Network Simplex Method

For network LP problems we have the standard LP

$$\begin{array}{rrcl} \min & c^T x & & \\ & Ax & = & b \\ & x & \geq & 0 \end{array}$$

where A is the node-arc incidence matrix of a directed graph G. The rows of A are linearly dependent, so we will assume $\sum_{i \in \mathcal{N}} b_i = 0$. We will also assume that G is connected and that $\mathcal{N} = \{1, 2, \ldots, n\}$. (If G is not connected we could instead solve problems separately for each connected component.) We'll see that then A has rank $n-1$. Let $\widetilde{A}$ be the $(n-1) \times |\mathcal{A}|$ matrix obtained from A by removing the last row, which we will refer to as the reduced node-arc incidence matrix of G. Similarly remove the last component of b to obtain $\widetilde{b}$, which yields the LP problem with constraints $\widetilde{A}x = \widetilde{b}$, which are equivalent to the original constraints.

What we would like to do now is relate the bases of $\widetilde{A}$ to graph-theoretic concepts in G, and in particular spanning trees.

Lemma 1. *Every tree on $n > 1$ nodes has at least two leaves, where a leaf is a node with degree 1.*

Proof. Suppose by way of contradiction that $G' = (\mathcal{N}, \mathcal{A}')$ is a tree on $n > 1$ nodes, with either zero leaves or just one leaf. Choose a node $i_0 \in \mathcal{N}$, which is the single leaf of the tree, or else an arbitrary node if there are no leaves. Start a walk at i_0, in which we leave every node (other than i_0) via a different arc than the one in which it was entered. This is possible because each of these nodes must have at least two incident arcs. Since there are a finite number of nodes, eventually a node must be repeated on this walk, say $i_k = i_\ell, k < \ell$, is the first repeat. Then the portion of the walk from i_k to i_ℓ is a cycle. □

Theorem 1. *For a directed graph $G' = (\mathcal{N}, \mathcal{A}')$, the following are equivalent:*

(a) G' is a tree (acyclic and connected);

(b) G' is acyclic and has $n-1$ arcs; and

(c) G' is connected and has $n-1$ arcs.

Note: This theorem may be thought of in terms of $(n-1)$-dimensional vectors, where acyclic is analogous to being linearly independent and being connected corresponds to spanning $\mathbb{R}^{n-1}$.

Proof. By induction. For $n = 1$ there is only one possible graph which is the single node. For this graph conditions (a), (b) and (c) trivially hold. Now assume the theorem is true for all graphs with fewer than n nodes, where $n > 1$.

(a) $\Rightarrow$ (b) Suppose G' is acyclic and connected. By Lemma 1, G' has a leaf. Remove this leaf and its single incident arc to get $\widetilde{G}'$, which is also acyclic and connected, and thus a tree. Hence by the inductive hypothesis, $\widetilde{G}'$ has $n-2$ arcs, and so G' has $n-1$ arcs. ✓

(b) $\Rightarrow$ (c) Let G' be acyclic with $n-1$ arcs. If G' is not connected, consider its connected components, $G'_1, \ldots, G'_p, p > 1$, each of which is connected and acyclic. By the induction hypothesis, each component G'_i has $|\mathcal{N}_i| - 1$ arcs, where $\mathcal{N}_i$ is its node set. So G' has $\sum_{j=1}^{p}(|\mathcal{N}'_j| - 1) = |\mathcal{N}| - p$ arcs, but this is less than $|\mathcal{N}| - 1$. ✓

(c) $\Rightarrow$ (a) Suppose G' is connected and has $n-1$ arcs, but contains a cycle. We can remove any arc that is part of this cycle and have the graph remain connected. By continuing to remove arcs from any existing cycles, we must eventually obtain a connected acyclic subgraph, which has $n-1$ arcs by the inductive hypothesis. However, before we removed any arcs the graph had $n-1$ arcs. ✓

□

Theorem 2. *Suppose $G' = (\mathcal{N}, \mathcal{A}')$ is a spanning tree of G, where $|\mathcal{N}| = n$. Let $\widetilde{A}$ be the reduced node-arc incidence matrix of G and $\widetilde{B}$ be the reduced node-arc incidence matrix of G'. Then the rows and columns of $\widetilde{B}$ can be permuted to make it triangular, with ± 1's on its diagonal.*

Proof. The theorem holds trivially for $n = 1$. For $n > 1$, Lemma 1 shows that G' has at least two leaves. Choose a node other than n and label it 1. Now remove node 1 and its incident arc, and find a leaf (other than n) of the resulting tree and label it 2. Continue in this fashion until all nodes have been labelled.

Next we label the arcs. We will label an arc (i, j) or (j, i) with $i < j$ as i. That is, we label arcs by the endpoint that is farther from the root n, which implies that all the labels are unique. This also means that all paths to the root n have increasing labels. An example is shown in Figure 1. Now consider a column of $\widetilde{B}$ corresponding to arc (i, j) or $(j, i) \in \mathcal{A}'$,

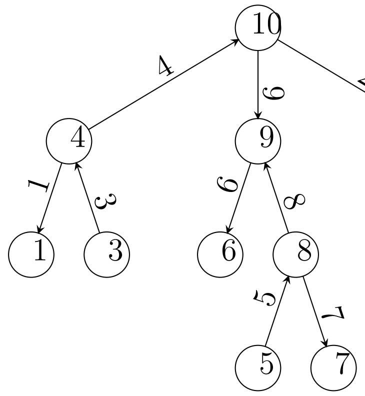

Figure 1: A sample tree labelled according to the conventions described in Theorem 2.

with $i < j$. There will be a $+1$ entry in the row corresponding to the tail, and a -1 entry in the row corresponding to the head, and all the other entries will be zero. Thus the only nonzero entries will be in rows i and j, and if $j = n$ there will be no entry there. Since $j > i$ we have that $\widetilde{B}$ is triangular. For the example corresponding to Figure 1 we have

$$\widetilde{B} = \begin{bmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & +1 & 0 & 0 & 0 & 0 & 0 & 0 \\ +1 & 0 & -1 & +1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & +1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & +1 & +1 & 0 \\ 0 & 0 & 0 & 0 & 0 & +1 & 0 & -1 & -1 \end{bmatrix}.$$

□

Corollary 1. *If G is connected, then it has a spanning tree and $\widetilde{A}$ has rank $n-1$.*

Proof. We can see G has a spanning tree by the same argument used to prove (c) $\Rightarrow$ (a) in Theorem 1. Then the corresponding columns of $\widetilde{A}$ form a nonsingular $(n-1) \times (n-1)$ matrix, which has full row rank. □

Corollary 2. *Every basis matrix of $\widetilde{A}$ corresponds to arcs of a spanning tree.*

Proof. Every spanning tree corresponds to a basis matrix of $\widetilde{A}$ by Theorem 2. Suppose we have a basis, which is $n-1$ columns from $\widetilde{A}$ that are linearly independent. The corresponding arcs cannot contain a cycle. (If they did, then we could add and subtract columns to get zero, which would imply they are not linearly independent.) Thus we have $n-1$ arcs with no cycle, which corresponds to a tree. □

Corollary 3. *If all b_i's are integer then every basic solution of $\widetilde{A}x = \widetilde{b}$ is integer-valued.*

Proof. We can solve $\widetilde{B}\bar{x}_{\widetilde{B}} = \widetilde{b}$ by adding and subtracting only, so the solution consists of all integers. □

Note that this last corollary shows that every basic solution to the assignment problem corresponds to a permutation.

Network Simplex Method

How do the steps of the usual primal simplex algorithm specialize in the network case? The basic feasible solution is given by $\bar{x} = \begin{pmatrix} \bar{x}_{\widetilde{B}} \\ \bar{x}_{\widetilde{N}} \end{pmatrix}$, where $\widetilde{B}\bar{x}_{\widetilde{B}} = \widetilde{b}$. So we can solve for $\bar{x}_{\widetilde{B}}$ easily. Indeed, we can calculate the flow on basic (tree) arcs sequentially, starting from the leaves of the tree and working up to the root.

We can compute the corresponding dual solution $\bar{y}$ from $\widetilde{B}^T\bar{y} = c_{\widetilde{B}}$ or $B^T\bar{y} = c_{\widetilde{B}}$ if we set $\bar{y}_n = 0$. When (i, j) is a tree arc we have

$$\begin{aligned} a_{ij}^T\bar{y} &= c_{ij} \\ \text{or } \bar{y}_i - \bar{y}_j &= c_{ij} \end{aligned}$$

For example, given $\bar{y}_{10} = 0$ and $c_{4,10}$, we get $\bar{y}_4 = c_{4,10} + \bar{y}_{10}$. We then solve sequentially, working down from the root and test the optimality of the obtained solution. If the solution is not optimal we find a non-tree arc, say (i, j) (which corresponds to the entering non-basic index "q") with $\bar{c}_{ij} = c_{ij} - \bar{y}_i + \bar{y}_j < 0$. We can also solve $\widetilde{B}\bar{a}_{ij} = \widetilde{a}_{ij}$ for $\bar{a}_{ij}$, which will correspond to ± 1 entries on the arcs in the unique path connecting i to j in the current basic spanning tree. There will be a $+1$ entry on arcs that are traversed forwards and a -1 entry on arcs that are traversed backwards to get from i to j.

Note that there are *polynomial-time* variants of the network simplex method: see e.g., J.B. Orlin, S.A. Plotkin, and E. Tardos, Mathematical Programming 60 (1993) 255–276 for the dual simplex algorithm, and J.B. Orlin, Mathematical Programming 78 (1997) 109–129 for the primal.