Transshipment Problem

CO 351: Network Flow Theory

David Duan, 20703592

Contents

1 Simplex Algorithm

1.1 Preliminaries

1.1.1 Standard Equality Form

An LP is in **SEF** when it is (1) a *maximization* problem with (2) *non-negative variables* and (3) *equality constraints*. The Simplex algorithm can solve any LP as long as it is in SEF.

To turn an LP into its SEF equivalent, (1) replace $\min f$ with $\max -f$; (2) use slack variables to replace inequalities by equalities; (3) express any free variable as the difference between two non-negative integers.

1.1.2 Basis and Basic Solutions

Consider an LP (P) $:= \{\max c^T x : Ax = b, x \geq \mathbf{0}\}$ where $A \in \mathbb{R}^{m\times n}$ with $\mathrm{rank}(A) = m$.

We write $A_J := [A_j]_{j\in J}$ for $J \subseteq [n]$. A set $B \subseteq [n]$ is a **basis** of A if A_B is *square* and *non-singular*. Equivalently, $|B| = m$ and $\mathrm{rank}(A_B) = m$. For $\ell \in B$, x_ℓ is called a **basic variable**.

Let $B \subseteq [n]$ be a basis for A and $N := [n] \setminus B$. Then $Ax = b$, $x_N = 0$ has a unique solution given by $x_B = A_B^{-1}b$, $x_N = 0$. We say that this is the **basic solution** corresponding to B. Equivalently, a basic solution x satisfies $Ax = b$ and $x_N = 0$. Note each basis has exactly one corresponding basic solution, but one basic solution may have many associated basis.

A basic solution x satisfying $x \geq \mathbf{0}$ is called a **basic feasible solution** (**BFS**).

1.1.3 Canonical Form

Given a basis B of A, an LP is in **canonical form** if $A_B = I$ and $c_j = 0$ for $j \in B$. This allows us to easily check whether the current BFS $\bar{x}$ is optimal for the given LP. If not, we will try a different basis. For (P) $:= \{\max c^T x : Ax = b, x \geq \mathbf{0}\}$, its equivalent in canonical form is $\max\{c_N^T x_N + \bar{z} : x_B + A_N x_N = b, x \geq \mathbf{0}\}$.

To turn an LP into canonical form, we need to rewrite the constraint and the objective function. The new constraint is given by $A_B^{-1}Ax = A_B^{-1}b$. The new objective function is given by $z = [c^T - y^T A]x + y^T b$ where $y = A_B^{-T}c_B$.

1.1.4 The Fundamental Theorem of Linear Programming

Let (P) denote an LP problem. Then either (P) is infeasible, or it is unbounded, or it is feasible with an optimal solution.

Moreover, suppose that (P) is in SEF and its constraint matrix has full row rank. Then if (P) has a feasible solution, then it has a feasible solution that is basic; if (P) has an optimal solution, then it has an optimal solution that is basic.

1.2 Two Phase Method

1.2.1 Phase I. Initialization

Given a general LP (P) $:= \{\max c^T x : Ax = b, x \geq \mathbf{0}\}$, we want to find an initial BFS to start Simplex. First, multiply equations so that b becomes non-negative. Next, construct an auxiliary LP $\min\{x_{n+1} + \cdots + x_{n+m} : (A \mid I)\, x = b, x \geq \mathbf{0}\}$ where $x_{n+1}, \ldots, x_{n+m}$ are auxiliary variables, and solve the auxiliary LP using Simplex. If $(x_1, \ldots, x_n, x_{n+1}, \ldots, x_{n+m})^T$ is the optimal solution to the auxiliary problem, then LP is feasible if and only if $x_{n+1} = \cdots = x_{n+m} = 0$.

An LP in SEF is either infeasible, unbounded, or feasible with a basic optimal solution.

1.2.2 Phase II. Loop

Given an LP with a feasible basis B, we want to output an optimal solution or report that the LP is unbounded. First, we rewrite the LP in canonical form for the basis B:

$$\begin{array}{lll} \max & c^T x \\ s.t. & Ax = b \\ & x \geq \mathbf{0} \end{array} \quad \Longrightarrow \quad \begin{array}{ll} \max & z = c_N^T x_N + \bar{z} \\ s.t. & x_B + A_N x_N = b \\ & x \geq \mathbf{0} \end{array}$$

At this stage, B is a feasible basis, $N = \{j \notin B\}$; the LP is in canonical form for B, and $\bar{x}$ is a basic solution (Remark: $A_B = I$). Next, we find a better basis B' or get required outcome:

1. If $c_N \leq \mathbf{0}$, STOP. The basic solution $\bar{x}$ is optimal. (Proof A)
2. Pick $k \notin B$ such that $c_k > 0$ and set $x_k = t$.
3. Pick $x_B = b - tA_k$.
4. If $A_k \leq 0$, STOP. The LP is unbounded. (Proof B)
5. Choose $t = \min\{b_i / A_{ik} \mid \forall i : A_{ik} \geq 0\}$.
6. Let x_r be a basis variable forced to 0.
7. Obtain the new basis by having k enter and r leave.

Bland's rule states that, if we have a choice for element entering or leaving the basis, always pick the smallest one. If we use Bland's rule, then the Simplex algorithm always terminates.

1.3 Tableau Method

See example.

2 Basic Graph Theory

A **directed graph** (**digraph**) $D = (N, A)$ consists of a set of **nodes** N and a set of **arcs** A.

For an arc $uv \in A$, u is the **tail** and v is the **head**.

The **out degree** of a node $v \in N$, denoted $d(v)$, is the number arcs whose tail is v.

The **in degree** of a node $v \in N$, denoted $d(\bar{v})$ is the number of arcs whose head is v.

A **directed walk** (**diwalk**) is a sequence of nodes $v_1, \ldots, v_k$ where v_i, v_{i+1} is an arc.

A **directed path** (**dipath**) is a diwalk with no repeated node.

A **directed cycle** (**dicycle**) is a diwalk $v_1, \ldots, v_{k-1}, v_k, v_1$, where $v_1, \ldots, v_k$ are distinct.

The **cut induced by** $S \subseteq N$ is the set of arcs leaving S, i.e., $\delta(S) := \{uv \in A : u \in S, v \notin S\}$.

We also define the set of arcs going into S: $\delta(\bar{S}) := \{uv \in A : u \notin S, v \in S\}$.

If $s \in S$ and $t \notin S$, we call $\delta(S)$ an s, t**-cut**.

The **underlying graph** of $D = (N, A)$ is obtained by removing the directions on each arc.

A **cycle** in a digraph is the node and arcs corresponding to a cycle in the underlying graph.

If we fix an orientation on the cycle, then arcs in this direction are called **forward arcs**, and other arcs are called **backward arcs**.

An undirected graph $G = (V, E)$ is **connected** if there exists a path between every pair of nodes $i, j \in V(G)$. A digraph is **connected** if the underlying graph is connected.

A set of arcs $T \subseteq E(G)$ is said to be a **spanning tree** of an undirected graph G if T is connected and $|T| = |E(G)| - 1$. A **spanning tree** of a digraph is a spanning tree of the underlying graph.

Let $T = (V, E)$ be a digraph. The following are equivalent:

1. T is a tree, i.e., T is connected, and T has no cycles.
2. There exists exactly one path between every pair of nodes in T.
3. T is connected and $|E| = |V| - 1$.
4. T has no cycles and $|E| = |V| - 1$.

Let $T = (V, E)$ be a (directed) tree. Then

1. Adding a new arc $uv \notin E$ results in a digraph $T + uv$ that has exactly one cycle C.
2. Removing any one arc pq from the unique cycle C of $T + uv$ results in a tree $T + uv - pq$.

3 Transshipment Problem

3.1 Overview

3.1.1 Problem

Given a problem

- Digraph $D = (N, A)$ - N is a set of nodes and A is a set of arcs.
- Node demands $b \in \mathbb{R}^N$ - negative = supply node, positive = demand node.
- Arc costs $w \in \mathbb{R}^A$ - can be positive, negative, or zero.

We want to find a flow of minimum cost:

- Flow $x \in \mathbb{R}^A$ - payload to each arc, satisfying the node demands and $x \geq \mathbf{0}$.
- Cost $w^T x = \sum_{ij \in A} w_{ij} x_{ij} \in \mathbb{R}$ - total costs by using arcs to transport units.

Remark. We assume that D is connected (i.e., it contains a spanning tree) and the node demands sum to zero (i.e., $b(N) = \sum_{v \in N} b_v = 0$). These assumptions are not essential and we may remove them if needed. See later sections for more explanation.

3.1.2 LP Formulation

The primal LP for TP is as follows.

$$\begin{array}{ll} \min & w^T x \\ s.t. & x(\delta(\bar{v})) - x(\delta(v)) = b_v \quad (\forall v \in N) \\ & x \geq \mathbf{0} \end{array}$$

For each node in the flow, the in-flow minus the out-flow equals its demand, i.e.,

$$x(\delta(\bar{v})) - x(\delta(v)) = \sum_{iv \in A} x_{iv} - \sum_{vj \in A} x_{vj} = b_v \quad (\forall v \in N).$$

- $x(\delta(\bar{v})) = \sum_{iv \in A} x_{iv}$: In-flow for a node $v \in N$.
- $x(\delta(v)) = \sum_{vj \in A} x_{vj}$: Out-flow for a node $v \in N$.

3.1.3 Incidence Matrix

Alternatively, we could express the constraints using an incidence matrix, where each row corresponds to a node and each column corresponds to an arc.

$$
\begin{aligned}
\min \quad & w^T x \\
s.t. \quad & Mx = b_v \quad (\forall v \in N) \\
& x \geq \mathbf{0} \\
\text{where} \quad & \\
& M \in \{-1, 0, 1\}^{|N| \times |A|} \\
& M_{v,ij} = \begin{cases} -1 & i = v, \text{ i.e., the arc leaves } v \\ 1 & j = v, \text{ i.e., the arc goes into } v \\ 0 & \text{else} \end{cases}
\end{aligned}
$$

3.1.4 Dual LP

Recall to derive (P_{max}) given (P_{min}), we flip signs of variables for signs of constraints and keep signs of constraints the same for the signs of variables. Thus, $x \geq \mathbf{0} \implies M^T y \leq w$ and $Mx = b_v \implies y$ free.

$$
\begin{aligned}
\max \quad & b^T y \\
s.t. \quad & M^T y \leq w_{uv} \quad (\forall uv \in A) \\
& y \text{ free}
\end{aligned}
$$

Since each column of M contains exactly one entry of 1 (the row corresponding to the arc's head or destination) and exactly one entry of -1 (the row corresponding to the arc's source or tail), each row in M^T contains exactly these two non-zero entries. Thus, we can also write dual constraints explicitly for each arc $uv \in A$:

$$
\begin{aligned}
\max \quad & b^T y \\
s.t. \quad & y_v - y_u \leq w_{uv} \quad (\forall uv \in A) \\
& y \text{ free}
\end{aligned}
$$

3.1.5 Complementary Slackness Conditions

Let $\bar{x}$ and $\bar{y}$ be feasible solutions for a pair of primal and dual LPs (P_{max}) and (P_{min}). Recall the solutions are optimal if the CS conditions hold:

- For all variables x_j of (P_{max}), $\bar{x}_j = 0$ or the jth constraint of (P_{min}) is satisfied with equality for $\bar{y}$;
- For all variables y_i of (P_{min}), $\bar{y}_i = 0$ or the ith constraint of (P_{max}) is satisfied with equality for $\bar{x}$.

Since all primal constraints are equalities, if a flow x is optimal, then the CS conditions say that there must exist a (dual potential) y such that $x_{uv} = 0$ or $y_v - y_u = w_{uv}$ for each $uv \in A$.

If the (dual potential) is then feasible (for all other non-active conditions, the constraints still hold), then we have verified optimality of the solution for the primal.

Intuitively, for an arc $uv \in A$, we either don't use it at all ($x_{uv} = 0$), or we use it to its full potential ($y_v - y_u = w_{uv}$). See economical interpretation in section 3.3.1 for more explanation.

3.2 Cycle, Spanning Tree, and Basis

Recall that an LP is called *infeasible* if it has no feasible solutions, and *unbounded* if it has feasible solutions with objective values that are arbitrarily large (for maximization) or arbitrarily small (for minimization). Review section 1.2 for more information.

By definition, each column of the incidence matrix M contains exactly one entry of 1 and one entry of -1. Thus, the sum of all rows of the incidence matrix is the zero vector $\mathbf{0} \in \mathbb{R}^{|A|}$. Recall from linear algebra, the vectors in a set $T = \{v_1, \ldots, v_n\}$ are linearly independent if the equation $\sum c_i v_i = 0$ can only be satisfied by $c_i = 0$ for all i. This means $\text{rank}(M) \leq |N| - 1$. We will now show that, assume D is connected, we have $\text{rank}(M) = |N| - 1$.

Proposition The columns of M corresponding to a cycle are linearly dependent.

Proof. Let C be a cycle and M be its incidence matrix. Let F, B be the forward and backward arcs of C, respectively. Now, suppose M' is obtained by multiplying columns of (backward) arcs in B by -1. Then M' is the incidence matrix of a cycle where every arc is a forward, or equivalently, the new graph C' is a directed cycle.

Since $\deg(v) = \deg(\bar{v}) = 1$ for all $v \in V$, in each row of M', there exists exactly one entry of 1 and one entry of -1. Thus, the sum of the entries of each row is 0. This implies that the sum of columns of M' is zero (you can use an example to verify). Since this is a non-trivial linear combination of columns of M which equals to zero, M is linearly dependent. □

Proposition Let D be a digraph and M be its incidence matrix. Let B be a subset of the arcs whose corresponding columns in M are linearly dependent. Prove that B contains a cycle. (Note that this cycle is not necessarily a directed cycle.)

Lemma. If every vertex of a graph has degree at least 2, then it contains a cycle. (See Appendix for Proof.)

Proof. Since the columns of B in M are linearly dependent, there exists non-trivial linear combination of these columns that equal to 0. Let $B' = \{e_1, \ldots, e_k\}$ be the subset of B whose columns receive non-zero coefficients in the linear combination, i.e., there exist non-zero c_i's where $c_1 M_{e_1} + \cdots + c_k M_{e_k} = 0$. For each node v, $c_1 M_{v,e_1} + \cdots + c_k M_{v,e_k} = 0$, where M_{v,e_i} represents the entry for v in the column for e_i.

Let S be the subset of nodes v where $M_{v,e_i} \neq 0$ for at least one i. For a node v in S, if $M_{v,e_i} \neq 0$, then $c_i M_{v,e_i} \neq 0$. Hence for each $v \in S$, there exist at least 2 arcs in B' whose columns in M have non-zero entries for v. Each of these arcs is incident with v. Therefore, the subgraph with S as the nodes and B' as the arcs have the property that every node has degree ≥ 2. By the Lemma, it contains a cycle. □

Recall that a basis is the maximal linearly independent set of vectors. In our case, a basis corresponds to a spanning tree of a graph, as it is maximal with no cycles. From these two propositions, we derive the following theorem:

Theorem Let M be the incidence matrix of a digraph $D = (N, A)$. Then a set of $|N| - 1$ columns of M is a basis if and only if the corresponding $|N| - 1$ arcs make up a spanning tree.

Intuitively, since you can never get a non-trivial linear combination of $\tilde{M}$ obtained by removing an arbitrary row from M equal to zero, $\text{rank}(M) = \text{rank}(\tilde{M}) = |N| - 1$.

We now link TP back to LP.

Given a spanning tree T of D, the unique solution x of $Mx = b$ such that $x_{ij} = 0$ for all arcs ij not in T is called a *tree solution.* (Note that arcs ij in T may have $x_{ij} > 0, x_{ij} = 0$, or $x_{ij} < 0$.) Thus, a tree solution is a basic solution of (P). A *flow* $x \in \mathbb{R}_+^A$ such that $x_{ij} = 0$ for all arcs ij not in T is called a *tree flow.* In other words, a tree solution x that is non-negative is called a tree flow. Thus, a tree flow is a basic feasible solution of (P).

3.3 Network Simplex

In the Simplex method, at all times we maintain a basic feasible solution for the LP, and a solution for the dual that satisfies all constraints (except the non-negative constraints), and the complementary slackness conditions. Each step fixes some of the non-negativity constraints, so eventually both dual and primal solutions are satisfied, and by CS conditions, both solutions must be optimal.

3.3.1 Node Potential and Reduced Cost

We have the following definitions:

1. The dual solution $y \in \mathbb{R}^N$ is called a *node potential.*
2. Given a node potential y, the *reduced cost* of arc uv is $\bar{w}_{uv} = w_{uv} + y_u - y_v$.
3. A node potential is *feasible* if $\bar{w}_{uv} \geq 0$ for all $uv \in A$.

We now provide some intuition -- an economic interpretation -- for these definitions. Suppose we are transporting a lot of commodities, where the dual potentials represent the prices for buying/selling of these commodities at the nodes.

Let $y_u = 80, y_v = 120, w_{uv} = 70$. The reduced cost on arc uv is $w_{uv} + y_u - y_v = 70 + 80 - 120 = 30$. We can view this as buying a unit at u costs 80, transporting it to v costs 70, and selling it at v earns 120, thus in total we suffer a loss of 30.

Let $y_u = 30, y_v = 80, w_{uv} = 40$. The reduced cost on arc uv is $w_{uv} + y_u - y_v = 40 + 30 - 80 = 10$. We can view this as buying a unit at u costs 40, transporting it to v costs 30, and selling it at v earns 80, thus in total we make a profit of 10.

From these two examples, we see that we want to transport through $e \in A$ if and only if $\bar{w}_e < 0$. In Simplex, when we find $\bar{w}_e < 0$, we want to increase flow on this arc, i.e., pick e to enter our basis, as it saves us money.

3.3.2 Unboundedness

In Network Simplex, unboundedness occurs when no leaving arc can be found, so that you can increase parameters (entering variable) indefinitely.

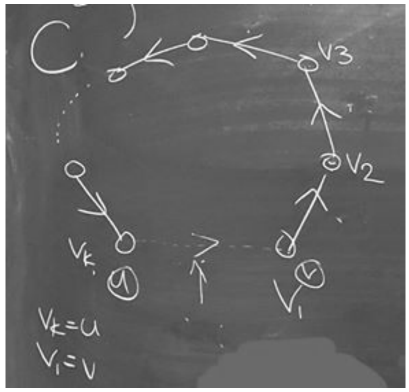

Consider the diagram above. Suppose all arcs $v_i v_{i+1}$ are in the basis and we pick uv as the entering arc. Let y_{v_1} be the dual potential at v_1. Since all arcs $v_i v_{i+1}$ here are in the basis, their reduced cost must be zero. Then

$$\begin{aligned} y_{v_2} &= y_{v_1} + w_{v_1 v_2} \\ y_{v_3} &= y_{v_2} + w_{v_2 v_3} = y_{v_1} + w_{v_1 v_2} + w_{v_2 v_3} \\ &\cdots \\ y_{v_k} &= y_{v_1} + \cdots + w_{v_{k-1} v_k} \end{aligned}$$

Since $\bar{w}_{uv} = y_{v_k} + w_{uv} = y_{v_1} = (y_{v_1} + \cdots + w_{v_{k-1}v_k}) + w_{uv} - y_{v_1}$,

$$\begin{aligned} \bar{w}_{uv} &= y_{v_k} + w_{uv} = y_{v_1} \\ &= (y_{v_1} + \cdots + w_{v_{k-1}v_k}) + w_{uv} - y_{v_1} \\ &= w_{v_1 v_2} + w_{v_2 v_3} + \cdots + w_{v_{k-1}v_k} + w_{uv} =: w(C) \end{aligned}$$

Since uv was chosen as an entering arc, $\bar{w}_{uv} < 0$, and hence the sum of arc costs on cycle C is negative. We call such C a *negative dicycle.*

Theorem A feasible TP is unbounded if and only if there exists a negative dicycle.

($\Longleftarrow$) Let C be a negative dicycle, i.e., $w(C) < 0$. Let x^* be a feasible flow. Define flow x^C where

$$x_e^C = \begin{cases} t & e \in C \\ 0 & e \notin C \end{cases}$$

for $t \geq 0$. So $x^C(\delta(\bar{v})) = x^C(\delta(v)) = 0$ for each $v \in N$. (Easy to check.)

Since C is a dicycle, $x^* + x^C$ is a feasible flow:

$$(x^* + x^C)(\delta(\bar{v})) - ((x^* + x^C)(\delta(v)) = [x^*(\delta(\bar{v})) - x^*(\delta(v))] + [x^C(\delta(\bar{v})) - x^C(\delta(v))] = b_v + 0 = b_v.$$

The objective value is $w^T(x^* + x^C) = w^T x^* + w^T x^C = w^T x^* + t \cdot w(C)$, where $w^T x^*$ is a fixed finite number. Then we have $w^T(x^* + x^C) \to -\infty$ as $t \to \infty$ as $w(C) < 0$.

For example, suppose b, c, d, f forms a negative dicycle C. Then we can push as much of goods on this dicycle since it will cancel out eventually. $t = 100$ in this case.

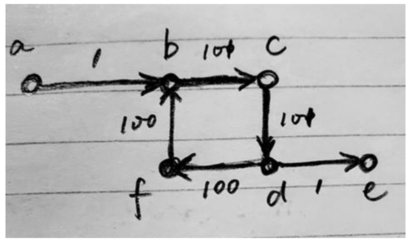

($\Longrightarrow$) Obtain a new digraph D' by adding a new node z and arcs zv for all $v \in N$. Set the cost of new arcs to be 0. For each node in D', let y_v be the minimum cost among all possible directed z, v-paths. This minimum exists since there is at least one such dipath, and the number of such dipaths is finite.

First, we show if D does not contain a negative dicycle, then y is a feasible potential for D'. Suppose not, i.e., there exist $pq \in A'$ where $\bar{w}_{pq} = w_{pq} + y_p - y_q < 0$. Consider a z, p-dipath P that has minimum cost. By assumption, $w(P) = y_p$. If q is not on P, then $P + pq$ is a z, q-dipath with cost $y_p + w_{pq}$, which is strictly less than y_q by assumption. This contradicts the fact that y_q is the minimum cost of all z, q-dipaths.

Now assume q is on P, and let P_1 be the part of P from z to q, and P_2 be the part of P from q to p. Since P_1 is a z, q-dipath and a minimum cost z, q-dipath has cost y_q, we have $w(P_1) \geq y_q$. Also, $w(P) = w(P_1) + w(P_2) = y_p$. Then $w(P_2) = y_p - w(P_1) \leq y_p - y_q$. By assumption, $y_p - y_q < -w_{pq}$, so $w(P_2) < -w_{pq}$, or $w(P_2) + w_{pq} < 0$. So $P_2 + pq$ is a dicycle whose cost is negative, contradicting the assumption that there is no negative dicycle. Hence y is a feasible potential for D'.

If D does not contain a negative dicycle, then there exists a feasible potential for D'. The same potential applied to only nodes in D is also feasible, hence the dual LP is feasible. Therefore, the original TP is bounded by the objective value of the dual feasible solution. $\square$

3.4 Network Simplex Algorithm

3.4.1 Detailed Algorithm

Given a connected digraph $D = (N, A)$ with arc costs w and node demands b,

1. Find a spanning tree T of D such that the associated tree solution x is non-negative. If no such tree can be found, STOP. The given transshipment problem has no feasible solution.
   a. To find a tree solution x for a given T, either use inspection, or solve the system $M_T x_T = b$.
2. Using T and the arc costs w, find $y \in \mathbb{R}^N$ such that each arc $ij \in T$ has $\bar{w}_{ij} := w_{ij} + y_i - y_j = 0$. y need not be feasible for the dual LP.
   a. To find y, use inspection, or solve the system $M_T' y = w_T$, where M_T' denotes the transpose of M_T.

3. While there exists arc $ij \in A$ such that $\bar{w}_{uv} < 0$, do
   a. Find such an arc uv so $\bar{w}_{uv} < 0$.
   b. Let Q be the oriented cycle obtained from the unique cycle in $T + uv$ by choosing the orientation such that uv is a forward arc.
   c. If all the arcs in Q are forwarded, STOP. The problem is UNBOUNDED as we have found a negative dicycle.
   d. Let pq be a reverse arc of Q such that $x_{pq} = \min\{x_{ij} \mid ij \text{ is a reverse arc of } Q\}$ and let $\gamma = x_{pq}$.
   e. Push γ units of flow along Q, i.e., for each forward arc ij in Q, increase x_{ij} by γ, and for each reverse arc ij in Q, decrease x_{ij} by γ.
   f. Replace T with $T + uv - pq$.
   g. Recompute $y \in \mathbb{R}^N$ using the new spanning tree.
4. Go back to 2.

If the algorithm terminates after the while loop (with $\bar{w}_{ij} \geq 0$ for all $ij \in A$), then x and y are feasible solutions to the primal and dual LPs, respectively, and x, y satisfy the CS conditions. Hence, x and y are primal solutions to the primal and dual LPs, respectively. □

3.4.2 Short Algorithm for Problem Solving

Given a connected digraph $D = (N, A)$ with arc costs w and node demands b,

1. Find a spanning tree T with a feasible tree flow x by inspection.
2. Calculate dual potentials for all $v \in N$.
   a. Pick an arbitrary node, say a, set $y_a = 0$.
   b. Solve $\bar{w}_{uv} = w_{uv} + y_u - y_v$ to get potentials for all other nodes.
   c. Shortcut: source + traffic = destination.
3. Calculate reduced costs for all non-basic arcs.
   a. $\bar{w}_{uv} = w_{uv} + y_u - y_v$.
   b. Use a three-column table for this step. Easy to check solutions.
4. Update the spanning tree.
   a. Find a non-basic arc uv with negative reduced cost.
   b. If no such arc exists, stop. The current solution is optimal
   c. Form and orient a cycle C by adding uv to T.
   d. If all arcs in C are forward, stop. The problem is unbounded.
   e. Find a backward arc pq with minimum flow in C.
   f. Push x_{pq} amount of units along C.

g. Update the spanning tree to $T - uv + pq$.

h. Go back to 2.

3.5 Network Simplex Initialization

Recall in standard Simplex, we find an initial feasible solution by constructing an auxiliary LP: (1) multiply -1 to appropriate rows to make $b \geq 0$, then (2) add one auxiliary variable for each constraint, and (3) solve the auxiliary LP using Simplex.

For Network Simplex, we add an extra row $\sum_{v \in |N|} z_v = 0$ containing all the auxiliary variables. Then the constraints behave like an incidence matrix corresponding to an auxiliary graph, where each supply node has an arc going into the auxiliary node z, and each demand node has an arc coming from z.

Since the objective function is $\min \sum_{v \in |N|} z_v$, each original arc has a cost of 0 and each auxiliary arc 1. This gives us an obvious starting feasible flow: transport all supplies to z then distribute to demand nodes from z.

3.5.1 Auxiliary Graph

Given a digraph $D = (N, A)$ and node demands $b \in \mathbb{R}^N$, the *auxiliary TP* $D' = (N', A')$ with $b' \in \mathbb{R}^{N'}$ and arc cost $w' \in \mathbb{R}^{A'}$ is

- $N' = N \cup \{z\}$ where z is a new node.
- $A' = A \cup \{vz \mid v \in N, b(v) < 0\} \cup \{zv \mid v \in N, b(v) > 0\}$.
- $b'(z) = 0$, $b'(v) = b(v)$ for all $v \in N$.
- $w'(e) = 0 \iff e \in A$ and $w'(e) = 1 \iff e \in A' \setminus A$.

As seen above, the auxiliary TP has a natural feasible flow (sending everything through z).

3.5.2 Feasibility Characterization

Theorem A TP is feasible if and only if its auxiliary TP has optimal value zero.

Intuition. Simplex initialization. □

Proposition Suppose TP with digraph $D = (N, A)$ is infeasible. If we partition N into $S_\_$ and S_+ according to the dual potentials of nodes, then

1. There does not exist any arc from $S_\_$ to S_+.
2. Any arc from S_+ to $S_\_$ is non-basic.

Proof. Assume that our TP is not feasible, i.e., the auxiliary TP has an optimal solution with strictly positive optimal value.

Let x^* be an optimal solution. Set $y_z = 0$. We can partition the nodes in N according to their *dual potentials*, which are either $+1$ or -1 (use definition for dual potentials and definition for w' to check this statement). Let $S\_$ be the set of nodes whose value (dual potential) is -1 and S_+ the set of nodes whose value is $+1$.

We now answer two questions. First, can there be an arc from $S\_$ to S_+? Next, if e is an arc from S_+ to $S\_$, do we know anything about x_e?

Let $u \in S\_$ and $v \in S_+$. If arc uv goes from $S\_$ to S_+, $\bar{w}_{uv} = w_{uv} + y_u - y_v = 0 - 1 - 1 = -2 < 0$. This contradicts the optimality of our solution (recall Simplex terminates when all reduced costs are positive). Thus, all arcs between $S\_$ and S_+ are in the same direction: from S_+ to $S\_$. Next, for an arc vu from S_+ to $S\_$, $\bar{w}_{uv} = 0 - (-1 - 1) = 2 \neq 0$. Then uv is not basic, i.e., not in our spanning tree, thus $x_{vu} = 0$. $\square$

Restricting our attention to arcs with non-zero flow, we get a digraph where all arcs are either leaving $S\_$ to z or leaving z to S_+.

Theorem A TP with digraph $D = (N, A)$ and node demand $b \in \mathbb{R}^N$ is infeasible if and only if there exists $S \subseteq N$ such that $b(S) < 0$ and $\delta(S) = \varnothing$. [Alternatively, we can say $b(S) > 0$ and $\delta(\bar{S}) = \varnothing$].

Intuition. If there exists a set of nodes with negative net demand and no arcs leaving, or a set of nodes with positive net demand and no arcs coming in, then the TP is infeasible.

Proof.

($\Longrightarrow$): Assume our TP is infeasible. By theorem, the AUX TP has an optimal solution with strictly positive optimal value. Let x be an optimal tree flow and y be the corresponding feasible potential with $y_z = 0$, where z is the auxiliary node. Partition N into $S\_$ and S_+ based on the dual potential. We claim that $S\_ := \{v \in N : y_v = -1\}$ satisfies the conditions $b(S\_) < 0$ (the net flow of $S\_$ is negative) and $\delta(S\_) = \varnothing$ (there is no arc leaving $S\_$).

- If $\delta(S\_) \neq \varnothing$, then there exists an arc uv such that $u \in S\_$, $v \in S_+$, but $\bar{w}_{uv} = w_{uv} + y_u - y_v = 0 - 1 - 1 = -2 < 0$. But this contradicts optimality. Thus, $\delta(S\_) = \varnothing$.
- Next, for any arc $uv \in \delta(\bar{S}\_)$, $\bar{w}_{uv} = w_{uv} + y_u - y_v = 0 + 1 + 1 = 2 \neq 0$ implies that uv is non-basic and hence $x_{uv} = 0$. Since some flow goes to z,

$$0 > x(\delta(\bar{S}\_)) - x(\delta(S\_)) = \sum_{v \in S\_} (x(\delta(\bar{v})) - x(\delta(v))) = \sum_{v \in S\_} b_v = b(S\_).$$

($\Longleftarrow$): Let $S \subseteq N$ where $b(S) < 0$ and $\delta(S) = \varnothing$. Suppose for a contradiction that there exists a feasible flow x. By feasibility,
$0 > b(S) = \sum_{v \in S} b_v = \sum_{v \in S} (x(\delta(\bar{v})) - x(\delta(v))) = x(\delta(\bar{S})) - x(\delta(S))$. But
$x \geq 0 \implies x(\delta(\bar{S})) \geq 0$, a contradiction. $\square$

4 Summary

4.1 LPs and CS Conditions

$$\begin{array}{ll} \min & w^T x \\ s.t. & x(\delta(\bar{v})) - x(\delta(v)) = b_v \quad (\forall v \in N) \\ & x \geq \mathbf{0} \end{array}$$

$$\begin{array}{ll} \max & b^T y \\ s.t. & y_v - y_u \leq w_{uv} \quad (\forall uv \in A) \end{array}$$

- $\bar{x}, \bar{y}$ are optimal solutions for the primal and dual if and only if $\bar{x}_{uv} = 0 \vee \bar{y}_v - \bar{y}_u = w_{uv}$ for each $uv \in A$.

4.2 Key Results

- Let M be the incidence matrix of a digraph $D = (N, A)$.
  - A set of columns are linearly dependent if and only if they correspond to a cycle.
  - A set of $|N| - 1$ columns of M is a basis if and only if the corresponding $|N| - 1$ arcs make up a spanning tree.
- A TP with digraph $D = (N, A)$ and node demand $b \in \mathbb{R}^N$ is **infeasible** if and only if there exists $S \subseteq N$ such that $b(S) < 0$ and $\delta(S) = \varnothing$.
- A feasible TP is **unbounded** if and only if there exists a negative dicycle.

5 Appendix

5.1 Additional Proofs

Lemma Show that if $G = (V, E)$ is a (undirected) graph in which every vertex has degree at least $\mathbf{2}$, then G contains a cycle C as a subgraph.

Proof. Suppose that G does not contain a cycle. Consider any connected component H of G. Since every vertex of G has degree at least $\mathbf{2}$, H has at least three vertices. Since H is connected and contains no cycles, H is a tree. Since H is a tree with at least two vertices, H has at least two vertices of degree one. But this contradicts the hypothesis that every vertex of G has degree at least $\mathbf{2}$. In fact, every connected component of G must contain a cycle. $\square$