

$$N_1 \geq \frac{O(1)\sigma^2}{\varepsilon^2}\left[n \log\left(\frac{O(1)DL}{\varepsilon}\right) + \log\left(\frac{1}{\alpha}\right)\right],$$

with total number of scenarios $N = N_1^2$ (Shapiro, 2006).

If we measure computational complexity, of the "true" problem, in terms of the number of scenarios required to approximate true distribution of the random data process with a reasonable accuracy, the conclusion is rather pessimistic. In order for the optimal value and solutions of the SAA problem to converge to their true counterparts all sample sizes $N_2, ..., N_T$ should tend to infinity. Furthermore, available estimates of the sample sizes required for a first stage solution of the SAA problem to be ε-optimal for the true problem, with a given confidence (probability), sums up to a number of scenarios which grows as $O(\varepsilon^{-2(T-1)})$ with decrease of the error level $\varepsilon > 0$. This indicates that from the point of view of the number of scenarios, complexity of multistage programming problems grows exponentially with increase of the number of stages.

Curse of dimensionality

One of the main difficulties in solving the dynamic programming equations (of the SAA problem) is how to represent the cost-to-go functions in a computationally feasible way.

For dimension of x_t say greater than 3 and large number of stages it is practically impossible to solve the dynamic programming equations with high accuracy. Several alternatives were suggested in recent literature.

Approximate dynamic programming

Basic idea is to approximate the cost-to-go functions by a class of computationally manageable functions. Since functions $\mathcal{Q}_t(\cdot)$ are convex it is natural to approximate these functions by piecewise linear functions given by maximum of cutting hyperplanes.

Stochastic Dual Dynamic Programming (SDDP) method (Pereira and Pinto, 1991). For trial decisions $\bar{x}_t$, $t = 1, ..., T-1$, at the backward step of the SDDP algorithm, piecewise linear approximations $\mathfrak{Q}_t(\cdot)$ of the cost-to-go functions $\mathcal{Q}_t(\cdot)$ are constructed by solving problems

$$\operatorname*{Min}_{x_t \in \mathbb{R}^{n_t}} (c_t^j)^{\mathsf{T}} x_t + \mathfrak{Q}_{t+1}(x_t) \text{ s.t. } B_t^j \bar{x}_{t-1} + A_t^j x_t = b_t^j, \ x_t \geq 0,$$

$j = 1, ..., N_t$, and their duals, going backward in time $t = T, ..., 1$.

Denote by v^0 and $\widehat{v}_N$ the respective optimal values of the true and SAA problems.

By construction

$$\mathcal{Q}_t(\cdot) \geq \mathfrak{Q}_t(\cdot),\ t = 2, ..., T.$$

Therefore the optimal value of

$$\operatorname*{Min}_{x_1 \in \mathbb{R}^{n_1}} c_1^{\mathsf{T}} x_1 + \mathfrak{Q}_2(x_1) \ \text{ s.t. } A_1 x_1 = b_1,\ x_1 \geq 0$$

gives a lower bound for the optimal value $\widehat{v}_N$ of the SAA problem.

We also have that

$$v^0 \geq \mathbb{E}[\widehat{v}_N].$$

Therefore *on average* $\widehat{v}_N$ is also a lower bound for the optimal value of the true problem.

The approximate cost-to-go functions $\mathfrak{Q}_2, ..., \mathfrak{Q}_T$ and a feasible first stage solution $\bar{x}_1$ define a feasible policy. That is for a realization (sample path) $\xi_1, ..., \xi_T$ of the data process, $\bar{x}_t = \bar{x}_t(\xi_{[t]})$ are computed recursively in $t = 2, ..., T$ as a solution of

$$\underset{x_t \geq 0}{\text{Min}} \, c_t^\top x_t + \mathfrak{Q}_{t+1}(x_t) \text{ s.t. } B_t \bar{x}_{t-1} + A_t x_t = b_t.$$

In the *forward step* of the SDDP algorithm M sample paths (scenarios) are generated and the corresponding $\bar{x}_t$, $t = 2, ..., T$, are used as trial points in the next iteration of the backward step.

Note that the functions $\mathfrak{Q}_2, ..., \mathfrak{Q}_T$ and $\bar{x}_1$ define a feasible policy also for the *true* problem.

Convergence of the SDDP algorithm

It is possible to show that, under mild regularity conditions, the SDDP algorithm converges as the number of iterations go to infinity. That is, the computed optimal values and generated policies converge w.p.1 to their counterparts of the considered SAA problem. However, the convergence can be very slow and one should take such mathematical proofs very cautiously.

Moreover, it should be remembered that the SAA problem is just an approximation of the "true" problem. It is possible to show that, in a certain probabilistic sense, the SAA problem converges to the "true" problem as all sample sizes N_t, $t = 2, ..., T$, tend to infinity.

Stopping criteria

The policy value $\mathbb{E}\left[\sum_{t=1}^{T} c_t^{\mathsf{T}} \bar{x}_t(\xi_{[t]})\right]$ can be estimated in the forward step of the algorithm. That is, let $\xi_2^i, ..., \xi_T^i$, $i = 1, ..., M$, be sample paths (scenarios) generated at a current iteration of the forward step, and

$$\vartheta_i := \sum_{t=1}^{T} (c_t^i)^{\mathsf{T}} \bar{x}_t^i, \ i = 1, ..., M,$$

be the corresponding cost values. Then $\mathbb{E}[\vartheta_i] = \mathbb{E}\left[\sum_{t=1}^{T} c_t^{\mathsf{T}} \bar{x}_t(\xi_{[t]}^i)\right]$, and hence

$$\bar{\vartheta} = \frac{1}{M} \sum_{i=1}^{M} \vartheta_i$$

gives an unbiased estimate of the policy value.

Also

$$\widehat{\sigma}^2 = \frac{1}{M-1}\sum_{i=1}^{M}(\vartheta_i - \bar{\vartheta})^2$$

estimates variance of the sample $\vartheta_1, ..., \vartheta_M$. Hence

$$\bar{\vartheta} + z_\alpha \widehat{\sigma}/\sqrt{M}$$

gives an *upper* bound for the policy value with confidence of about $100(1-\alpha)\%$. Here z_α is the corresponding critical value.

At the same time this gives an upper bound for the optimal value of the corresponding multistage problem, SAA or the "true" problem depending from what data process the random scenarios were generated.

The Brazilian hydro power operation planning problem

The Brazilian power system generation is hydro dominated (about 75% of the installed capacity) and characterized by large reservoirs presenting multi-year regulation capability, arranged in complex cascades over several river basins. The hydro plants use store water in the reservoirs to produce energy in the future, replacing fuel costs from the thermal units. Since the water inflows depend on rainfalls, the amount of future inflows is uncertain and cannot be predicted with a high accuracy.

The purpose of hydrothermal system operation planning is to define an operation strategy which, for each stage of the planning period, given the system state at the beginning of the stage, produces generation targets for each plant.

The Brazilian hydro power operation planning problem is a multistage, large scale (more than 200 power plants, of which 141 are hydro plants), stochastic optimization problem. On a high level, planning is for 5 years on monthly basis together with 5 additional years to smooth out the end of horizon effect. This results in 120-stage stochastic programming problem. Four energy equivalent reservoirs are considered, one in each one of the four interconnected main regions, SE, S, N and NE. The resulting policy obtained with the aggregate representation can be further refined, so as to provide decisions for each of the hydro and thermal power plants.

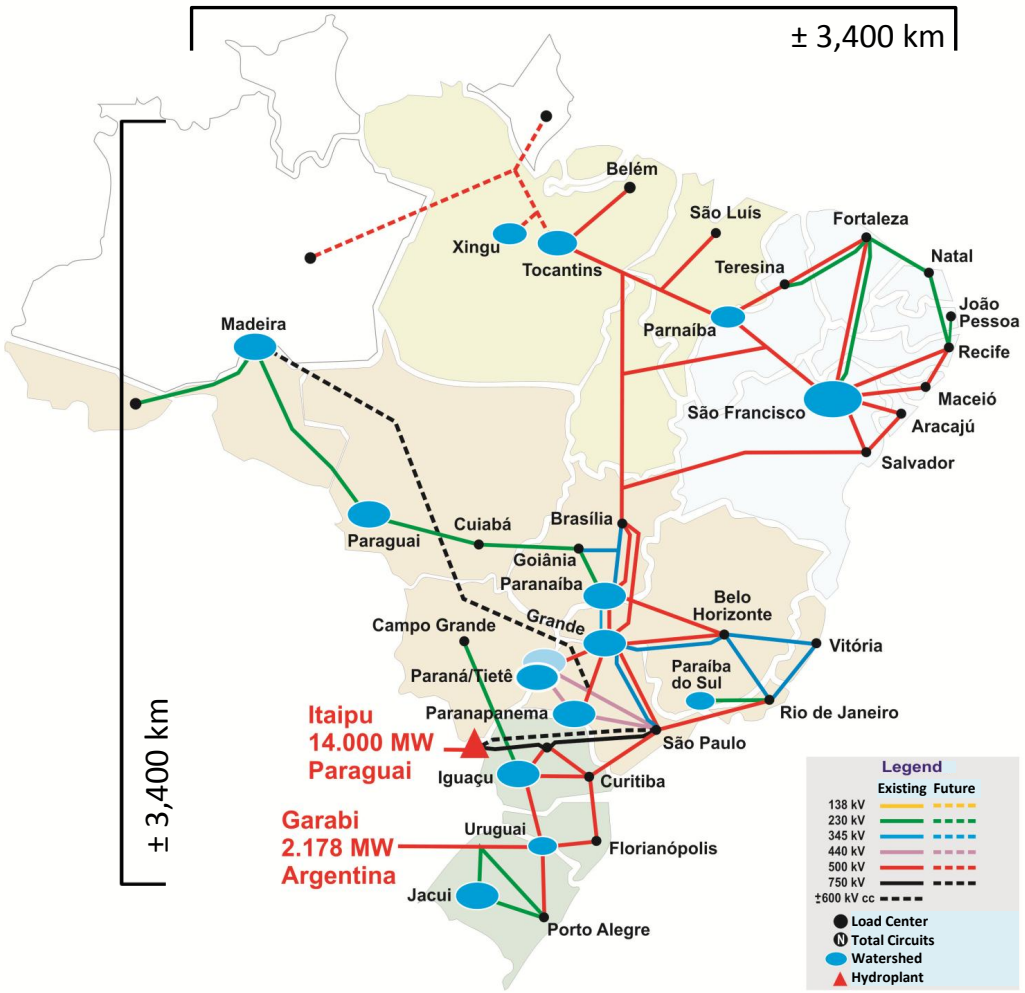
± 3,400 km
± 3,400 km
Belém
São Luís
Fortaleza
Xingu
Tocantins
Teresina
Natal
João Pessoa
Madeira
Parnaíba
Recife
Maceió
São Francisco
Aracajú
Salvador
Cuiabá
Brasília
Paraguai
Goiânia
Paranaíba
Belo Horizonte
Campo Grande
Grande
Vitória
Paraná/Tietê
Paraíba do Sul
Itaipu
14.000 MW
Paraguai
Paranapanema
Rio de Janeiro
São Paulo
Iguaçu
Curitiba
Garabi
2.178 MW
Argentina
Uruguai
Florianópolis
Jacui
Porto Alegre
Legend
Existing Future
138 kV
230 kV
345 kV
440 kV
500 kV
750 kV
±600 kV cc
Load Center
Total Circuits
Watershed
Hydroplant

Typical example of behavior of the lower and upper bounds produced by the SDDP algorithm for an SAA problem (Shapiro, Tekaya, Paulo da Costa, Pereira, 2013).

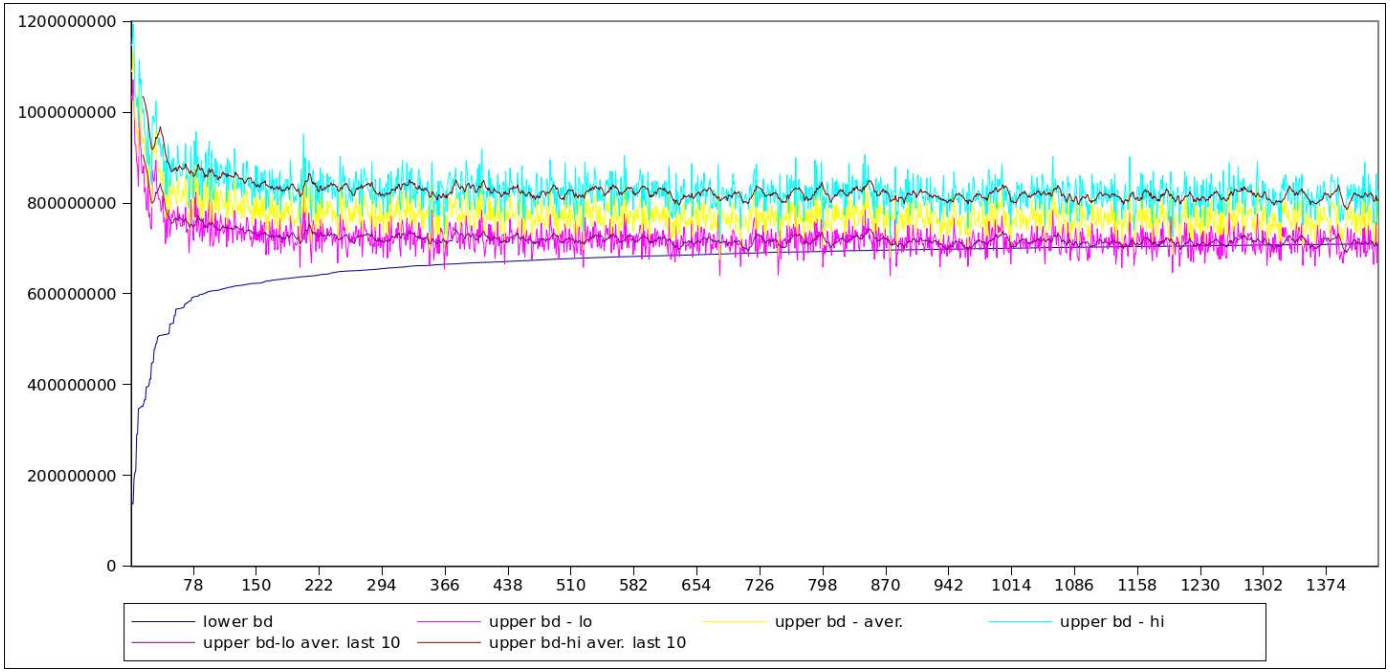

8 state variables, 120 stages, 1 cut per iteration

Theoretical analysis and numerical experiments indicate that computational complexity of the SDDP algorithm grows fast with increase of the number of state variables. The optimality gap jumped from 4% to 20% when the number of state variables was increased from 4 to 8 as a result of considering an autoregressive model.

Sensitivity to initial conditions

Individual stage costs for the risk neutral approach in two cases: all the reservoirs start at 25% or at 75% of the maximum capacity. The yellow curve denotes the 75% initial reservoir level and the dark green denotes the 25% initial level.

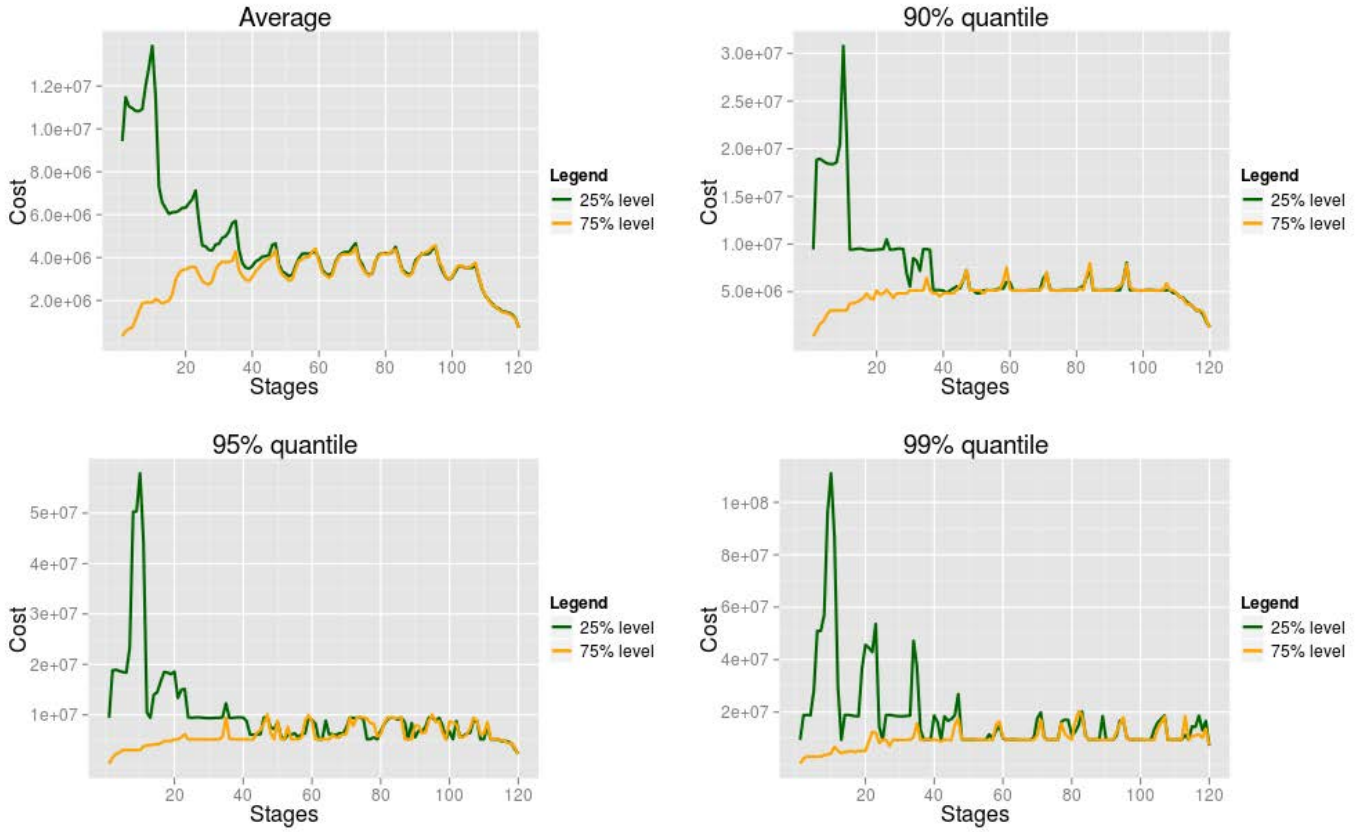
Average
Cost
1.2e+07
1.0e+07
8.0e+06
6.0e+06
4.0e+06
2.0e+06
20
40
60
80
100
120
Stages
Legend
25% level
75% level
90% quantile
Cost
3.0e+07
2.5e+07
2.0e+07
1.5e+07
1.0e+07
5.0e+06
20
40
60
80
100
120
Stages
Legend
25% level
75% level
95% quantile
Cost
5e+07
4e+07
3e+07
2e+07
1e+07
20
40
60
80
100
120
Stages
Legend
25% level
75% level
99% quantile
Cost
1e+08
8e+07
6e+07
4e+07
2e+07
20
40
60
80
100
120
Stages
Legend
25% level
75% level

Variability of SAA problems

Table shows the 95% confidence interval for the lower bound and average policy value at iteration 3000 over a sample of 20 SAA problems. Each of the policy value observations was computed using 2000 scenarios. The last 2 columns of the table shows the range divided by the average of the lower bound (where the range is the difference between the maximum and minimum observation) and the standard deviation divided by the average value. This problem has relatively low variability (approx. 4%) for both of the lower bound and the average policy value.

	95% C.I. left ($\times 10^9$)	Average ($\times 10^9$)	95% C.I. right ($\times 10^9$)	range average	sdev. average
Lower bound	22.290	22.695	23.100	15.92%	4.07%
Average policy	27.333	27.836	28.339	17.05%	4.12%

SAA variability for risk neutral SDDP

Risk averse approach

How to control risk, i.e., to reduce chances of extreme costs, at every stage of the time process.

Value-at-Risk of a random outcome (variable) Z at level $\alpha \in (0,1)$:

$$\text{V@R}_\alpha(Z) = \inf\{t : F_Z(t) \geq 1-\alpha\},$$

where $F_Z(t) = \Pr(Z \leq t)$ is the cdf of Z. That is, $\text{V@R}_\alpha(Z)$ is the $(1-\alpha)$-quantile of the distribution of Z.

Note that $\text{V@R}_\alpha(Z) \leq c$ is equivalent to $\Pr(Z > c) \leq \alpha$. Therefore it could be a natural approach to impose constraints (chance constraints) of $\text{V@R}_\alpha(Z) \leq c$ for $Z = cost$, chosen constant c and significance level α at every stage of the process.

There are two problems with such approach. It is difficult to handle chance constraints numerically and could lead to infeasibility problems.

Average Value-at-Risk (also called *Conditional Value-at-Risk*)

$$\mathsf{AV@R}_\alpha(Z) = \inf_{t\in\mathbb{R}} \left\{ t + \alpha^{-1}\mathbb{E}[Z-t]_+ \right\}$$

Note that the minimum in the above is attained at $t^* = \mathsf{V@R}_\alpha(Z)$. If the cdf $F_Z(z)$ is continuous, then

$$\mathsf{AV@R}_\alpha(Z) = \mathbb{E}\Big[Z | Z \geq \mathsf{V@R}_\alpha(Z) \Big].$$

It follows that $\mathsf{AV@R}_\alpha(Z) \geq \mathsf{V@R}_\alpha(Z)$. Therefore the constraint $\mathsf{AV@R}_\alpha(Z) \leq c$ is a conservative approximation of the chance constraint $\mathsf{V@R}_\alpha(Z) \leq c$.

In the problem of minimizing expected cost $\mathbb{E}[Z]$ subject to the constraint $\mathsf{AV@R}_\alpha(Z) \leq c$, we impose an infinite penalty for violating this constraint. This could result in infeasibility of the obtained problem. Instead we can impose a finite penalty and consider problem of minimization of $\mathbb{E}[Z] + \kappa \mathsf{AV@R}_\alpha(Z)$ for some constant $\kappa > 0$. Note that this is equivalent to minimization of $\rho(Z)$, where

$$\rho(Z) = (1 - \lambda)\mathbb{E}[Z] + \lambda \mathsf{AV@R}_\alpha(Z)$$

for $\lambda \in (0, 1)$ and $\kappa = \frac{\lambda}{1-\lambda}$.

This leads to the following (nested) formulation of risk averse multistage problem.

$$\begin{aligned}\operatorname*{Min}_{A_1x_1=b_1,x_1\geq 0} \; & c_1^{\mathsf{T}}x_1+\rho_{2|\xi_1}\Big[\inf_{\substack{B_2x_1+A_2x_2=b_2\\ x_2\geq 0}} c_2^{\mathsf{T}}x_2+\ldots \\ & +\rho_{T-1|\xi_{[T-2]}}\Big[\inf_{\substack{B_{T-1}x_{T-2}+A_{T-1}x_{T-1}=b_{T-1}\\ x_{T-1}\geq 0}} c_{T-1}^{\mathsf{T}}x_{T-1} \\ & +\rho_{T|\xi_{[T-1]}}\big[\inf_{\substack{B_Tx_{T-1}+A_Tx_T=b_T\\ x_T\geq 0}} c_T^{\mathsf{T}}x_T\big]\Big]\Big],\end{aligned}$$

with

$$\rho_{t|\xi_{[t]}}(\cdot):=(1-\lambda)\mathbb{E}_{|\xi_{[t]}}[\cdot]+\lambda\mathsf{AV@R}_{\alpha|\xi_{[t]}}(\cdot)$$

being conditional analogue of $\rho(\cdot)$.

We can write the risk averse multistage programming problem as

$$\begin{array}{cl} \underset{x_1, x_2(\cdot), \ldots, x_T(\cdot)}{\text{Min}} & \bar{\rho}\Big[F_1(x_1) + F_2(x_2(\xi_{[2]}), \xi_2) + \cdots + F_T\left(x_T(\xi_{[T]}), \xi_T\right)\Big] \\ \text{s.t.} & x_1 \in \mathcal{X}_1,\ x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t),\ t = 2, \ldots, T, \end{array}$$

where $F_t(x_t, \xi_t) = c_t^{\mathsf{T}} x_t$ and

$$\mathcal{X}_t(x_{t-1}, \xi_t) = \{x_t : B_t x_{t-1} + A_t x_t = b_t,\ x_t \geq 0\}.$$

$$\begin{aligned} \bar{\rho}(Z_1 + \ldots + Z_T) &= \rho_{|\xi_1}\Big(\rho_{|\xi_{[2]}}\Big(\cdots \rho_{|\xi_{[T-1]}}(Z_1 + \ldots + Z_T)\Big)\Big) \\ &= Z_1 + \rho_{|\xi_1}\Big(Z_2 + \rho_{|\xi_{[2]}}\Big(+ \cdots \rho_{|\xi_{[T-1]}}(Z_T)\Big)\Big) \end{aligned}$$

is the corresponding composite risk measure. The optimization is performed over (nonanticipative) policies $x_1, x_2(\xi_{[2]}), \ldots, x_T(\xi_{[T]})$ satisfying the feasibility constraints.

With some modifications the SDDP algorithm can be applied to the above multistage problem.

Remarks

Unfortunately there is no easy way for evaluating value of the risk objective of generated policies, and hence constructing a corresponding upper bound. Some suggestions were made in the recent literature. However, in larger problems the optimality gap (between the upper and lower bounds) never approaches zero in any realistic time. Therefore stopping criteria based on stabilization of the lower bound (and may be optimal solutions) could be reasonable. Also it should be remembered that there is no intuitive interpretation for the risk objective $\bar{\rho}(cost)$ of the total cost. Rather the goal is to control risk at every stage of the process.

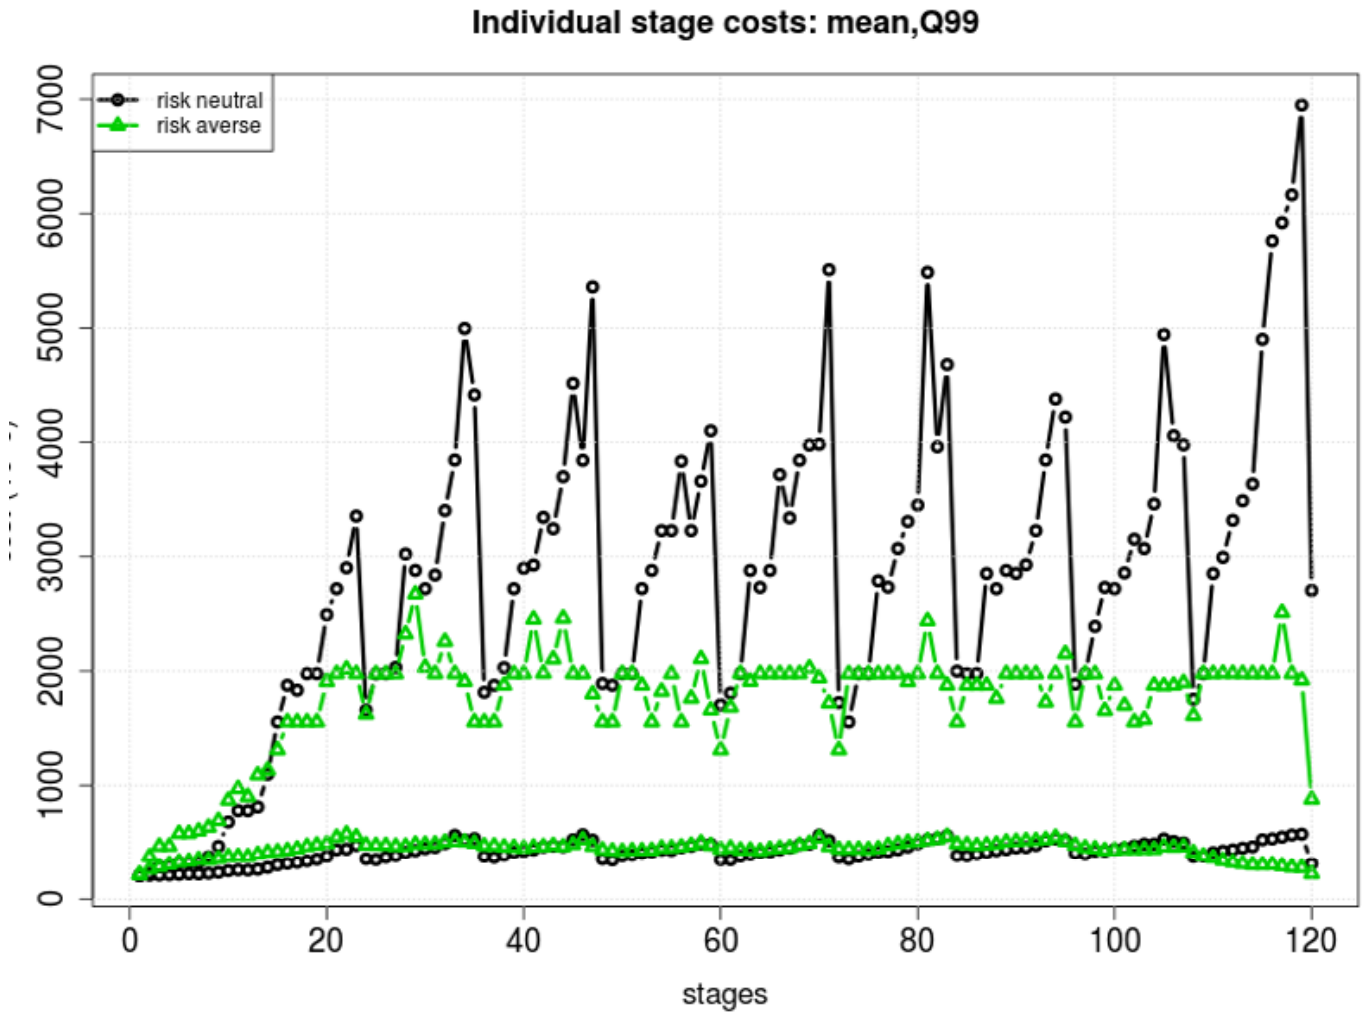
Individual stage costs: mean,Q99
risk neutral
risk averse
0
1000
2000
3000
4000
5000
6000
7000
0
20
40
60
80
100
120
stages