物流系统分析
Logistics System Analysis

第 7 周　一到一配送问题（3）–网络设计问题
One-to-One Distribution – Network Design

葛乾

西南交通大学　系统所 & 交通工程系

目 录 |Outline

Network Design Issues

In all the scenarios discussed so far, the items followed a predetermined path. Real logistics problems, however, often involve the choice of alternative routes (e.g., alternative ways of shipping) between origins and destinations, in addition to the choice of when and how much to dispatch. In some instances one may even be interested in whether certain routes should be provided at all; or even in the design of an entirely new physical distribution network.

We also found that there were economies of scale in flow; i.e., the optimal cost per item decreased with D'. In the following lectures, we will have to consider logistics problems with multiple destinations, where an item's route is not predetermined and cost decreases with flow. We discuss here some key features of these problems, and conclude this lecture with a comparison of detailed and non-detailed approaches for logistic system design.

目 录 |Outline

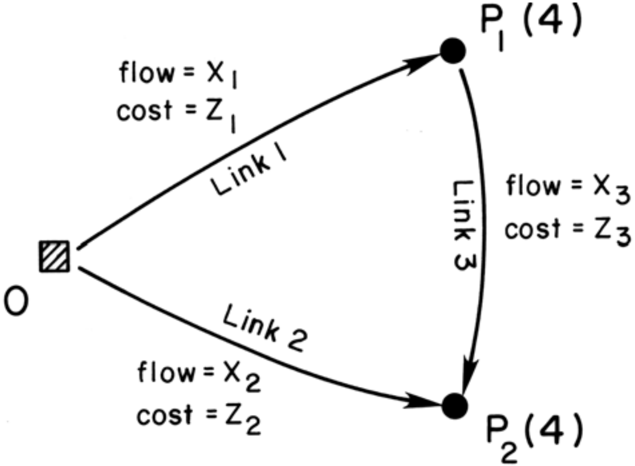
P1 (4)
flow = X1
cost = Z1
Link 1
Link 3
flow = X3
cost = Z3
O
Link 2
flow = X2
cost = Z2
P2 (4)

- A simple example with one origin and two destinations effectively illustrates the properties of optimal system designs with and without flow economies of scale.
- The origin, O, produces items of type $i (i = 1, 2)$ for destination P_i at a constant rate, given by the parenthetical (括号的) numbers in the figure: $D'_1 = D'_2 = 4$ items per unit time. The combined production rate at the origin is $D'_1 + D'_2 = 8$ items/unit time. The arrows in the figure depict possible shipment trips; these transportation links are numbered 1, 2, 3. While all the items traveling to P_1, must travel directly between O and P_1, the items traveling to P_2 may go either directly or via P_1.
- Let us assume that a fraction (to be decided) x, of the items for P_2 are sent via P_1 and the rest are shipped directly. This establishes a flow $x_1 = 4(1 + x)$ on link 1 (OP_1), a flow $x_3 = 4x$ on link 3 (P_1P_2) and a flow $x_2 = 4(1 - x)$ on link 2 (OP_2).

We also assume that the total cost on the network can be expressed as a sum of link costs, and that these depend only on their own flows. This is a reasonable assumption if no attempt is made to coordinate the shipping schedules on the three links, as then the prorated cost to link should be close to the EOQ expression with demand rate equal to the link flow. Thus, if we let $z_i(x_i)$ denote the cost per item on link i when the flow is x_i, the total system cost per unit time is:
With economies of scale, the functions $x_i z_i(x_i)$ increase at a decreasing rate (are concave) as in

$$TC = \sum_{i=1}^{3} x_i z_i(x_i).$$

Economy of scale

With economies of scale, the functions $x_i z_i(x_i)$ increase at a decreasing rate (are concave). Because the x_i's are linear in the split x, the total cost is a concave function of the split – this (concave) dependence of cost on splits (decision variables) also holds for general networks. Suppose, for example, that

$$z_1 = x_1^{-1/2}, z_2 = 3x_2^{-1/2}, \text{ and } z_3 = 1;$$
$$x_1 z_1 = x_1^{1/2}, x_2 z_2 = 3x_2^{1/2}, \text{ and } x_2 z_3 = x_3;$$

Then, as a function of x, the total cost is calculated as:

$$TC = 2(1+x)^{1/2} + 6(1-x)^{1/2} + 4x$$

The total cost is a concave function of the split, x. For our data the optimal solution is $x^* = 1 \rightarrow$ everything should be shipped through P_1. The total cost is 6.8. Although shipping everything direct may be better for different data, clearly one would never want to split the flow to P_2 among the two routes (OP_2 and OP_1P_2).

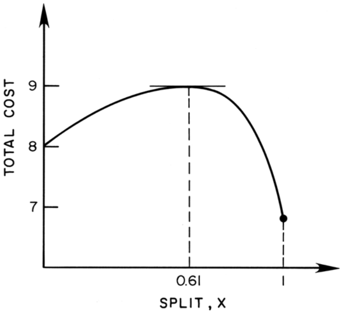

- A similar "all-or-nothing" principle holds for networks with multiple ODs if the total cost is a concave function of all the link flows. In that case all the flow from any origin to any destination should be allocated to only one route. This is not difficult to see: one can define a split between any two routes joining an origin and a destination, and since the link flows are linear in that split, the total cost is concave in the split; thus, only one of the routes can carry flow. Networks with diseconomies of scale behave in an opposite manner.
- In that case the total cost function is convex in the splits and there is an incentive to spread out the flow among routes. In fact, if for a one origin and one destination network, there exist several routes with identical cost functions (with dis-economies); it is not difficult to prove that the total flow should be evenly divided among all the routes.

- Networks with flow economies of scale also respond in a different manner to changes in conditions. While, with diseconomies, a small improvement to one of the routes would lead to a small change in the optimal flow distribution, with economies, the optimal flows either stay the same or change by a discrete amount. This can be seen with the example. As long as $z_3 < [2-2^{-1/2}] \approx 1.3$, x^* equals 1, but if z_3 is increased beyond this value ever so slightly, the solution jumps to $x* = 0$.
- This is typical of concave cost problems: minor changes to the input data can induce large changes in the optimal solution. Fortunately, the cost does not behave in such manner; despite the jump in our example the cost is a continuous function of z_3.

$$TC^* = \begin{cases} 8 & \text{if } z_3 \leqslant 2 - 2^{-1/2} \\ 8^{1/2} + 4z_3 & \text{if } z_3 > 2 - 2^{-1/2} \end{cases}$$

目 录 |Outline

The nature of the solution is not the only difference between networks with economies and diseconomies; the way to find it is also different. While networks with diseconomies are well behaved optimization problems without local minima that are not global, networks with economies are not.

Local searches

Except for technical details, all local search algorithms work in the same manner.

1. The total cost is evaluated for an initial feasible solution, described by a set of variables that uniquely identify the decisions; e.g., the set of splits for all origin destination pairs.
2. A small cost-reducing perturbation to the feasible solution (e.g., a differential change to the splits) is then sought. If not found, the search stops because the initial solution is a local minimum; i.e., a solution that cannot be improved without substantial changes. Otherwise, an improved larger perturbation obtained from the original small perturbations is identified, and then used to construct a new improved feasible solution.

The process is then iterated (seeking small cost-reducing perturbations to the new solution, etc.) until no significant improvements result.

Problems of local search algorithms

- Although local search algorithms can be used to find near optimal solutions for large detailed networks with convex costs, the same procedures fail with concave cost networks. The task is then much more complicated, and the network sizes that can be handled by numerical methods much smaller $\rightarrow$ 对于大网络的凸成本函数，局部搜索算法可以求得近似最优解；但是不适用于凹目标函数
- Local search techniques work acceptably for networks with scale diseconomies, because in those instances any local minimum is a global minimum. Unfortunately, this is not the case with economies of scale. Our simple problem has two local minima: $x = 0$ and $x = 1$. If a local search algorithm is applied to our example, any starting solution with $x < 0.61$ (the maximum in the figure) will converge suboptimally to $x = 0$.
- While for our simple example this can be corrected simply by starting with different x's, the task is daunting for large, highly detailed networks. In that case, the number of potential traps for a local search –all local minima regardless of cost –increases exponentially with the amount of detail.

An example

The items from a large number N of origins are shipped to one destination using two transportation modes (1 and 2). We use x_i to denote the split of production from origin i sent on mode 1, and assume that (to satisfy an agreement with the providers of type-1 transportation) each x_i must satisfy $x_i > h_i$ for some constant $h_i > 0$. Transportation by mode 2 is assumed to be more attractive, but limited in capacity; that is, the sum of the x_i's must exceed a value h.
For a set of splits to be feasible, thus, the following must be true:

$$\sum_{i=1}^{N} x_i \geqslant h, \text{ and } h_i \leqslant x_i \leqslant 1, \forall i$$

We seek the set of feasible splits that minimize the total cost, or equivalently the penalty paid because not all the items can be shipped by mode 2.

Solution of the example

The penalty paid for each origin is assumed to increase with x_i, except at certain values where a fixed amount ϵ_i is saved – perhaps because shipments can then be multiples of a box, requiring less handling*. To simplify the exposition, let us assume that there is only one such value δ_i for every origin, and that away from this value the penalty equals x_i; otherwise the penalty is $x_i - \epsilon_i$. If we define $\epsilon_i(x_i)$ to be: ϵ_i if $x_i = \delta_i$ and 0 otherwise, then the combined penalty for all the origins can be expressed as:

$$\sum_{i=1}^{N}[x_i - \epsilon_i(x_i)]$$

Note that each one of the terms in this summation for which $\delta_i > h_i$ exhibits two local minima in the range of feasibility $[h_i, 1]$: $x_i = h_i$ and $x_i = \delta_i$.

*假设起点处的惩罚费用随选择第一种方式的比例增大而增大，仅当该比例等于某个特定值时，可节省 ϵ_i 的费用。或因配送可由多个箱子完成，节省了处理费用。

Solution of the example

Any combination of x's, each equaling either h_i or δ_i, and satisfying the constraints is a local minimum, which could stop a search. If the δ_i and the h_i are uniformly distributed between 0 and 1, and h is small, there will be $O(2^{N/2})$ local optima. With so many traps, local search algorithms are doomed to failure for this problem –not because the penalty is discontinuous, but because it is not convex. A different method must be used.

Certainly, one could search exhaustively over all the possible solutions with a combinatorial tool such as branch and bound, but these methods can only handle problems of small size –typically with $O(10^2)$ decision variables or less.

Heuristic method

Alternatively, one could try to exploit the peculiar mathematical structure of the total penalty –or whichever problem is at hand –to develop a suitable algorithm. If successful, the approach would find a solution with all its detail. In our case, the optimization problem can be reduced to a knapsack problem that can be solved easily; in other instances it may be possible to decompose the problem into a collection of small easy problems. Very often, however, a simple solution method cannot be found.

In our case, this would happen if there were more than one (ϵ_i, δ_i) for each origin. Traditionally one then turns to *ad hoc*(为特定目的的) intuitive solution methods (known as heuristics) which one hopes will yield reasonable solutions.

Simplifying the problem

- There is also another approach. If while inspecting the formulation, or even better in the process of formulating the problem, one realizes that certain details are of little importance one should leave them out. Our example illustrates how removing minor details can turn a nightmare into an easy problem.
- If the ϵ_i's are so small that the $\epsilon_i(x_i)$ can be neglected, then the objective function reduces to $\sum_i x_i$. Former sources of difficulty, the ϵ_i and δ_i no longer enter the formulation. With less detail, the problem becomes well behaved (convex), and even admits a closed form solution; e.g., if $\sum_{i=1} h_i \geqslant h$ then the optimal splits are $x_i = h_i$ and the total cost is $\sum_{i=1} h_i = N\bar{h}$.
- Note that the optimal cost is given by an average (there is no need to know precisely each individual h_i in order to estimate the optimal cost), and that the optimal solution can be described with the simple rule "make every split as small as possible", which can be stated without making reference to the h_i's.

目 录 |Outline

In the following lectures we will seek solutions to logistic problems using as little details as possible, describing (as in the example) the solution in terms of guidelines which are developed based on broad averages instead of detailed data. We recognize that the solutions obtained from such guidelines may benefit from fine-tuning once detailed data become available; but also note that incorporating all the details into the model early will increase the effort for gathering data and may even get in the way of obtaining a good solution.

Similarity of the logistics system with trees

Observation of mother nature's logistics networks suggests that many logistics systems can be designed in this manner. Trees can be viewed as a logistic system for carrying nutrients from the soil to an above-ground region (the leaves) to meet the sun's rays. While every individual tree of a species is distinct from other individuals, we also see that the members of a species share many common characteristics on average. There is order at the macroscopic level. This is not surprising, since members of the species have adapted to similar environmental conditions, also filling the same niche (发挥相同作用) in the eco-system.

Similarity of the logistics system with trees

The detailed characteristics of an individual tree are (like our logistic systems) developed from two levels of data in two different ways:

- Members of the same species share a genetic code, which has evolved in response to the typical or average conditions that can be expected. This code is analogous to the guidelines of a simple model; e.g., "make each split as small as possible."
- In response to the detailed conditions of its location, a tree develops an individuality within the guidelines of the genetic code, better to exploit the local conditions. This would be analogous to the fine tuning that could have taken place if the ϵ_i, h_i, and δ_i had been given in our example.

The same could be said for other logistic systems encountered in nature, such as the circulating and nervous systems of the human body.

Similarities

On further inspection we notice that, not only average characteristics, but some specific traits are also the same for all individuals, (e.g., some tree species have always one trunk, all humans have one aorta artery （主动脉）, etc.). It is as if nature had decided that these items of commonality are optimal for almost any conditions that can be encountered; therefore, that part of the design is not open to fine tuning. Perhaps the same can be said of logistics systems.

The logistics systems of nature also have economies of scale. It takes less energy to move a certain flow through one single pipe than through two pipes with one-half the cross section. As in our networks with concave costs, there is an incentive to consolidate flow into single routes that can handle great volumes efficiently. Nature has responded to this challenge by evolving hierarchical systems of conveyance, such as the three hierarchy network.

The 'logistics system' of a tree

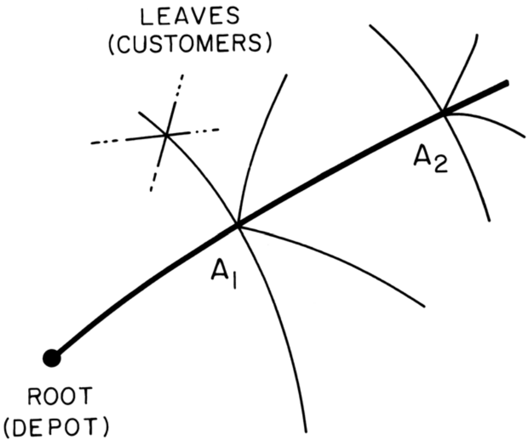

Scientists have begun to realize that apparently very complex ("fractal") structures, such as a fern leaf (蕨类植物的叶子), can be replicated and/or described with just a few rules and parameters. For the example of the figure, the separations between "nodes" (e.g., A1 and A2) for each hierarchy might be found to be relatively constant, perhaps varying with the distance from the root, as might be the number of branches at every node and the relative size of the main and secondary branches at nodes of the same hierarchy. The latter may also vary with the distance from the "root."†

A physical distribution network should probably be organized in a similar way with the root becoming the depot, the leaves the customers, and the nodes intermediate transshipment centers or terminals. Physical distribution networks that serve similar purposes, just as in nature, should likely share the same hierarchical organization and overall traits even if the specific details differ. As in nature, it should be possible to describe their near optimal configuration with just a few simple rules and parameters.

† 两个节点之间的间隔也许恒定或者随着到根部的距离变化。每个节点处的分支数及节点内部相同层次的主次分支的相对大小也呈现相似规律。后者也会随着距离根部的距离而变化。

Plans for further lectures

In this spirit, our lectures that follow will try to get at the "genetic code" of logistics systems; i.e., describe how general classes of logistics networks should be organized, with guidelines for obtaining an optimal structure developed without using detailed data. Building on the simple EOQ model, we gradually consider more complex systems.

- 3-4 lectures for "one-to-many" problems, i.e., problems with a single hierarchy consisting of one origin and many destinations (or the reverse)
- 3-4 lectures for "one-to-many" problems with transshipments (multiple hierarchies)
- 3-4 lectures for the "many-to-many" problems.

Readings

- Daganzo. Logistics System Analysis. Ch.3. Page 75-83

Any questions?