Chapter 7

Minimum Cost Flows

7.1 Introduction

Let $G = (N, A)$ be a directed graph with a cost c_{ij} and a capacity u_{ij} associated with every arc $(i, j) \in A$. We associate, with each node $i \in N$, a number b_i that indicates its supply (if $b_i > 0$) or demand (if $b_i < 0$). Nodes featuring $b_i = 0$ are called transshipment nodes. The Minimum Cost Flow Problem (MCFP) calls for finding a min-cost flow to satisfy the demands from the supplies without exceeding arc capacities. The MCFP can be stated as

$$z(\boldsymbol{x}) = \min \sum_{(i,j)\in A} c_{ij} x_{ij} \tag{7.1a}$$

$$\text{s.t.} \sum_{j \in N : (i,j) \in A} x_{ij} - \sum_{j \in N : (j,i) \in A} x_{ji} = b_i \quad \forall i \in N \tag{7.1b}$$

$$0 \le x_{ij} \le u_{ij} \quad \forall (i, j) \in A \tag{7.1c}$$

Let $n = |N|$ and $m = |A|$. We make the following assumptions:

Assumption 7.1. *All data (cost, supply/demand, and capacity) are integral.*

Assumption 7.2. *The graph is directed.*

Assumption 7.3. *The supplies/demands at the nodes satisfy the condition* $\sum_{i \in N} b_i = 0$, *and the MCFP has a feasible solution.*

Assumption 7.4. *All arc costs are non-negative.*

The algorithms to solve the MCFP rely on the concept of residual graphs, as defined in Chapter 5. The residual graph $G(\boldsymbol{x})$ corresponding to a flow $\boldsymbol{x}$ is defined as follows. We replace each arc $(i, j) \in A$ by two arcs (i, j) and (j, i). The arc (i, j) has cost c_{ij} and residual capacity $r_{ij} = u_{ij} - x_{ij}$, and the arc (j, i) has cost $c_{ji} = -c_{ij}$ and residual capacity $r_{ji} = x_{ij}$. The residual graph consists of arcs with positive residual capacity only.

7.2 Applications

7.2.1 Distribution Problems

A car manufacturer has several manufacturing plants and produces several car models at each plant that it then ships to geographically dispersed retail centers throughout the country. Each plant has a maximum production capacity for each car model. Each retail center requests a specific number of cars of each model. The firm must determine the production plan of each model at each plant and a shipping pattern that satisfies the demands of each retail center and minimizes the overall cost of production and transportation knowing that there are bounds on the number of cars of each model that can be shipped from each plant and each retailer. Two types of costs must be considered: a production cost for each car model and each plant, and a transportation cost for each car model between each plant/retailer pair. We assume that the total production capacity of all the plants for each car model exceeds the sum of the retailer demands of the corresponding car model.

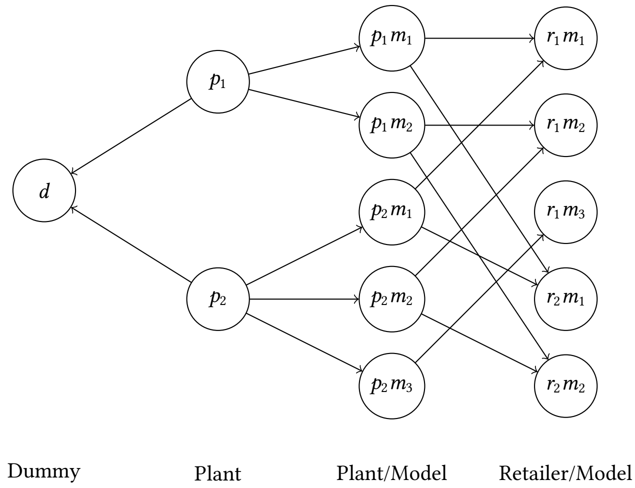

Figure 7.1: Production-distribution graph

Figure 7.1 illustrates a situation with two manufacturing plants, two retailers, and three car models, where the third car model cannot be produced at the first plant and the first retailer does not require any car of the third model. This graph has four types of nodes: (i) plant nodes, representing various plants; (ii) plant/model nodes, corresponding to each model made at a plant; (iii) retailer/model nodes, corresponding to the models required by each retailer; and (iv) a dummy node to handle the extra production capacity of the plants.

The b_i value of each plant node is equal to the sum of the production capacities of the car models that can be produced there. Plant/model nodes are transshipment nodes. The b_i value of each retailer/model node is equal to the opposite of the request of the corresponding retailer of the corresponding car model. The b_i of the dummy node is equal to the sum of all the requests of the retailers minus the sum of the production capacities of the plants.

The graph contains three types of arcs:

Production arcs These arcs connect a plant node to a plant/model node; the cost of this arc is the cost of producing the model at that plant. We place upper bounds on these arcs to control for the max production of each particular car model at the plants.

Transportation arcs These arcs connect plant/model nodes to retailer/model nodes; the cost of such an arc is the total cost of shipping one car from the manufacturing plant to the retail center. These arcs have upper bounds imposed on their flows to model contractual agreements with shippers or capacities imposed on any distribution channel.

Dummy arcs These arcs plant nodes to the dummy nodes. These arcs have zero costs and infinite capacity.

The production and shipping schedules for the automobile company correspond in a one-to-one fashion with the feasible flows in this graph. Consequently, a min-cost flow yields an optimal production and shipping schedule.

7.2.2 Balancing of Schools

Suppose that a municipality has a set S of schools. We divide the municipality into a set L of districts, and let b_i and g_i denote the number of boys and girls at district $i \in L$. Let d_{ij} be a distance measure that approximates the distance any student at district $i \in L$ must travel if he or she is assigned to school $j \in S$. School $j \in S$ can enroll u_j students. Finally, let $\overline{p}$ denote an upper bound on the percentage of boys assigned to each school (we choose these numbers so that school $j \in S$ has same percentage of boys as does the municipality). The objective is to assign students to schools in a manner that maintains the stated boys/girls balance and minimizes the total distance traveled by the students. For the sake of simplicity, we assume that $\sum_{i \in L}(b_i + g_i) = \sum_{j \in S} u_j$.

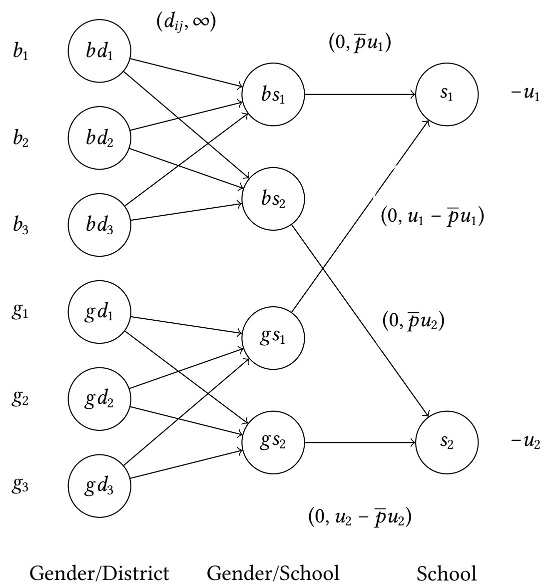

Figure 7.2: Graph for the school balancing problem

We model this problem as a MCFP. Figure 7.2 shows the graph for a three-district, two-school problem. We model each district $i \in L$ as two nodes bd_i and gd_i (which are the supply nodes) and each school $j \in S$ as three nodes bs_j, gs_j, and s_j - nodes s_j are the demand nodes. The decision variables are the number of boys assigned from district $i \in L$ to school $j \in S$ (which we represent by an arc from node bd_i to node bs_j) and the number of girls assigned from district $i \in L$ to school $j \in S$ (which we represent by an arc from node gd_i to node gs_j). These arcs are uncapacitated, and we set their per unit flow cost equal to d_{ij}. For each $j \in S$, we connect the nodes bs_j and gs_j to the school node s_j. The flow on the arcs (bs_j, s_j) and (gs_j, s_j) denotes the total number of boys and girls assigned to school j. Since each school must satisfy upper bounds on the number of boys it enrolls, we set the arc capacity of the arc (bs_j, s_j) equal to $\overline{p}u_j$. For each school, arc (gs_j, s_j) has a capacity of $u_j - \overline{p}u_j$.

Please notice that the proposed solution, represented in Figure 7.2, is not the only correct way of modeling the problem as a MCFP.

7.2.3 Optimal Loading of Hopping Airline

A commuter airline uses a plane, with a capacity to carry p passengers, on a "hopping flight", as shown in Figure 7.3. The hopping flight visits the cities 1, 2, 3, ..., n, in a fixed sequence. The plane can pick up passengers at any node and drop them off at any other node. Let b_{ij} denote the number of passengers available at node i who want to go to node j, and let f_{ij} denote the fare per passenger from node i to node j. The airline would like to determine the number of passengers that the plane should carry between the various origins to destinations to maximize the total fare per trip while never exceeding the plane capacity.

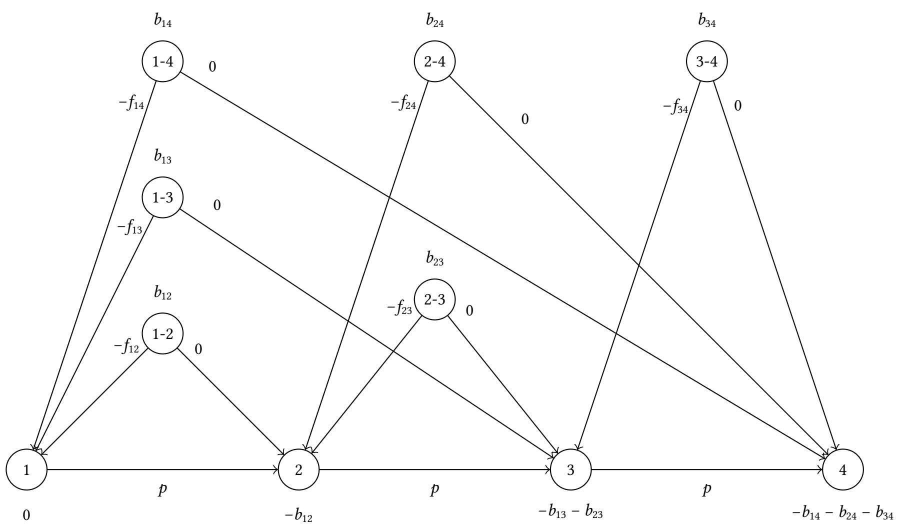

Figure 7.3: Formulating the hopping plane flight problem as a MCFP: labels next to nodes indicate b_i, labels above arcs the arc cost c_{ij}, labels below arcs the arc capacity u_{ij}

Figure 7.3 shows a MCFP formulation of this hopping plane flight problem. The graph contains data for only those arcs with nonzero costs and with finite capacities: any arc without an associated cost has a zero cost; any arc without an associated capacity has an infinite capacity. Consider, for example, node 1. Three types of passengers are available at node 1, those whose destination is node 2, node 3, or node 4. We represent these three types of passengers by the nodes $1-2$, $1-3$, and $1-4$ with supplies b_{12}, b_{13}, and b_{14}. A passenger available at any such node, say $1-3$, either boards the plane at its origin node by flowing through the arc $(1-3, 1)$, and thus incurring a cost of $-f_{13}$ units, or never boards the plane which we represent by the flow through the arc $(1-3, 3)$.

Exercise 7.1.

> Centralized teleprocessing graphs often contain many (as many as tens of thousands) geographically dispersed terminals. These terminals need to be connected to a CPU either by direct lines or though concentrators. Each concentrator is connected to the CPU through a high-speed, cost-effective line that is capable of merging data flow streams from different terminals and sending them to the CPU. Suppose that the concentrators are in place and that each concentrator can handle at most K terminals. For each terminal j, let c_{oj} denote the cost of laying down a direct line from the CPU to the terminal, and let c_{ij} denote the line construction cost for connecting concentrator i to terminal j. The decision problem is to construct the min-cost graph for connecting the terminals to the CPU. Formulate this problem as a MCFP.

Construct the graph $G = (\{s\} \cup N_c \cup N_t, A)$, where the node s represents the CPU and has a supply of $|N_t|$ units (i.e., $b_s = |N_t|$), each node $i \in N_c$ denotes the concentrator i ($b_i = 0$), and each node $j \in N_t$ denotes the terminal j and has a demand of 1 unit (i.e., $b_j = -1$). The arc set A contains three types of arcs: (i) arcs $\{(s, j) \mid j \in N_t\}$ of cost c_{oj} and capacity 1 representing which terminals are directly connected to the CPU; (ii) arcs $\{(i, j) \mid i \in N_c, j \in N_t\}$ of cost c_{ij} and capacity 1 representing the terminals connected to the CPU via the concentrators; (iii) arcs $\{(s, i) \mid i \in N_c\}$ of cost 0 and capacity K representing the number of terminals connected to the CPU via concentrator i. A min-cost flow in this graph determines a min-cost graph for connecting the terminals to the CPU.

Exercise 7.2.

> A library facing insufficient primary storage space for its collection is considering the possibility of using secondary facilities to store portions of its collection. These options are preferred to an expensive expansion of primary storage. Each secondary storage facility has limited capacity and a particular access costs for retrieving information. Through appropriate data collection, we can determine the usage rates for the information needs of the users. Let q_j denote the capacity of storage facility j, and v_j denote the access cost per unit item from this facility. In addition, let a_i denote the number of items of a particular class i requiring storage and let u_i denote the expected rate (per unit time) that we will need to retrieve books from this class. Our goal is to store the books in a way that will minimize the expected retrieval cost. Show how to formulate the problem of determining an optimal policy as a MCFP.

Construct the graph $G = (C \cup F \cup \{t\}, A)$, where C contains one node corresponding to each book class, F contains one node corresponding to each storage facility, and t is a dummy demand node. The arc set A consists of two arc sets, A_1 and A_2 (i.e., $A = A_1 \cup A_2$) defined as follows: $A_1 = \{(i, j) \mid i \in C, j \in F\}$ and $A_2 = \{(j, t) \mid j \in F\}$. The cost of each arc $(i, j) \in A_1$ is equal to $u_i \cdot v_j$, and the capacity of these arcs is infinite. The cost of each arc $(j, t) \in A_2$ is zero, but their capacity is equal to q_j. The supply of each node $i \in C$ is equal to a_i (i.e., $b_i = a_i$) while the demand of node t is equal to $\sum_{i \in C} a_i$ (i.e., $b_i = -\sum_{i \in C} a_i$). Nodes of the set F are transshipment nodes.

A minimum cost flow is found in the network thus constructed. The flow on the arc $(i, j) \in A_1$ for $i \in C$ and $j \in F$ denotes the number of books in the ith class assigned to the jth facility.