Linear Programming

David C. Parkes

Division of Engineering and Applied Science,

Harvard University

CS 286r–Spring 2002

Introduction

- LP is the problem of optimizing a linear function subject to linear inequality constraints.
- Provides the foundation for the theory and computational methods of mathematical optimization.
- Today (e.g. Vazirani 02) provides a unified framework for the construction of approximation algorithms to combinatorial optimization problems.

In addition, for our purposes, LP duality provides an **important connection via primal-dual algorithms to the design of iterative auctions**.

A Simple Example

Consider the following minimization problem:

$$\begin{aligned} &\min 7x_1 + x_2 + 5x_3 \\ \text{s.t.} \quad & x_1 - x_2 + 3x_3 \geq 10 \\ & 5x_1 + 2x_2 - x_3 \geq 6 \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

$x = (7/4, 0, 11/4)$; value 26 optimal.

Def. A *feasible* solution satisifes the constraints. [e.g. $x = (4, 4, 4)$]. A feasible problem has a feasible solution.

Def. The *optimal solution* is the solution with the minimal objective value from all feasible solutions.

Def. An *unbounded problem* has a feasible solution with unbounded value.

Computing a Good Lower-Bound

(a) can take the first bound. Since $7x_1 + x_2 + 5x_3 \geq x_1 - x_2 + 3x_3 \geq 10$, then 10 is one lower-bound;

(b) can take the sum of the first two bounds. Since $7x_1 + x_2 + 5x_3 \geq (x_1 - x_2 + 3x_3) + (5x_1 + 2x_2 - x_3) \geq 16$, then 16 is a stronger lower-bound.

Generalize: determine positive multipliers, s.t. coeff. x_i in sum is dominated by the coeff. in the obj. fun; choose multipliers to maximize the RHS.

$$
\begin{aligned}
\max \quad & 10y_1 + 6y_2 \\
\text{s.t.} \quad & y_1 + 5y_2 \leq 7 \\
& -y_1 + 2y_2 \leq 1 \\
& 3y_1 - y_2 \leq 5 \\
& y_1, y_2 \geq 0
\end{aligned}
$$

This is the *dual program* to the first LP; $y = (2, 1)$, value 26 optimal.

Basic LP Duality Results

- Systematic method to construct the dual of any linear program; and dual of the dual is the primal.
- *Weak duality*. Every feasible dual solution provides a lower bound on the optimum value of the primal. (c.f. for primal) *[provide bounds]*
- *Optimality.* Feasible primal and dual solutions with the same value are both optimal.
- Strong duality. Feasible primal and dual solutions are optimal if and only if they have the same value. [*stopping criteria*]

Duality

$$\min \sum_{j=1}^{n} c_j x_j \qquad \text{[P]}$$

$$\text{s.t.} \sum_{j=1}^{n} a_{ij} x_j \geq b_i, \quad \forall i = 1, \ldots, m$$

$$x_j \geq 0, \quad \forall j = 1, \ldots, n$$

$$\max \sum_{i=1}^{m} b_i y_i \qquad \text{[D]}$$

$$\text{s.t.} \sum_{i=1}^{m} a_{ij} y_i \leq c_j, \quad \forall j = 1, \ldots, n$$

$$y_i \geq 0, \quad \forall i = 1, \ldots, m$$

Equivalently:

$$\min_x \{c^T x \ : \ Ax \geq b, x \geq 0\} \qquad \text{(P)}$$

$$\max_y \{b^T y \ : \ A^T y \leq c, y \geq 0\} \qquad \text{(D)}$$

Weak Duality

Thm. If $\overline{x}$ is a feasible solution to (P) and $\overline{y}$ is a feasible solution to (D) then

$$c^T\overline{x} \geq b^T\overline{y}$$

Proof. Follows immediately from feasibility of the solutions. Dual feasibility implies $c^T\overline{x} \geq (A^T\overline{y})^T\overline{x}$, and primal feasibility implies $b^T\overline{y} \leq (A\overline{x})^T\overline{y}$. Then, observe that $(A^T\overline{y})^T\overline{x} = (A\overline{x})^T\overline{y}$.

Optimality Property

Thm. If $\overline{x}$ is a feasible primal solution and $\overline{y}$ is a feasible dual solution, and $c^T x = b^T y$ then $\overline{x}$ and $\overline{y}$ are optimal.

Proof. Immediate from the weak duality property, since a dual feasible solution is a lower bound on the optimal primal solution and this bound is attained by the given feasible primal solution; c.f. for the dual problem.

Unboundedness Property

Thm. If the primal (dual) problem has an unbounded solution then the dual (primal) problem is infeasible.

Proof. Immediate from the weak duality property, this must be true for the primal since any feasible solution to the dual would provide a lower bound on the optimal primal solution; c.f. for the dual.

Thm. (P) has a finite optimal solution if and only if (D) does.

Example

primal feasible and unbounded:

$$
\begin{aligned}
\min -4y_1 + 2y_2 \\
\text{s.t.} \quad -y_1 + y_2 \geq 2 \\
-y_1 + y_2 \geq 1 \\
y_1, y_2 \geq 0
\end{aligned}
$$

dual infeasible:

$$
\begin{aligned}
\max 2x_1 + x_2 \\
\text{s.t.} \quad -x_1 - x_2 \leq -4 \\
x_1 + x_2 \leq 2 \\
x_1, x_2 \geq 0
\end{aligned}
$$

Strong-Duality Property

Thm. Feasible primal, $\overline{x}$, and dual, $\overline{y}$, solutions are optimal if and only if $c^T\overline{x} = b^T\overline{y}$.

Proof. The proof is constructive, via the *Simplex algorithm*.

Essentially, one shows that at the solution to the Simplex algorithm (which computes an optimal primal solution whenever the problem is feasible and unbounded) there is also enough information to construct a feasible dual solution with the same value.

Complementary Slackness

Thm. Feasible primal, $\overline{x}$ and dual, $\overline{y}$, are optimal if and only if

$$\overline{y}_i > 0 \Rightarrow \sum_{j=1}^{n} a_{ij}\overline{x}_j = b_i \qquad \text{(P-CS)}$$

$$\overline{x}_j > 0 \Rightarrow \sum_{i=1}^{m} a_{ij}\overline{y}_i = c_j \qquad \text{(D-CS)}$$

Proof. via the Strong duality property.

Auction Context

- Present an *iterative mechanism design* paradigm:
  – CS conditions can be *checked* by agents
  – dual and primal feasibility is easy; one set of CS conditions is easy
  – announce a dual and primal solution, and have agents confirm that the rest of the CS conditions hold
- The challenge is to make this incentive compatible
  – propose mechanisms that, (a) terminate with efficient allocations, and (b) agents have simple "truth-revealing" strategies
- Connections with VCG mechanism can be leveraged

$\Rightarrow$ primal-dual algorithms, that terminate with Vickrey outcomes, and elicit incremental preference elicitation from agents to compute solutions.

Vickrey Connections

- In special cases, the Vickrey payments are supported in an optimal dual solution
  – at the CE prices that maximize the total payoff to the agents
- **[centralized computation]** permits the Vickrey payments to be computed as the solution to a pair of LP formulations; instead of N formulations.
- **[iterative mechanism design]** adjust dynamically towards CE prices, asking agents for “best-response” in each round, terminate with Vickrey outcome
  – best-response is an *ex post* Nash eq.

$\Rightarrow$ one interesting challenge in comp. MD is to provide LP formulations (perhaps approximations) for combinatorial optimization problems that support Vickrey payments.

Example: Simple Allocation Problem

(a) from the Integer program

$$\max \sum_i v_i x_i \qquad \text{[IP}_1\text{]}$$

$$\text{s.t.} \quad \sum_i x_i \leq 1$$

$$x_i \in \{0, 1\}$$

(b) construct the LP relaxation

$$\max \sum_i v_i x_i \qquad \text{[LP}_1\text{]}$$

$$\text{s.t.} \quad \sum_i x_i \leq 1$$

$$x_i \geq 0$$

(has the "integrality property") (c) construct the dual problem

$$\min \pi \qquad \text{[DLP}_1\text{]}$$

$$\text{s.t.} \quad \pi \geq v_i, \quad \forall i$$

$$\pi \geq 0$$

and notice that optimal dual, $\pi^* = \max v_i \quad (*)$.

(d) construct complementary-slackness conditions:

$$x_i > 0 \Rightarrow \pi = v_i, \quad \forall i \tag{CS1}$$

$$\pi > 0 \Rightarrow \sum x_i = 1 \tag{CS2}$$

Consider whether a feasible dual solution, $\overline{\pi}$, and feasible primal solution, $\overline{x}$, are optimal.

(i) by (CS2), we must allocate the item to exactly one agent if $\overline{\pi} \geq \max_i v_i > 0$

(ii) by (CS1), and dual feasibility ($\pi \geq v_i, \quad \forall i$), the item must be allocated to the agent with the greatest value.

A Useful Reformulation

(a) Introduce redundant constraints into the primal, that provide *economic content.*

$$\max \sum_i v_i x_i \qquad \text{[LP}_2\text{]}$$

$$\text{s.t.} \quad \sum_i x_i \leq 1$$

$$x_i \leq 1$$

$$x_i \geq 0, \quad \forall i$$

(b) Construct the dual problem

$$\min_{p,\pi_i} p + \sum_i \pi_i \qquad \text{[DLP}_2\text{]}$$

$$\text{s.t.} \quad p + \pi_i \geq v_i, \quad \forall i$$

$$p, \pi_i \geq 0$$

Notice that $\pi_i^* = \max[v_i - p, 0] \quad (**)$, and that price, $p \geq 0$, defines a feasible dual solution.

Interpret p as *price*, and π_i as *agent surplus*.

(c) construct complementary-slackness conditions

$$x_i > 0 \Rightarrow p + \pi_i = v_i \tag{CS1}$$

$$p > 0 \Rightarrow \sum x_i = 1 \tag{CS2}$$

$$\pi_i > 0 \Rightarrow x_i = 1 \tag{CS3}$$

Question: *given prices $\overline{p}$, is there an $\overline{x}, \overline{\pi}_i$, that satisfies CS conditions?*

First, (dual feas.) $\pi \geq \max[0, v_i - p]$, so (CS3), if $v_i > p$ then $x_i = 1$, which with primal feas., requires that $p \geq v_{(2)}$. Second, by (CS2), whenever the price, p, is positive, the item must be allocated to some agent. Third, by (CS1), as long as $v_i \geq p$ then $\pi_i = v_i - p$ is dual feasible.

Primal-Dual Algorithms

1. Maintain a feasible dual, $\overline{y}$.

2. Compute a solution to a *restricted primal problem*, given the dual $\overline{y}$; find a $\overline{x}$ that is feasible and satisfies CS conditions.

3. Either: (a) compute some $\overline{x}$ that minimizes the "violation" of CS conditions; or (b) compute some $\overline{x}$ that satisfies CS and minimizes violation of feasibility.

4. Adjust dual solution, based on dual of restricted primal.

In the *auction context*, variation (a) is more useful, and *agent bids* provide suff. information to solve restricted primal and adjust the dual across rounds.

Primal-Dual Methodology

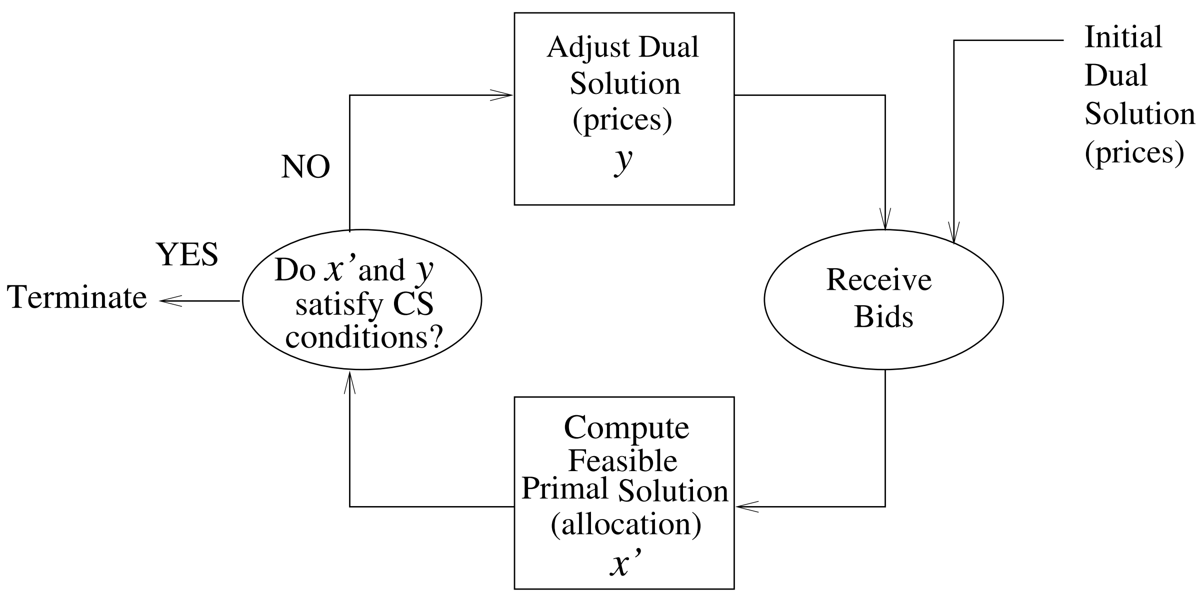

Example: English Auction

1. initial price, $p = 0$.
2. In each round, announce price p, and let $B \subseteq I$ denote set of agents that bid.
3. If $|B| > 1$, allocate item to some $i \in B$, and increase price to $p + \epsilon$.
4. If $|B| = 1$, then allocate to agent and terminate.

Terminates with CE price, $p^* = p_{\text{vick}}$, and the efficient allocation.

Proof: (a) Efficiency

Assume MBR, and show terminates with feasible primal and dual that satisfy CS conditions.

$$x_i > 0 \Rightarrow p + \pi_i = v_i \qquad \text{(CS1)}$$

$$p > 0 \Rightarrow \sum x_i = 1 \qquad \text{(CS2)}$$

$$\pi_i > 0 \Rightarrow x_i = 1 \qquad \text{(CS3)}$$

Lemma. If agents follow MBR, then (CS1) and (CS2) hold in any round.

Lemma. If agents follow MBR, the auction terminates.

Thm. If agents follow MBR, the auction terminates with the efficient allocation.

Proof. Termination implies just one bidder left, now (CS3) holds.

Proof: (b) Incentives

Now, show that MBR is an ex post *Nash eq. of the auction.*

Lemma. The minimal price across all optimal dual solutions equals the Vickrey payment.

Prop. If agents follow MBR, the auction terminates with the Vickrey payment.

Prop. MBR is an *ex post* Nash eq. of the auction.

Proof. W.o.l.g. suppose other agents follow truthful MBR. Consider agent 1. Suppose 1 follows some strategy s', and auction terminates with $x_{i'} = 1$ and price p'. By case analysis, show this outcome is always consistent with an outcome of the auction for MBR for *some* value $\hat{v}_1$. Therefore, agent 1 "selects" the Vickrey outcome for some (perhaps untruthful) bid, $\hat{v}_1$, and (weakly) prefers the actual Vickrey outcome.

Comp. Equil. Interpretation

Def. Price $\overline{p}$ and allocation $\overline{x}$ are in *competitive equilibrium* (CE) if $(\overline{p}, \overline{x})$ satisfy:

(1) the allocation, $\overline{x}$, maximizes the surplus of every bidder at the price, $\overline{p}$ [**from (CS1) and (CS3)**]

(2) the allocation, $\overline{x}$, maximizes the surplus of the seller at the price, $\overline{p}$ [**from (CS2)**]

Thm. An allocation is efficient if and only if there exist CE prices to support the allocation.

Note. This, with non-linear and non-anonymous prices, extends to *combinatorial allocation problems*; however not then the case that CE prices support Vickrey outcome.

Convex Polytopes

Def. A **polytope**, Q, is the convex hull of a finite set of points.

$$\mathrm{CH}(\{x^i\}) = \{\overline{x} \,:\, \overline{x} = \sum_i \alpha_i x_i, \sum_i \alpha_i = 1, \alpha_i \geq 0\}$$

Def. A **convex** polytope is one in which $\mathrm{CH}(Q) = Q$.

Def. Consider that the set of points at the *boundary* of a system of linear inequalities *and* consistent with all inequalities, define a *polyhedron*, P.

Lemma. A system of linear inequalities define a convex polytope.

Def. A **supporting hyperplane** of polyhedron, P, is a hyperplane H that touches P s.t. all of P is contained in a halfspace of H.

Def. A **face** is a subset of a polyhedron P that is the intersection of P with a supporting hyperplane of P.

Def. A **vertex** is a face that intersects P at exactly one place.

Simplex Method

Thm. If a LP has an optimal solution, then it has an optimal *extremal* solution. [from convexity, linearity]

but, there are *many extremal points*; e.g. the Upper Bound Thm. states that the number of vertices can be as large as $O(k^{\lfloor d/2 \rfloor})$, given $d+1$ variables and k constraints.

Simplex: "walk" the extremal vertices, terminate when there is no local improving direction. ["veritable workhorse" of linear programming.]

Although there is no known implementation of simplex with polynomial time complexity (in k, d), both empirical and probabilistic analyses indicate that *the number of iterations in the simplex method is just slightly more than linear in the dimension of the primal polyhedron.*

Ellipsoid Method

- Ellipsoid (Shor 70; Khachiyan 79)
  – given a "strong separation oracle" can determine a solution to a system of linear inequalities in polynomial number of calls.

Def. `Strong_Sep` $((Q, y)$ Given vector $y \in \mathbb{R}^d$, decide whether $y \in Q$, and if not find a hyperplane that separates y from Q.

But, disappointing gap in practice between theoretical promise and practical speed.

Interior Point Methods

Closed the gap between a method that was good in a theoretical sense but poor in practice (ellipsoid), and another that was good in practice (and average) but poor in theoretical worst-case sense (simplex).

- Karmarkar, 1984; closed the door with a “breathtaking new scaling algorithm”
- Identified with a class of nonlinear prog. methods known as “logarithmic barrier methods”
- Ideas is to move in middle of feasible region, in the direction of gradient of obj. function, but bias towards center to avoid “jamming” into corners
  – in high dimensions, corners make up most of the interior of a polyhedron.

Summary

- Large LP's can be solved quickly
- LP duality provides a mathematical framework to implement efficient allocations with incomplete information about agent preferences.
- Primal-dual algorithms provide a constructive method to design iterative mechanisms.
- Terminating with Vickrey payments inherits useful incentive properties from the VCG.