IE 495 – Lecture 11

The LShaped Method

Prof. Jeff Linderoth

February 19, 2003

Before We Begin

- HW#2
- \$300 $\rightarrow$ \$0
  - ◇ `http://www.unizh.ch/ior/Pages/Deutsch/Mitglieder/Kall/bib/ka-wal-94.pdf`
- Great source of recent papers in stochastic programming.
  - ◇ `http://www.speps.info`
  - ◇ Login: Password:

Outline

- Small amount of review
- The LShaped algorithm
  - Feasibility cuts
  - Formal description
  - Programming in AMPL
  - "Proof" of correctness
  - Multicut-method

LShaped Method

$$\min_{x \in \Re^n_+} \{c^T x + \mathcal{Q}(x) | Ax = b\}$$

- We know that a subgradient of $\mathcal{Q}(x)$ $\hat{x}$ looks like...

$$u = -\sum_{s \in S} p_s T_s^T \lambda_s^* \in \partial \mathcal{Q}(\hat{x}),$$

- where λ^* is an optimal dual solution to the recourse problem in scenario s:

$$\lambda_s^* = \arg\max_{\lambda} \{\lambda^T (h_s - T_s \hat{x}) : \lambda^T W \leq q\}.$$

LShaped Method

- So that by the subgradient inequality...

$$\mathcal{Q}(x) \geq \mathcal{Q}(\hat{x}) + u^T(x - \hat{x})$$

- In other words $\mathcal{Q}(\hat{x}) + u^T(x - \hat{x})$ is a supporting hyperplane of $\mathcal{Q}$ at $\hat{x}$.
- This insight is used to build up an (increasingly better) approximation of $\mathcal{Q}(x)$.

LShaped Method

- Imagine that we had L subgradients of $Q(x)$
- $u_1 \in \partial \mathcal{Q}(x_1), u_2 \in \partial \mathcal{Q}(x_2), \ldots u_l \in \partial \mathcal{Q}(x_l)$
- Then...

minimize

$$c^T x + \theta$$

subject to

$$\begin{aligned} Ax &= b \\ \theta &\geq \mathcal{Q}(x_l) + u_l^T (x - x_l) \qquad \forall l = 1, 2, \ldots L \end{aligned}$$

Good Ol' Farkas

- What if for some realization $\hat{\omega}$, we cannot solve the LP necessary to evaluate $\mathcal{Q}(\hat{x})$?
  - Then our problem does *not* have complete recourse or relatively complete recourse

$$Q(\hat{x}, \hat{\omega}) = \min_{y \in \Re^p_+} \{q^T y : Wy = h(\hat{\omega}) - T(\hat{\omega})\hat{x}\} = \infty$$

- By our favorite Theorem of the Alternative...
- $\{y \in \Re^p_+ | Wy = h - T\hat{x}\} = \emptyset$

$\Rightarrow$ $\exists \sigma \in \Re^m$ such that $W^T \sigma \leq 0$ and $(h - T\hat{x})^T \sigma > 0$.

Feasibility Cuts

- But for any *feasible* x, we know that there is at least one $y \geq 0$ such that $Wy = h - Tx$.
- Combining this with our Farkas knowledge gives...
  - ◇ $\sigma^T(h - Tx) = \sigma^T Wy \leq 0$
    - – $(\sigma^T W \leq 0, y \geq 0)$.
- This inequality $\sigma^T h \leq \sigma^T Tx$ must hold for all feasible x.
- It doesn't hold for our current iterate $\hat{x}$.
  - ◇ Remember Farkas: $(h - T\hat{x})^T \sigma > 0$

Feasibility Cuts

- So if we just knew the values for σ, we would be able to add the inequality $\sigma^T(h(\hat{\omega}) - T(\hat{\omega})x) \leq 0$ to our "master problem", and we would be assured of never getting this infeasible $\hat{x}$ again.
- Where do we get σ?
  - ◇ When the (primal) simplex method tells you that the problem is infeasible, then (if the dual is feasible), the dual is unbounded.
  - ◇ An LP is unbounded if there is some feasible direction (or "ray") that is improving. This "improving" ray is the σ we are looking for.
  - ◇ Most LP solvers will return this ray if asked.

Don't Believe Me

LP's (to justify previous)

LShaped Method – Step 0

- With θ_0 a lower bound for $\mathcal{Q}(x) = \sum_{s \in S} p_s Q(x, \omega)$,
- Let $\mathcal{B}_0 = \{\Re_n^+ \times \{\theta\} | Ax = b\}$
- Let $\mathcal{B}_1 = \{\Re_n^+ \times \{\theta\} | \theta \geq \theta_0\}$

LShaped Method – Step 1

- Solve the *master problem*:

$$\min\{c^T x + \theta | (x, \theta) \in \mathcal{B}_0 \cap \mathcal{B}_1\}$$

- yielding a solution $(\hat{x}, \hat{\theta})$.

Lshaped Method – Step 2

- Evaluate $\mathcal{Q}(\hat{x}) = \sum_{s \in S} p_s Q(\hat{x}, \omega_s)$.
- If $\mathcal{Q}(\hat{x}) = \infty$,
  - ◇ There is some $\hat{\omega}$ such that $Q(\hat{x}, \hat{\omega}) = \infty$
- Add a *feasibility cut*:
  - ◇ $\mathcal{B}_1 = \mathcal{B}_1 \cap \{(x, \theta) | \sigma^T (h(\hat{\omega}) - T(\hat{\omega})x) \leq 0\}$
- Go to 1.

Step 2 (cont.)

- If $\mathcal{Q}(\hat{x}) < \infty$, then you were able to solve all s scenario LP's (with corresponding dual optimal solutions λ_s^*), and you get a subgradient:

$$u = -\sum_{s \in S} p_s \lambda_s^* T_s \in \partial \mathcal{Q}(\hat{x})$$

- If $\mathcal{Q}(\hat{x}) \leq \hat{\theta}$.
  - ◇ Stop, $\hat{x}$ is an optimal solution.
  - ◇ (Our approximation is exact and minimized).
- Otherwise,
  - ◇ $\mathcal{B}_1 = \mathcal{B}_1 \cap \{(x, \theta) : \theta \geq \mathcal{Q}(\hat{x}) + u^T(x - \hat{x})\}$.
- Go to 1.

Programming in AMPL

minimize

$$x_1 + x_2$$

subject to

$$\begin{aligned} \omega_1 x_1 + x_2 &\geq 7 \\ \omega_2 x_1 + x_2 &\geq 4 \\ x_1 &\geq 0 \\ x_2 &\geq 0 \end{aligned}$$

Why One θ?

- A key idea in the LShaped method is to underestimate $\mathcal{Q}(x)$ by an auxiliary variable θ.
- We get the underestimate by the subgradient inequality.
- $\mathcal{Q}(x) = \sum_{s \in S} p_s Q(x, \omega_s)$
- For any scenario $s \in S$, $-T_s^T \lambda_s^* \in \partial Q(x, \omega_s)$, and some "fancy" convex analysis can show that

$$-\sum_{s \in S} p_s T_s^T \lambda_s^* \in \partial \mathcal{Q}(x)$$

$\Rightarrow$ We can equally well approximate (or underestimate) *each* $Q(x, \omega_s)$ by the auxilary variable(s) $\theta_s, s \in S$.

Multicut-LShaped Method – Step 0

- With θ_s^0 a lower bound for $Q(x, \omega_s)$,
- Let $\mathcal{B}_0 = \{\Re_n^+ \times \{\theta_1, \theta_2, \ldots, \theta_{|S|}\} | Ax = b\}$
- Let $\mathcal{B}_1 = \{\Re_n^+ \times \{\theta 1, \theta_2, \ldots, \theta_{|S|}\} | \theta_s \geq \theta_s^0 \quad \forall s \in S\}$

Multicut-LShaped Method – Step 1

- Solve the *master problem*:

$$\min\{c^T x + \sum_{s \in S} p_s \theta_s | (x, \theta_1, \theta_2, \ldots \theta_{|S|}) \in \mathcal{B}_0 \cap \mathcal{B}_1\}$$

- yielding a solution $(\hat{x}, \hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_{|S|})$.

Lshaped Method – Step 2

- Evaluate $\mathcal{Q}(\hat{x}) = \sum_{s \in S} p_s Q(\hat{x}, \omega_s)$.
- If $\mathcal{Q}(\hat{x}) = \infty$, which means that there is some $\hat{\omega}$ such that $Q(\hat{x}, \hat{\omega}) = \infty$, we add a *feasibility cut*:
  - ◇ $\mathcal{B}_1 = \mathcal{B}_1 \cap \{(x, \theta) | \sigma^T (h(\hat{\omega}) - T(\hat{\omega})x) \leq 0\}$
  - ◇ (Note that the inequality has no terms in θ_s – it is the same inequality as the LShaped method
- Go to 1.

Step 2 (cont.)

- If $\mathcal{Q}(\hat{x}) < \infty$, then you were able to solve all s scenario LP's (with corresponding dual optimal solution λ_s^*), and you get subgradients:

$$u = -T_s^T \lambda_s^* \in \partial Q(\hat{x}, \omega)$$

- If $Q(\hat{x}, \omega_s) \leq \theta_s \forall s \in S$, Stop. $\hat{x}$ is optimal.
- If $Q(\hat{x}, \omega_s) > \theta_s$
  - ◇ $\mathcal{B}_1 = \mathcal{B}_1 \cap \{(x, \theta_1, \theta_2, \ldots \theta_{|S|}) : \theta_s \geq Q(\hat{x}, \omega_s) + u^T(x - \hat{x})$.
- Go to 1.

A Whole Spectrum

- So far we have given an algorithms that give one cut per master iteration and $|S|$ cuts (potentially) per master iteration. We can do anything inbetween...
- Partition the scenarios into C "clusters" $\mathcal{S}_1,\ \mathcal{S}_2, \ldots \mathcal{S}_C$.

$$\mathcal{Q}_{[\mathcal{S}_k]}(x) = \sum_{s \in S_k} p_s Q(x, \omega_s)$$

The "Chunked" multicut method

$$\mathcal{Q}(x) = \sum_{k=1}^{C} \mathcal{Q}_{[\mathcal{S}_k]}(x).$$

$$\eta = \sum_{s \in S_k} p_s T_s^T \lambda_s^* \in \partial \mathcal{Q}_{[\mathcal{S}_k]}(x)$$

- We can do the same thing, just approximating $\mathcal{Q}_{[\mathcal{S}_k]}(x)$ by the subgradient inequalities.

Next time

- More LShaped...
  - Correctness/Convergence
  - Bunching
  - Regularizing the LShaped method
  - Parallelizing the LShaped method
  - Hand out a couple papers, and then that's it on LShaped for now.