

Stochastic Dual Dynamic Programming

Operations Research

Anthony Papavasiliou

Contents [§10.4 of BL], [Pereira, 1991]

- 1 Recalling the Nested L-Shaped Decomposition
- 2 Drawbacks of Nested Decomposition and How to Overcome Them
- 3 Stochastic Dual Dynamic Programming (SDDP)
- 4 Example

Table of Contents

- 1 Recalling the Nested L-Shaped Decomposition
- 2 Drawbacks of Nested Decomposition and How to Overcome Them
- 3 Stochastic Dual Dynamic Programming (SDDP)
- 4 Example

The Nested L-Shaped Decomposition Subproblem

For each stage $t = 1, \dots, H - 1$, scenario $k = 1, \dots, \mathcal{K}^t$

$$\begin{aligned} NLDS(t, k) : & \min (c_k^t)^T x_k^t + \theta_h^t \\ \text{s.t. } & W^t x_k^t = h_k^t - T_k^{t-1} x_{a(k)}^{t-1}, (\pi_k^t) \\ & D_{k,j}^t x_k^t \geq d_{k,j}^t, j = 1, \dots, r_k^t, (\rho_k^t) \end{aligned} \quad (1)$$

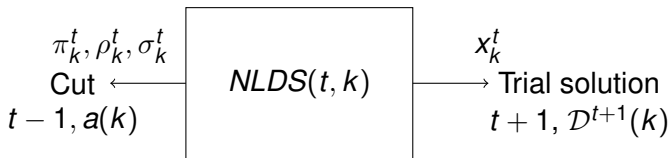
$$\begin{aligned} & E_{k,j}^t x_k^t + \theta_k^t \geq e_{k,j}^t, j = 1, \dots, s_k^t, (\sigma_k^t) \\ & x_k^t \geq 0 \end{aligned} \quad (2)$$

- \mathcal{K}^t : number of distinct scenarios at stage t
- $a(k)$: ancestor of scenario k at stage $t - 1$
- $x_{a(k)}^{t-1}$: current solution from $a(k)$
- Constraints (1): feasibility cuts
- Constraints (2): optimality cuts

Nested L-Shaped Method

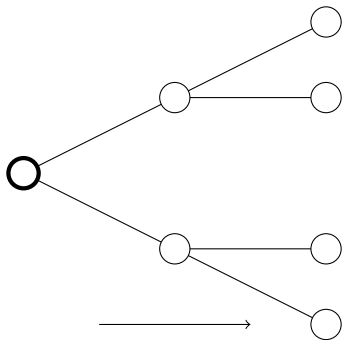
Building block: $NLDS(t, k)$: problem at stage t , scenario k

- Repeated application of the L-shaped method
- Variants depending on how we traverse the scenario tree



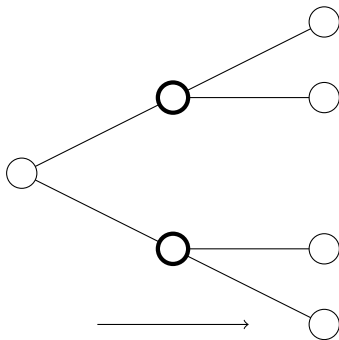
- $a(k)$: ancestor of scenario k
- $\mathcal{D}^{t+1}(k)$: descendants of scenario k in period $t + 1$

Example



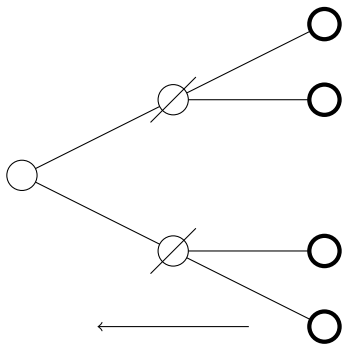
- Node: $(t = 1, k = 1)$
- Direction: forward
- Output: x_1^1

Example



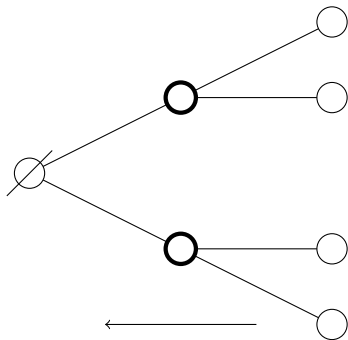
- Nodes: $(t = 2, k), k \in \{1, 2\}$
- Direction: forward
- Output: $x_k^2, k \in \{1, 2\}$

Example



- Nodes: $(t = 3, k), k \in \{1, 2, 3, 4\}$
- Direction: backward
- Output: $(\pi_k^3, \rho_k^3, \sigma_k^3), k \in \{1, 2, 3, 4\}$

Example



- Nodes: $(t = 2, k), k \in \{1, 2\}$
- Direction: backward
- Output: $(\pi_k^2, \rho_k^2, \sigma_k^2), k \in \{1, 2\}$

Feasibility Cuts

If $NLDS(t, k)$ is infeasible, solver returns $\pi_k^t, \rho_k^t \geq 0$

- $(\pi_k^t)^T (h_k^t - T_k^{t-1} x_{a(k)}^{t-1}) + (\rho_k^t)^T d_k^t > 0$
- $(\pi_k^t)^T W^t + (\rho_k^t)^T D_k^t \leq 0$

The following is a valid feasibility cut for $NLDS(t-1, a(k))$:

$$D_{a(k)}^{t-1} x \leq d_{a(k)}^{t-1}$$

where

$$\begin{aligned} D_{a(k)}^{t-1} &= (\pi_k^t)^T T_k^{t-1} \\ d_{a(k)}^{t-1} &= (\pi_k^t)^T h_k^T + (\rho_k^t)^T d_k^t \end{aligned}$$

Solve $NLDS(t, k)$ for $j = 1, \dots, \mathcal{K}^{t-1}$, then compute

$$E_j^{t-1} = \sum_{k \in \mathcal{D}^t(j)} \frac{p_k^t}{p_j^{t-1}} (\pi_k^t)^T T_k^{t-1}$$

$$e_j^{t-1} = \sum_{k \in \mathcal{D}^t(j)} \frac{p_k^t}{p_j^{t-1}} [(\pi_k^t)^T h_k^t + \sum_{i=1}^{r_k^t} \rho_{ki}^t d_{ki}^t + \sum_{i=1}^{s_k^t} \sigma_{ki}^t e_{ki}^t]$$

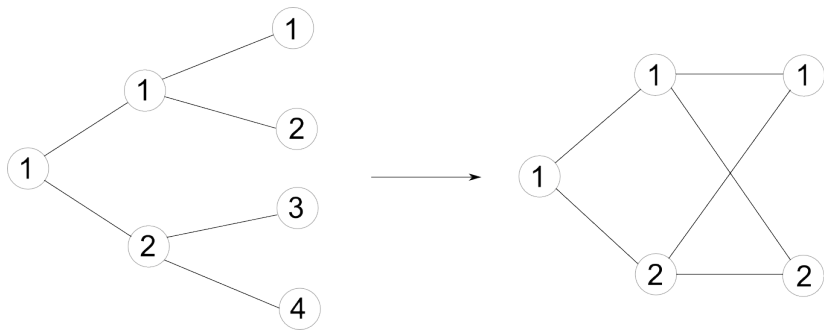
$\mathcal{D}^t(j)$: period t descendants of a scenario j at period $t - 1$

Note: $\frac{p_k^t}{p_j^{t-1}} = p(k, t | j, t - 1)$

Table of Contents

- 1 Recalling the Nested L-Shaped Decomposition
- 2 Drawbacks of Nested Decomposition and How to Overcome Them
- 3 Stochastic Dual Dynamic Programming (SDDP)
- 4 Example

Recombining Scenario Tree



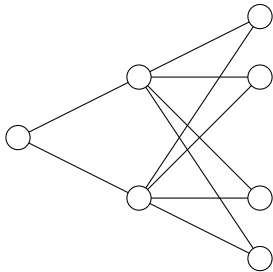
- When can we recombine nodes?
- When can we assign the same value function $V^{t+1}(x)$ to each node k of stage t ?

Nested Decomposition Is Non-Scalable

Assume

- H time steps, M^t discrete outcomes in each stage
- No infeasibility cuts

$$M^1 = 1 \quad M^2 = 2 \quad M^3 = 4$$



- Forward pass: $M^1 + M^1 \cdot M^2 + \dots = \sum_{t=1}^H \prod_{j=1}^t M_j$
- Backward pass: $\sum_{t=2}^{H-1} \prod_{j=1}^t M_j$

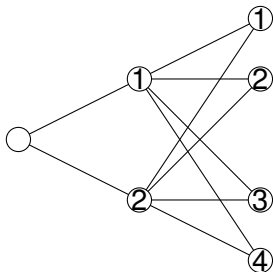
Was Nested Decomposition any Good?

Alternative to nested decomposition is extended form

- Extended form will not even load in memory
- Nested decomposition will load in memory, but will not terminate (for large problems)

Nested Decomposition lays the foundations for SDDP

Enumerating Versus Simulating



- Enumeration: $\{(1, 1), (1, 2), (1, 3), (1, 4), (2, 1), (2, 2), (2, 3), (2, 4)\}$
- Simulation (with 3 samples): $\{(1, 3), (2, 1), (1, 4)\}$

Making Nested Decomposition Scalable

Solution for forward pass

- In the forward pass, we *simulate* instead of *enumerating*
- This results in a probabilistic upper bound / termination criterion

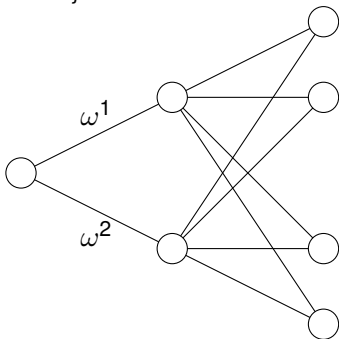
Solutions for backward pass

- In the backward pass, we share cuts among nodes of the same time period
- This requires an assumption of **serial independence**

Serial Independence

Serial independence: probability of realization ξ_j^t is constant from all possible $(t - 1)$ -stage scenarios

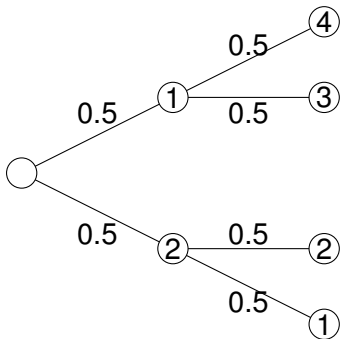
$$P(\xi_k^3 = c_k^3 | \xi_j^2 = c_j^2) = p_k, j \in 1, \dots, M^2, k \in 1, \dots, M^3$$



Problem is identical from $t = 2$ whether we observe ω^1 or ω^2

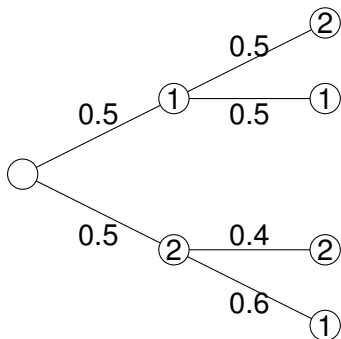
Example of Serial Independence (I)

- Value in circles: realization of ξ_k^t
- Value in edges: transition probabilities



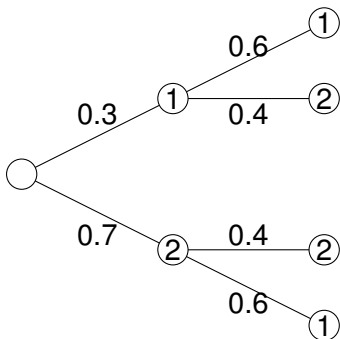
Is this tree serially independent?

Example of Serial Independence (II)



Is this tree serially independent?

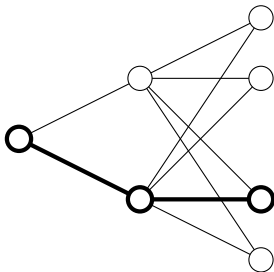
Example of Serial Independence (III)



Is this tree serially independent?

Implications for Forward Pass

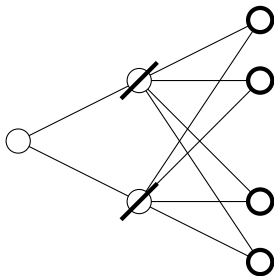
At each forward pass we solve $H - 1$ NLDS problems



For K Monte Carlo simulations, we solve $1 + K \cdot (H - 1)$ linear programs

Implications for Backward Pass

Serial independence implies same value function for all nodes of stage $t \Rightarrow$ cut sharing



For a given trial sequence x_k^t , we solve $\sum_{t=2}^H M^t$ linear programs, for K trial sequences we solve $K \sum_{t=2}^H M^t$ linear programs

Serial Independence is Helpful, Not Necessary

We can use dual multipliers in stage $t + 1$ for cuts in stage t even *without* serial independence

However, each node in stage t has a different value function

- More memory
- More optimality cuts needed because we are approximating more value functions

With serial independence, we can

- get rid of the scenario tree
- work with continuous distribution of ξ^t

Table of Contents

- 1 Recalling the Nested L-Shaped Decomposition
- 2 Drawbacks of Nested Decomposition and How to Overcome Them
- 3 Stochastic Dual Dynamic Programming (SDDP)**
- 4 Example

- Solve $NLDS(1, 1)$. Let x_1^1 be the optimal solution. Initialize $\hat{x}_i^1 = x_1^1$ for $i = 1, \dots, K$
- Repeat for $t = 2, \dots, H, i = 1, \dots, K$
 - Sample a vector h_i^t from the set $h_k^t, k = 1, \dots, M^t$
 - Solve the $NLDS(t, i)$ with trial decision \hat{x}_i^{t-1}
 - Store the optimal solution as \hat{x}_i^t

SDDP Backward Pass

- Repeat for $t = H, H - 1, \dots, 2$
 - Repeat for $i = 1, \dots, K$
 - Repeat for $k = 1, \dots, M^t$
Solve $NLDS(t, k)$ with trial decision \hat{x}_i^{t-1}
 - Compute

$$E^{t-1} = \sum_{k=1}^{M^t} p_k^t \pi_{k,i}^t T_k^{t-1}, e^{t-1} = \sum_{k=1}^{M^t} p_k^t (\pi_{k,i}^t h_k^t + \sigma_{k,i}^t e_k^t)$$

- Add the optimality cut

$$E^{t-1}x + \theta \geq e^{t-1}$$

to every $NLDS(t-1, k)$, $k = 1, \dots, M^{t-1}$

Central Limit Theorem

Suppose $\{X_1, X_2, \dots\}$ is a sequence of independent identically distributed random variables with $\mathbb{E}[X_i] = \mu$ and $\text{Var}[X_i] = \sigma^2 < \infty$. Then

$$\sqrt{n} \left(\left(\frac{1}{n} \sum_{i=1}^n X_i \right) - \mu \right) \xrightarrow{d} N(0, \sigma^2).$$

Probabilistic Upper Bound

Suppose we draw a sample k of $(\xi_k^t), t = 1, \dots, H$ and solve $NLDS(t, k)$ for $t = 1, \dots, H$

- This gives us a vector $x_k^t, t = 1, \dots, H$
- We can compute a cost for this vector $z_k = \sum_{t=H} c_k^t x_k^t$
- If we repeat this K times, we get a distribution of independent, identically distributed costs $z_k, k = 1, \dots, K$
- By the Central Limit Theorem, $\bar{z} = \frac{1}{K} \sum_{k=1}^K z_k$ converges to a Gaussian with standard deviation estimated by

$$\sigma = \sqrt{\left(\frac{1}{K^2}\right) \sum_{k=1}^K (\bar{z} - z_k)^2}$$

- Each $(x_k^t, t = 1, \dots, H)$ is feasible but not necessarily optimal, so \hat{z}_K is an estimate of an upper bound

Bounds and Termination Criterion

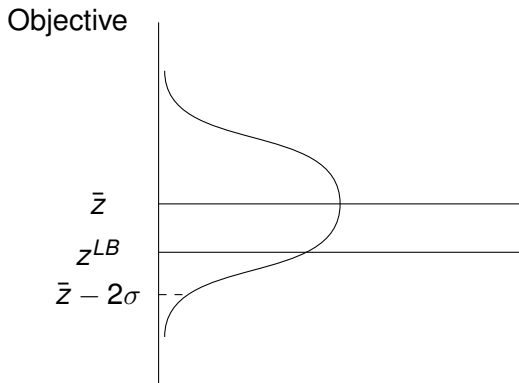
After solving NLDS(1, 1) in a forward pass we can compute a lower bound z^{LB} as the objective function value of NLDS(1, 1)

After completing a forward pass, we can compute

$$\begin{aligned}z_k &= \sum_{t=1}^H c_k^t \hat{x}_k^t \\ \bar{z} &= \frac{1}{K} \sum_{k=1}^K z_k \\ \sigma &= \sqrt{\frac{1}{K^2} \sum_{k=1}^K (z_k - \bar{z})^2}\end{aligned}$$

Terminate if $z^{LB} \in (\bar{z} - 2\sigma, \bar{z} + 2\sigma)$, which is the 95.4% confidence interval of \bar{z}

Graphical Illustration of Termination Criterion



Size of Monte Carlo Sample

How can we ensure 1% optimality gap with 95.4% confidence?

- Choose K such that $2\sigma \simeq 0.01 \cdot \bar{z}$
- Mean and variance depend (asymptotically) on the statistical properties of the process, not K

$$\bar{z} = \frac{1}{K} \sum_{k=1}^K z_k$$
$$s = \sqrt{\frac{1}{K} \sum_{k=1}^K (z_k - \bar{z})^2} \Rightarrow \sigma = \frac{1}{\sqrt{K}} s$$

- Set

$$K \simeq \left(\frac{2 \cdot s}{0.01 \cdot \bar{z}} \right)^2$$

Full SDDP Algorithm

- Initialize: $\bar{z} = \infty, \sigma = 0$
- Forward pass, store z^{LB} and \bar{z} . If $z^{LB} \in (\bar{z} - 2\sigma, \bar{z} + 2\sigma)$ terminate, else go to backward pass
- Backward pass
- Go to forward pass

Table of Contents

- 1 Recalling the Nested L-Shaped Decomposition
- 2 Drawbacks of Nested Decomposition and How to Overcome Them
- 3 Stochastic Dual Dynamic Programming (SDDP)
- 4 Example

Example

Consider the following problem

- Produce air conditioners for 3 months
- 200 units/month at 100 \$/unit
- Overtime costs 300 \$/unit
- Known demand of 100 units for period 1
- Equally likely demand, 100 or 300 units, for periods 2, 3
- Storage cost is 50 \$/unit
- All demand must be met

- x_k^t : regular production
- y_k^t : number of stored units
- w_k^t : overtime production
- d_k^t : demand

What does the scenario tree look like?

Extended Form

$$\min x^1 + 3w^1 + 0.5y^1 + \sum_{k=1}^2 p_k^2(x_k^2 + 3w_k^2 + 0.5y_k^2) +$$

$$\sum_{k=1}^4 p_k^3(x_k^3 + 3w_k^3)$$

$$\text{s.t. } x^1 \leq 2$$

$$x^1 + w^1 - y^1 = 1$$

$$y^1 + x_k^2 + w_k^2 - y_k^2 = d_k^2$$

$$x_k^2 \leq 2, k = 1, 2$$

$$y_{a(k)}^2 + x_k^3 + w_k^3 - y_k^3 = d_k^3$$

$$x_k^3 \leq 2$$

$$x_k^t, w_k^t, y_k^t \geq 0, k = 1, \dots, \mathcal{K}^t, t = 1, 2, 3$$

Optimal solution:

- Stage 1: $x^1 = 2, y^1 = 1$
- Stage 2, scenario 1: $x_1^2 = 1, y_1^2 = 1$
- Stage 2, scenario 2: $x_2^2 = 2, y_2^2 = 0$
- Stage 3, scenario 1: $x_1^3 = 0$
- Stage 3, scenario 2: $x_2^3 = 2$
- Stage 3, scenario 3: $x_3^3 = 1$
- Stage 3, scenario 4: $x_4^3 = 2, l_4^3 = 1$

What is the cost for each path?

SDDP Upper Bound Computation

```
param CostRecord{1..MCCount, 1..IterationCount};

let {m in 1..MCCount, i in 1..IterationCount}
CostRecord[m, i] := sum{j in Decisions, t in 1..H}
c[j]*xTrialRecord[j, t, m, i];

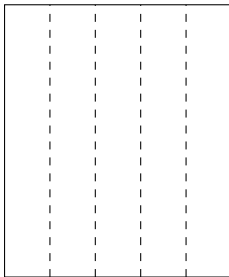
let {m in 1..MCCount} CostSamples[m] := CostRecord[m,
IterationCount];

let CostAverage := sum{m in 1..MCCount} CostSamples[m]
/ MCCount;

let CostStDev := sqrt(sum{m in 1..MCCount}
(CostSamples[m] - CostAverage)^2 / MCCount^2);
```

Thinking About the Data

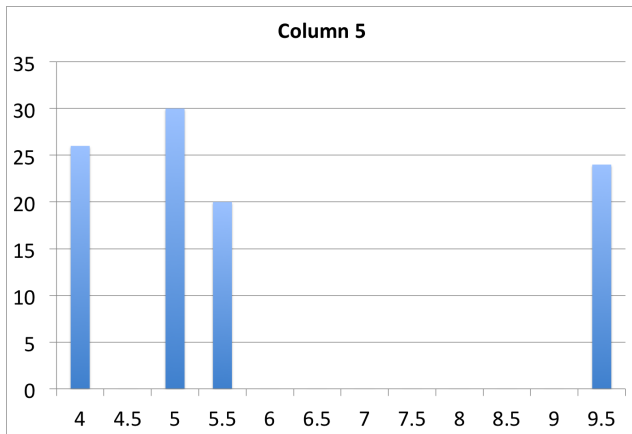
```
CostRecord{1..MCCount, 1..IterationCount}
```



- What is the distribution of each column?
- How does (k, i) entry depend on $(k + a, i)$ entry?
- Which column is more likely to have a lower average?
- Which data has a Gaussian distribution?

Distribution of Last Column

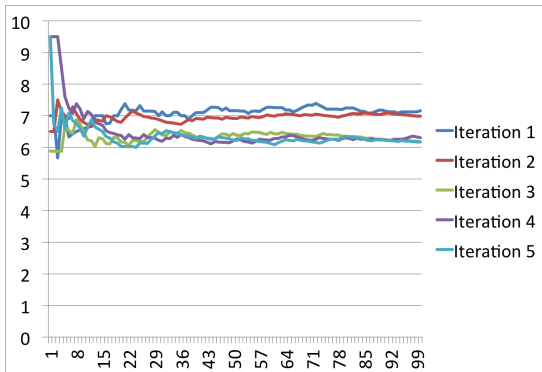
$$\bar{z} = 6.17, s = 2.02$$



Not a Gaussian distribution

Moving Average for 5 Iterations

Plot: $(N, \sum_{n=1}^N \frac{CostSamples_n}{N})$, MCCCount = 100, IterCount = 5



CostStDev = 0.2007: sample standard deviation of last column of CostRecord

Note: average cost decreases as iterations increase

How Many Monte Carlo Samples?

$$K \simeq \left(\frac{2 \cdot s}{0.01 \cdot \bar{z}} \right)^2 = \left(\frac{2 \cdot 2.02}{0.01 \cdot 6.17} \right)^2 = 4287$$