Linear Programming

Linear programming is the problem of minimizing or maximizing a linear function subject to linear constraints. The function being optimized is called the objective function.

Example $\quad \min \; 3x_1 + 4x_2$

subject to

$$\left.\begin{aligned} x_1 + 2x_2 &\geq 5 \\ x_1 + x_2 &\geq 4 \\ x_1, x_2 &\geq 0 \end{aligned}\right\}$$

Any solution that satisfies all the constraints is called a feasible solution.

In the above problem, the optimal solution is $x_1 = 3$ and $x_2 = 1$. So the optimal value of this LP is $3 \times 3 + 4 \times 1 = 13$.

A feasible solution that optimizes the objective function is called an optimal solution.

Is there any way we can quickly convince someone that 13 is the best answer here?

$$1 \cdot (x_1 + 2x_2) + 2 \cdot (x_1 + x_2) \geq 1 \cdot 5 + 2 \cdot 4 = 13$$

$$\text{Thus} \quad 3x_1 + 4x_2 \geq 13$$

Is it always possible to find such multipliers?

Let us look at another example.

$$\min \; 7x_1 + 3x_2$$

subject to

$$\begin{aligned} -x_1 + x_2 &\geq 10 \\ 2x_1 + 5x_2 &\geq 15 \\ x_1, x_2 &\geq 0 \end{aligned}$$

Let y_1 be the multiplier for the first inequality and let y_2 be the multiplier for the second inequality.

Let λ_1, λ_2 be the multipliers for $x_1 \geq 0$, $x_2 \geq 0$, respectively.

Is there a setting of $\underbrace{y_1, y_2, \lambda_1, \lambda_2}_{\text{non-negative}}$ such that

$$y_1(-x_1 + x_2) + y_2(2x_1 + 5x_2) + \lambda_1 x_1 + \lambda_2 x_2$$
$$= 7x_1 + 3x_2.$$

Then we know that $7x_1 + 3x_2 \geq 10y_1 + 15y_2$.

So the optimal value is bounded from below by $10y_1 + 15y_2$. We want to find the best setting of y_1, y_2 to get as good a lower bound as possible.

This problem of determining the best values of y_1, y_2 is the dual LP.

$$\max \; 10y_1 + 15y_2$$

subject to

$$-y_1 + 2y_2 + \lambda_1 = 7$$
$$y_1 + 5y_2 + \lambda_2 = 3$$
$$y_1, y_2, \lambda_1, \lambda_2 \geq 0$$

Another way of writing these constraints

$$-y_1 + 2y_2 \leq 7$$
$$y_1 + 5y_2 \leq 3$$
$$y_1, y_2 \geq 0$$

Primal LP

$$\min \sum_{j=1}^{k} c_j x_j$$

subject to

$$A \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix} \geq \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}$$

A: $n \times k$ matrix

$$x_1, \ldots, x_k \geq 0$$

Dual LP

$$\max \sum_{i=1}^{n} b_i y_i$$

subject to

$$A^T \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} \leq \begin{bmatrix} c_1 \\ \vdots \\ c_k \end{bmatrix}$$

A^T: Transpose of A

$$y_1, \ldots, y_n \geq 0$$

Weak duality theorem.

If $(x_1, \ldots, x_k)$ is primal feasible and $(y_1, \ldots, y_n)$ is dual feasible then $\sum_{j=1}^{k} c_j x_j \geq \sum_{i=1}^{n} b_i y_i$.

Proof. Consider $[y_1 \cdots y_n]\, A \begin{pmatrix} x_1 \\ \vdots \\ x_k \end{pmatrix}$.

This equals $\sum_{i=1}^{n} \sum_{j=1}^{k} a_{ij} x_j y_i$.

This sum can be written as

$$\sum_{i=1}^{n} \left(\sum_{j=1}^{k} a_{ij} x_j \right) y_i \geq \sum_{i=1}^{n} b_i y_i$$

(since $(x_1, \ldots, x_k)$ is primal feasible)

$$\parallel$$

$$\sum_{j=1}^{k} \left(\sum_{i=1}^{n} a_{ij} y_i \right) x_j \leq \sum_{j=1}^{k} c_j x_j \qquad \blacksquare$$

(since $(y_1, \ldots, y_n)$ is dual feasible)

An important question. Is there a gap between the primal optimal value and the dual optimal value?

Strong duality theorem. If $(x_1^*, \ldots, x_k^*)$ is an optimal solution for the primal LP then the dual LP also has an optimal solution $(y_1^*, \ldots, y_n^*)$ and $\sum_{j=1}^{k} c_j x_j^* = \sum_{i=1}^{n} b_i y_i^*$.

There are three possibilities for the primal LP:
1) Infeasible 2) Unbounded 3) Finite optimal value.

Similarly, there are three possibilities for the dual LP:

1') Infeasible 2') Unbounded 3') Finite optimal value.

1) $\Rightarrow$ the optimal value of the primal LP is the minimum of the empty set $= \infty$.

Similarly 1') $\Rightarrow$ the optimal value of the dual LP is the maximum of the empty set $= -\infty$.

2) $\Rightarrow$ primal OPT $= -\infty$ and 2') $\Rightarrow$ dual OPT $= \infty$.

Observe that 2) $\Rightarrow$ 1') by weak duality.

Similarly observe that 2') $\Rightarrow$ 1).

We can also have 1) & 1') occurring together. Then there is ∞ gap between primal OPT and dual OPT.

An example where both primal LP & dual LP are infeasible.

$$\begin{aligned} \min\ & x_1 - 2x_2 \\ \text{subject to}\ & \\ & x_1 - x_2 \geq 2 \\ & -x_1 + x_2 \geq -1 \\ & x_1, x_2 \geq 0 \end{aligned} \qquad\qquad \begin{aligned} \max\ & 2y_1 - y_2 \\ \text{subject to}\ & \\ & y_1 - y_2 \leq 1 \\ & -y_1 + y_2 \leq -2 \\ & y_1, y_2 \geq 0. \end{aligned}$$

Thus there are four possibilities in total:

1) and 1'); 2) and 1'); 1) and 2');

$\underbrace{\text{3) and 3')}}$

this is strong duality, i.e., 3) $\Rightarrow$ 3') and 3') $\Rightarrow$ 3).

A proof of strong duality

It will be convenient to model this as a 2-player game. The 2 players will be called the MIN player and the MAX player.

Game 1: The min player plays first. He chooses values for the variables that he controls.

The max player plays second. He chooses values for the variables that he controls.

Value of the game $= f(x_1, \ldots, x_k, y_1, \ldots, y_n, \lambda_1, \ldots, \lambda_k)$

where $x_1, \ldots, x_k$ are the variables that the min player controls

and $y_1, \ldots, y_n, \lambda_1, \ldots, \lambda_k$ are the variables that the max player controls.

We want the value of this game to be OPT_{primal} if the values $x_1^\circ, \ldots, x_k^\circ$ chosen by the min player are primal feasible. Else we want the value of the game to be ∞.

How do we enforce this?

The max player controls $n+k$ variables

$$\underbrace{y_1, \ldots, y_n}_{\text{one for each constraint of the primal LP}} + \underbrace{\lambda_1, \ldots, \lambda_k}_{\text{one variable for each constraint } x_i \geq 0}$$

all the y_i's and λ_j's have to be ≥ 0.

Let $f(x_1, \ldots, x_k, y_1, \ldots, y_n, \lambda_1, \ldots, \lambda_k)$

$$= c_1 x_1 + \cdots + c_k x_k + \sum_i y_i \left(b_i - \sum_j a_{ij} x_j\right) - \sum_j \lambda_j x_j$$

If the values $x_1^\circ, \ldots, x_k^\circ$ chosen by the min player are primal feasible then the max player sets all y_i's & λ_j's to 0.

If some constraint of the primal LP is violated then the max player makes the appropriate y_i or $\lambda_j \to \infty$. This makes the game $\to \infty$.

Claim 1. Value of Game 1 is OPT_{primal}.

Proof. Value of Game 1 $\leq OPT_{primal}$. This is because every value attained by the primal LP can also be achieved in Game 1 by the min player taking the very same values of $x_1, \dots, x_k$. Thus by taking the optimal solution $(x_1^*, \dots, x_k^*)$ of the primal LP, the value of Game 1 becomes OPT_{primal}.

Suppose value of Game 1 $< OPT_{primal}$. Then some primal LP constraint got violated by the values $x_1^\circ, \dots, x_k^\circ$ chosen by the min player. However this makes the value of Game 1 $\to \infty$. But $OPT_{primal} < \infty$. This contradicts that value of Game 1 $< OPT_{primal}$. ▨

Game 2. Now the max player plays first. He chooses $y_1, \dots, y_n, \lambda_1, \dots, \lambda_k$: all these values have to be ≥ 0.

The min player plays second - he chooses $x_1, \dots, x_k$.

Value of the game $= f(x_1, \dots, x_k, y_1, \dots, y_n, \lambda_1, \dots, \lambda_k)$.

Claim 2. Value of Game 2 is OPT_{dual}.

Proof. Observe that $f(x_1, \dots, x_k, y_1, \dots, y_n, \lambda_1, \dots, \lambda_k)$ can also be written as $\sum_i b_i y_i + \sum_j x_j (c_j - \sum_i a_{ij} y_i - \lambda_j)$.

The min player plays after the max player and the min player is unconstrained – he can choose $x_1, \ldots, x_k$ to be non-negative or negative (recall that we said the max player is constrained to choose only non-negative values for his variables).

– How can the max player make the min player totally irrelevant?

The only way the max player can make the min player totally irrelevant is by taking

$$\sum_i a_{ij} y_i + \lambda_j = c_j \text{ for each } j.$$

That is, $\sum_i a_{ij} y_i \le c_j$ for $j = 1, \ldots, k$.

Since $y_1, \ldots, y_n$ anyway satisfy $y_i \geq 0 \ \forall i$, the max player's strategy is to choose $y_1, \ldots, y_n$ to be dual feasible.

* If the max player violates any of the dual LP constraints then the min player makes the value of game 2 $\longrightarrow -\infty$.

Thus the best strategy for the max player is to take $(y_1, \ldots, y_n)$ to be the optimal solution $(y_1^*, \ldots, y_n^*)$ of the dual LP and to take $\lambda_j = c_j - \sum_i a_{ij} y_i^*$ for $j = 1, \ldots, k$.

This makes the value of Game 2 $\sum_i b_i y_i^* = OPT_{dual}$. ▨

Let us now show a new proof of weak duality.

Claim 3. Value of Game 1 $\geq$ Value of Game 2.

Proof. Suppose the min player chooses values $x_1^\circ, \ldots, x_k^\circ$ in Game 1 and the max player chooses values $y_1^\circ, \ldots, y_n^\circ, \lambda_1^\circ, \ldots, \lambda_k^\circ$ in Game 1.

In Game 2, suppose the max-player chooses values $y_1', \ldots, y_n', \lambda_1', \ldots, \lambda_k'$ and the min-player chooses values $x_1', \ldots, x_k'$.

Then
$$\underbrace{\sum_j c_j x_j^\circ + \sum_i y_i^\circ \Big(b_i - \sum_j a_{ij} x_j^\circ\Big) - \sum_j \lambda_j^\circ x_j^\circ}_{\text{this is the value of Game 1}}$$

$$\geq \sum_j c_j x_j^\circ + \sum_i y_i' \Big(b_i - \sum_j a_{ij} x_j^\circ\Big) - \sum_j \lambda_j' x_j^\circ$$

since the max player could as well have chosen $(y_1', \ldots, y_n', \lambda_1', \ldots, \lambda_k')$ in Game 1 but he instead chose $(y_1^\circ, \ldots, y_n^\circ, \lambda_1^\circ, \ldots, \lambda_k^\circ)$. Recall that the max player desires to maximize the value of a game.

We can rewrite the above sum as

$$\sum_i b_i y_i' + \sum_j x_j^\circ \Big(c_j - \sum_i a_{ij} y_i' - \lambda_j'\Big)$$

$$\geq \sum_i b_i y_i' + \sum_j x_j' \Big(c_j - \sum_i a_{ij} y_i' - \lambda_j'\Big) \leftarrow \text{This is the value of Game 2.}$$

since in Game 2, the min player could as well have chosen $(x_1^\circ, \ldots, x_k^\circ)$ but he instead chose $(x_1', \ldots, x_k')$. Recall that the min player always desires to minimize the value of a game.

Hence value of Game 1 $\geq$ value of Game 2,

i.e, $OPT_{primal} \geq OPT_{dual}$. ▨

Observe that weak duality or what we proved in Claim 3 is very intuitive — the player who plays second has an advantage in any game since he can see the values chosen by the first player and choose his values accordingly. Since the max player plays second in Game 1 & the min player plays second in Game 2, we have value of Game 1 $\geq$ value of Game 2.

Proving Strong Duality

Our goal now is to show that there is an assignment of values to $y_1, \ldots, y_n, \lambda_1, \ldots, \lambda_k$ so that value of Game 2 $= \sum_j c_j x_j^*$ where $(x_1^*, \ldots, x_k^*)$ is the primal optimal solution.

Let X be the constraint matrix of the primal LP. That is, $X = \begin{bmatrix} A \\ I \end{bmatrix}$. Permute the rows of X into $\begin{bmatrix} X_= \\ X_> \end{bmatrix}$ so that the top k rows are satisfied as equalities by $(x_1^*, \ldots, x_k^*)$.

$X_=$ is made up some rows of A and the remaining are some unit vectors.

Assume without loss of generality that these rows of A are its first t rows and the remaining $k-t$ rows of $X_=$ are the unit vectors. So we have

$$\underbrace{\begin{bmatrix} X_= \end{bmatrix}}_{k \times k \text{ matrix}} \begin{bmatrix} x_1^* \\ x_2^* \\ \vdots \\ x_k^* \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_t \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Call the rows of $X_=$ $\underbrace{\vec{N_1}, \ldots, \vec{N_k}}$

these are the normals of the k hyperplanes whose intersection is $(x_1^*, \ldots, x_k^*)$.

The k vectors $\vec{N_1}, \ldots, \vec{N_k}$ are linearly independent.

So $(c_1, \ldots, c_k)$ can be expressed as a linear combination of $\vec{N_1}, \ldots, \vec{N_k}$. That is, $[c_1 \ldots c_k] = [\beta_1 \ldots \beta_k] \begin{bmatrix} \vec{N_1} \\ \vdots \\ \vec{N_k} \end{bmatrix}$

Let us take the following assignment of values to $y_1, \ldots, y_n, \lambda_1, \ldots, \lambda_k$:

- each y_i or λ_j corresponds to a constraint in the primal LP or equivalently, to a row in X.
  * take y_i or λ_j to be β_ℓ if the corresponding constraint is $\vec{N_\ell}$;
  * otherwise take it to be 0.

We need to show the following:

(i) all y_i's and λ_j's are ≥ 0;

(ii) the assignment of values given above makes the value of Game 2 $= \sum_j c_j x_j^*$.

Let us show (ii) first. We need to check that our assignment of values to λ_j's and y_i's satisfies

$$\sum_i a_{ij} y_i + \lambda_j = c_j \quad \text{for } j = 1, \ldots, k.$$

Consider the product of 3 matrices:

$$[\beta_1 \cdots \beta_k \underbrace{0 \ldots 0}_{n \text{ of them}}] \underbrace{\begin{pmatrix} \vec{N_1} \\ \vdots \\ \vec{N_k} \\ X_> \end{pmatrix}}_{\text{this is an } (n+k)\times k \text{ matrix}} \begin{bmatrix} x_1^* \\ \vdots \\ x_k^* \end{bmatrix}.$$

The product of the first two matrices is $[c_1 \ldots c_k]$. Hence the entire product is $\underline{\underline{\sum_j c_j x_j^*}}$.

We claim this product also equals $\sum_{i=1}^n b_i y_i$. This is by multiplying the second and third matrices first and then multiplying the first matrix with this product.

The product of the second and third matrices is

$$\begin{pmatrix} b_1 \\ \vdots \\ b_t \\ \hline 0 \\ \vdots \\ 0 \\ \hline * \\ \vdots \\ * \end{pmatrix} \begin{array}{l} \left.\vphantom{\begin{matrix} b_1 \\ \vdots \\ b_t \end{matrix}}\right\} t \text{ coordinates} \\ \left.\vphantom{\begin{matrix} 0 \\ \vdots \\ 0 \end{matrix}}\right\} k-t \text{ coordinates} \\ \left.\vphantom{\begin{matrix} * \\ \vdots \\ * \end{matrix}}\right\} n \text{ remaining coordinates} \end{array}$$

The product of the first matrix with this matrix $= \sum_{i=1}^{t} \beta_i b_i$

$= \sum_{i=1}^{n} y_i b_i$

(since the remaining $n-t$ y_i's are 0)

Hence assuming non-negativity of y_i's and λ_j's, we have shown that the value of Game 2 is $\underline{\sum_{j=1}^{k} c_j x_j^*}$.

Recall that $\vec{N_1}, \ldots, \vec{N_k}$ are the normals of hyperplanes whose intersection is $(x_1^*, \ldots, x_k^*)$. Since $\vec{N_1}, \ldots, \vec{N_k}$ are linearly independent, we said $(c_1 \ldots, c_k)$ can be expressed as a linear combination of $\vec{N_1}, \ldots, \vec{N_k}$.

Thus $[c_1 \ldots c_k] = [\beta_1 \ldots \beta_k] \begin{pmatrix} \vec{N_1} \\ \vdots \\ \vec{N_k} \end{pmatrix}$.

In fact, something stronger is true.

Claim. $[c_1 \ldots c_k]$ is a conic combination of $\vec{N_1}, \ldots, \vec{N_k}$.

That is, $[c_1 \ldots c_k]$ can be expressed as a non-negative linear combination of $\vec{N_1}, \ldots, \vec{N_k}$. Thus all y_i's and λ_j's are ≥ 0.

Why is the above claim true? We will prove this using a classical lemma from optimization. This is called Farkas' lemma.

Farkas' lemma. Exactly one of the following two statements is true: ④

(1) $(c_1, \dots, c_k)$ is a conic combination of $\vec{N_1}, \dots, \vec{N_k}$.

(2) $\exists$ a hyperplane $h_1 x_1 + \dots + h_k x_k = 0$ such that

$$[c_1 \dots c_k]\begin{bmatrix} h_1 \\ \vdots \\ h_k \end{bmatrix} < 0 \quad \text{and} \quad \begin{bmatrix} \vec{N_1} \\ \vdots \\ \vec{N_k} \end{bmatrix}\begin{bmatrix} h_1 \\ \vdots \\ h_k \end{bmatrix} \geq 0.$$

Farkas' lemma essentially says that if $(c_1, \dots, c_k)$ is not in the cone of $\vec{N_1}, \dots, \vec{N_k}$ then there is a separating hyperplane. Let us take Farkas' lemma for granted and complete our proof.

Assuming Farkas' lemma, if $(c_1, \dots, c_k)$ is not a conic combination of $\vec{N_1}, \dots, \vec{N_k}$ then there exists such a hyperplane $h_1 x_1 + \dots + h_k x_k = 0$.

Consider $(x_1^* + \varepsilon h_1, \dots, x_k^* + \varepsilon h_k)$ for a suitable $\varepsilon > 0$.

$$\sum_j c_j (x_j^* + \varepsilon h_j) = \sum_j c_j x_j^* + \varepsilon \cdot \sum_j c_j h_j$$

$$< \sum_j c_j x_j^* \quad \left(\text{since } [c_1 \dots c_k]\begin{bmatrix} h_1 \\ \vdots \\ h_k \end{bmatrix} < 0\right)$$

This will contradict the optimality of $(x_1^*, \dots, x_k^*)$ if we show $(x_1^* + \varepsilon h_1, \dots, x_k^* + \varepsilon h_k)$ to be primal feasible.

That is what we will show now. We will show there is a suitably small ε such that $(x_1^* + \varepsilon h_1, \dots, x_k^* + \varepsilon h_k)$ is primal feasible.

$$\begin{bmatrix} X_= \\ \\ X_\geq \end{bmatrix}\begin{pmatrix} x_1^*+\varepsilon h_1\\ \vdots \\ x_k^*+\varepsilon h_k\end{pmatrix} = \begin{pmatrix} b_1\\ \vdots\\ b_t\\ 0\\ \vdots\\ 0\\ \left.\begin{matrix} \\ \\ \end{matrix}\right\}\text{some slack}\end{pmatrix} + \begin{pmatrix} \left.\begin{matrix}\geq 0 \\ \vdots \\ \geq 0\end{matrix}\right\}\text{k of them} \\ * \\ \vdots \\ *\end{pmatrix}$$

The top k values in the rightmost column are the dot products of $\vec{N_1}, \ldots, \vec{N_k}$ with $(\varepsilon h_1, \ldots, \varepsilon h_k)$. We know all these k values are ≥ 0.

The $*$ values in the rightmost column are the dot products of $\vec{N}_{k+1}, \ldots, \vec{N}_{n+k}$ with $(\varepsilon h_1, \ldots, \varepsilon h_k)$ — these values can be negative. However $(x_1^*, \ldots, x_k^*)$ satisfies these constraints with slack.

So if we choose ε small enough, then the slack compensates for this negative value and so $(x_1^* + \varepsilon h_1, \ldots, x_k^* + \varepsilon h_k)$ is primal feasible.

Since the point $(x_1^* + \varepsilon h_1, \ldots, x_k^* + \varepsilon h_k)$ achieves a better objective function value than $(x_1^*, \ldots, x_k^*)$ — which contradicts the optimality of $(x_1^*, \ldots, x_k^*)$ — we can conclude that option (1) in Farkas' lemma holds. Thus $(c_1, \ldots, c_k)$ is a conic combination of $\vec{N_1}, \ldots, \vec{N_k}$. ∎

Let us use LP-duality to show a new proof of Max-flow min-cut theorem.

Max-flow as an LP

Our variables will be x_e for $e \in E$. Let us also add a new arc from t to s so that our flow becomes a circulation. That is, flow conservation will be obeyed at all vertices, including s and t.

Set $c_{(t,s)} = \infty$.

Primal LP

$$\max \; x_{(t,s)}$$

subject to

$$\sum_{e:\, e \text{ entering } u} x_e - \sum_{e:\, e \text{ leaving } u} x_e \leq 0 \quad \forall u \in V$$

$$x_e \leq c_e \quad \forall e \in E$$

$$x_e \geq 0 \quad \forall e \in E$$

Observe that setting $\sum_{e:\, e \text{ entering } u} x_e - \sum_{e:\, e \text{ leaving } u} x_e \leq 0$ for all u will imply $\sum_{e:\, e \text{ ent. } u} x_e = \sum_{e:\, e \text{ leav. } u} x_e$ for all u. (Why?)

Let us now write the dual LP. The variables are y_u for $u \in V$ and z_e for $e \in E$.

Dual LP

$$\min \sum_e c_e z_e$$

subject to

$$y_s - y_t \geq 1$$

$$y_v - y_u + z_{(u,v)} \geq 0 \quad \forall (u,v) \in E$$

$$y_u \geq 0, \quad z_e \geq 0 \quad \forall u \in V \; \forall e \in E$$

In order to understand the above LP better, let us add the following restriction to it:

let us constrain z_e, y_u for all edges e and vertices u to be either 0 or 1. So $y_s = 1$, $y_t = 0$.

For any edge (u,v) with $y_u = 1$ and $y_v = 0$, we are forced to set $z_{(u,v)} = 1$. Since we want to minimize $\sum_e c_e z_e$, we will set the remaining z_e values to 0.

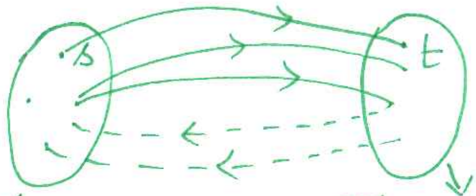

These are vertices whose y-value is 1.

These are vertices whose y-value is 0.

So the goal of the integer program is to come up with a partition of V into S and V-S such that $s \in S$, $t \in V-S$, and the sum of capacities of edges in $S \times (V-S)$ is minimized. Thus the integer program computes an s-t min-cut.

What about the dual LP? It computes a min fractional s-t cut. It assigns distance labels to arcs such that on any s-t path, the sum of distance labels is ≥ 1.

Let $s = u_0 \;—\; u_1 \;—\; u_2 \;\cdots\; u_{i-1} \overset{d_i}{—} u_i \;\cdots\; t$ be any s-t path.

$d_i \geq u_{i-1} - u_i$ for each i.

So $\sum_{i=1}^{k} d_i \geq u_0 - u_k \geq 1$

The constraint matrix of this LP has a special property — it is totally unimodular. This means any square submatrix of the constraint matrix has determinant in $\{0, \pm 1\}$.

This implies the feasible region is an integral polytope, i.e, the vertices or extreme points of this polytope have integral coordinates.

This means the optimal solution of the dual LP and the integer program are the same. That is, adding the integrality constraints does not change the optimal solution or the optimal value. Thus the dual LP computes an s-t min-cut. Since the primal LP computes a max-flow, LP duality implies max-flow = min-cut.

Total Unimodularity

Let us look at the constraint matrix of the max-flow LP. This matrix is as follows:

$$\begin{array}{c|c}
 & e_1 \quad \cdots \quad e_m \\
\begin{matrix} u_1 \\ \vdots \\ \\ \\ u_n \end{matrix} & \left[\begin{matrix} 1 & 0 & \\ 0 & -1 & \\ 0 & 0 & \cdots \\ -1 & 0 & \\ 0 & 1 & \end{matrix}\right] \\
\hline
\begin{matrix} e_1 \\ \vdots \\ e_m \end{matrix} & I
\end{array}$$

$\rightarrow$ Every column in the upper part of this matrix has one $+1$ and one -1 in it, all other entries are 0.

$\rightarrow$ Identity matrix

Take any $s \times s$ submatrix (where $1 \leq s \leq m$) of the above matrix. To compute its determinant, keep expanding along any row/column that has a single 1 or a single -1 till there is an all 0s row/column or every row/column has ≥ 2 non-zero entries. Otherwise we have a 1×1 matrix and the determinant of our submatrix $= (\pm 1) \cdot (\text{value of this entry}) \in \{0, \pm 1\}$.

Claim. The determinant of such a matrix (as underlined above) is 0. (Please do this as an exercise.)

Another application of strong duality

We will prove a result in graph theory now. Before we state this result, let us define a vertex cover: this is a subset of the vertex set such that every edge has at least one endpoint in this set.

König - Egerváry theorem. In any bipartite graph, the size of a maximum matching = the size of a minimum vertex cover.

It is important to restrict the graph to be bipartite. Consider

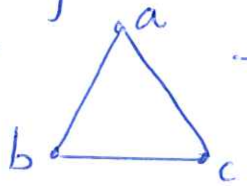

→ here the size of a maximum matching is 1 and the size of a minimum vertex cover is 2.

In any graph, it is easy to see that the size of a maximum matching ≤ the size of a minimum vertex cover.

What König-Egerváry theorem says is that the above constraint is tight for bipartite graphs. We will prove this using LP-duality.

Consider the following linear program in a bipartite graph $G = (A \cup B, E)$.

$$\max \sum_{e \in E} x_e$$

subject to

$$\sum_{e \in \delta(u)} x_e \leq 1 \quad \forall\, u \in A \cup B$$

$$x_e \geq 0 \quad \forall\, e \in E$$

Here $\delta(u)$ is the set of edges incident to vertex u.

If we add the constraints $x_e \in \mathbb{Z}$ then this integer program computes a maximum matching in G.
The claim is that the above constraint matrix is totally unimodular. So adding integrality constraints is redundant since the optimal solution is integral.

Let $v_1, \ldots, v_n$ be the vertices of G and let $e_1, \ldots, e_m$ be its edges. The constraint matrix looks as follows:

$$M = \begin{array}{c} \\ v_1 \\ \vdots \\ v_n \end{array} \begin{array}{c} e_1 \; \ldots \; e_m \\ \left[\begin{array}{ccc} & & \\ & & \\ & & \end{array} \right] \end{array}$$

$\rightarrow$ This is a 0/1 matrix with $M[v_i, e_j] = 1$ if v_i is one of the endpoints of e_j.
And $M[v_i, e_j] = 0$ otherwise

Consider any $k \times k$ submatrix.

- If there is an all 0s row or all 0s column in this submatrix then its determinant is 0.
- If there is a row or column with only one 1 then expand using that row/column.
  - Continue doing the above step till we are either left with a 1×1 matrix (so determinant $= \pm$ value of this element) or there is an all 0s row/column (in this case determinant is 0) or we have a matrix where every row has ≥ 2 1s and every column has exactly 2 1s.

Claim. In the last case, the determinant of this matrix is 0.

Proof. Observe that every row also has exactly two 1s. Since every column has exactly two 1s, the total number of 1s in this submatrix is $2t$ (where t is the number of columns in this submatrix) and since there are t rows and at least two 1s in each row, it follows that every row also has exactly 2 1s. So this matrix is a disjoint collection of cycles. Since the graph is bipartite, each cycle is of even length.

We will come up with a linear combination of the rows of this $t \times t$ matrix that is the all 0's vector. This means the rows of this matrix are linearly dependent — hence the determinant of this matrix is 0.

$$\begin{array}{c|c} & e_{j_1} \;\cdots\; e_{j_t} \\ \begin{array}{c} v_{i_1} \\ \vdots \\ v_{i_t} \end{array} & \left[\begin{array}{ccc} & & \\ & & \\ & & \end{array}\right] \end{array} \equiv \text{A disjoint collection of even length cycles.}$$

Multiply rows indexed by vertices in A with -1 and rows indexed by vertices in B with $+1$ and add all the rows. Observe that we get the all 0s vector. ▨

Let us now prove König - Egerváry theorem. Consider the LP dual to the bipartite matching LP. The variables here are y_u $\forall$ vertices u.

$$\min \sum_{u \in A \cup B} y_u$$

subject to

$$y_u + y_v \geq 1 \quad \forall (u,v) \in E$$

$$y_u \geq 0 \quad \forall u \in A \cup B.$$

This is the minimum fractional vertex cover LP.
However since the constraint matrix (this is the transpose of the constraint matrix of the primal LP) is totally unimodular, this is the minimum integral vertex cover LP – note that this is so only for bipartite graphs.

So we have: size of maximum matching $=$ size of maximum fractional matching
(due to total unimodularity)

and size of maximum fractional matching $=$ size of minimum fractional vertex cover
(due to LP duality)

and size of minimum fractional vertex cover $=$ size of minimum integral vertex cover
(due to total unimodularity)

Exercise. Show that every doubly stochastic matrix is a convex combination of permutation matrices.

- A doubly stochastic matrix is a square matrix with non-negative entries such that every row-sum is 1 and every column sum is 1.
- A permutation matrix is a square matrix with 0/1 entries such that every row has exactly one 1 in it and similarly, every column has exactly one 1 in it.
- A convex combination is a linear combination where all coefficients are non-negative and sum to 1.

The above statement is called Birkhoff-von Neumann theorem

Complementary Slackness

We have our primal LP on k variables $x_1, \dots, x_k$ and dual LP on n variables $y_1, \dots, y_n$.

$(x_1^*, \dots, x_k^*)$ is a primal optimal solution and $(y_1^*, \dots, y_n^*)$ is a dual optimal solution if and only if:

- for every $1 \le j \le k$: either $x_j^* = 0$ or $\sum_{i=1}^{n} a_{ij} y_i^* = c_i$
- for every $1 \le i \le n$: either $y_i^* = 0$ or $\sum_{j=1}^{k} a_{ij} x_j^* = b_i$

These conditions are called complementary slackness conditions. Why do they hold?

We have
$$\sum_i \sum_j a_{ij} x_j^* y_i^* = \sum_i y_i^* \Big(\sum_j a_{ij} x_j^*\Big) \geq \sum_i y_i^* b_i$$

We also have
$$\sum_j \sum_i a_{ij} x_j^* y_i^* = \sum_j x_j^* \Big(\sum_i a_{ij} y_i^*\Big) \leq \sum_j x_j^* c_j$$

Since $(x_1^*, \dots, x_k^*)$ is a primal optimal solution and $(y_1^*, \dots, y_n^*)$ is a dual optimal solution, $\sum_j c_j x_j^* = \sum_i b_i y_i^*$. Thus complementary slackness conditions have to hold.

Finding a min-cost perfect matching in bipartite graphs

Our input is a complete bipartite graph $G = (A \cup B, E)$ where every edge has a cost associated with it.

The problem is to find a perfect matching M in G whose sum of edge costs is minimum.

We can solve the above problem by solving the following LP:

$$\min \sum_{i,j} c_{ij} x_{ij}$$

subject to

$$\sum_j x_{ij} = 1 \quad \forall i \in A$$

$$\sum_i x_{ij} = 1 \quad \forall j \in B$$

$$x_{ij} \geq 0 \quad \forall i \in A, j \in B.$$

Since the above constraint matrix is totally unimodular (why?), there is an optimal solution to the above LP that is integral and we can assume that the LP-solver returns this solution. Any integral feasible point is the incidence vector of a perfect matching. Thus we have found a min-cost perfect matching.

Though we can solve an LP in polynomial time, the running time is quite high. We prefer combinatorial algorithms. We will now see a combinatorial algorithm based on linear programming. This algorithm is "primal-dual". Let us write down the dual LP.

$$\max \sum_i u_i + \sum_j v_j$$

subject to

$$u_i + v_j \leq c_{ij} \quad \forall i \in A, j \in B.$$

Note: We have no non-negativity constraints on u_i, v_j since the primal LP constraints are equalities.

Our goal is to find a dual feasible solution $(u_i)_{i \in A}$, $(v_j)_{j \in B}$ and a perfect matching M such that

$$\sum_{(i,j) \in M} c_{ij} = \sum_{i \in A} u_i + \sum_{j \in B} v_j.$$

Then weak LP duality implies that M is a primal optimal solution.

For the above equality to hold, the matching M has to contain only those edges (i,j) for which $c_{ij} = u_i + v_j$ (by complementary slackness).

The algorithm works as follows:

- start with any dual feasible solution, say $u_i = 0 \ \forall i \in A$ and $v_j = \min_i c_{ij} \ \forall j \in B$.

In a given iteration, the algorithm has a dual feasible solution $\vec{u}, \vec{v}$. Let G_0 be the subgraph whose edge set consists of edges (i,j) such that $c_{ij} = u_i + v_j$. We check if G_0 admits a perfect matching or not.

- Here we use the max-size matching algorithm in bipartite graphs seen earlier in the course.

If G_0 has a perfect matching then the edge incidence vector of this matching is a primal feasible solution and satisfies complementary slackness with the current dual solution. Hence this is a min-cost perfect matching in G.

Suppose G_0 does not have a perfect matching. Then the algorithm will update the dual solution such that the objective function value of the dual solution increases (recall that we are maximizing the dual).

Let $\delta = \min_{i \in A', j \in B'} (c_{ij} - u_i - v_j)$

where A' = the set of vertices in A reachable by an alternating path from an unmatched vertex in A.

a path whose edges alternate between being in M and not in M.

and B' = the set of vertices in B not reachable by an alternating path from an unmatched vertex in A.

By definition of A' and B', note that

$c_{ij} > u_i + v_j \quad \forall\, i \in A'$ and $j \in B'$. Thus $\underline{\delta > 0}$.

Let us update the dual solution as follows:

$$u_i = \begin{cases} u_i & \text{for } i \in A-A' \\ u_i + \delta & \text{for } i \in A' \end{cases} \quad \text{and} \quad v_j = \begin{cases} v_j & \text{for } j \in B' \\ v_j - \delta & \text{for } j \in B-B'. \end{cases}$$

Please check the following:

- the above values of $\vec{u}$ and $\vec{v}$ are dual feasible.
- the difference in the objective function values achieved by the new dual solution and the old one is $\delta\left(\frac{n}{2} - k\right)$, where k is the size of a maximum matching in G_0.

Since $k < n/2$, the value strictly increases.

We repeat the above procedure until the algorithm terminates. At that point, we have an incidence vector of a perfect matching and a dual feasible solution that satisfy complementary slackness. Thus we find a min-cost perfect matching.

We still have to show that the algorithm terminates. Observe that whenever we update the dual solution, at least one more vertex in B moves from B' to $B-B'$ and no vertex in $B-B'$ moves out. Thus the algorithm runs for at most $|B| = n/2$ iterations till the size of the matching increases. So in $O(n^2)$ iterations, we have a perfect matching.