物流系统分析
Logistics System Analysis

第 9 节　一到多配送问题 (2)–同质顾客的策略
One-to-Many Distribution— Strategies for Identical Customers

葛乾

西南交通大学　系统所 & 交通工程系

目录 |Outline

1 Review

2 Identical Customers and Fixed Vehicle Loads

3 Identical Customers and Vehicle Loads Not Given

4 Implementation Considerations

Quick flashbacks

- The total distance for systems with many vehicle tours

$$\text{total distance } \approx \int_{\boldsymbol{R}} [\frac{2}{C} r(\boldsymbol{x}) + k\delta^{-1/2}(\boldsymbol{x})] \delta(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}.$$

- The total distance for systems with few vehicle tours

$$\text{total distance} \approx k' N \delta^{-1/2} = k' \sqrt{N|\boldsymbol{R}|}.$$

其中 $\boldsymbol{R}$ 表示区域，$|\boldsymbol{R}|$ 为其面积；N 为总顾客数，C 为单位分区内的顾客数，$r(\boldsymbol{x})$ 为从配送中心到各点（最短路）距离的分布，kk' 为参数；$\delta(\boldsymbol{x})$ 为单位面积内的顾客数；$\delta(\boldsymbol{x})$ 为单位面积内的顾客数；$\delta^{-1/2}(\boldsymbol{x})$ 为单位面积内相邻顾客的距离；

目录 |Outline

1 Review

2 Identical Customers and Fixed Vehicle Loads
- Very cheap items: $c_i \ll c_r$
- More expensive items: $c_i \gg c_r$
- Inventory at the origin

3 Identical Customers and Vehicle Loads Not Given

4 Implementation Considerations

- We first consider strategies where the loads carried by each vehicle are given. Since one would then operate the smallest possible vehicles able to carry the loads, we will denote by $v_{\max}$ the load size used.
- Given $D_n(t) = D(t)$ for t in $[0, t_{\max}]$, we seek the **dispatching times** $\{t_l : l = 0, \ldots, L\}$ and **vehicle routes** which minimize the **total logistics cost**. We let $t_0 = 0$ and $t_l \leqslant t_{l+1}$.
- Because all the customers are alike, there is no compelling reason to treat some differently from others, and we shall assume that *every customer is visited with every dispatch l*. Under these conditions, the search for the t_l is facilitated considerably because, **the transportation cost only depends on the number of dispatches**, L.

Decomposition

- We now show that for a given number of dispatches L, the total transportation cost between $t = 0$ and $t = t_{\max}$ is independent of the headways: $H_l = t_l - t_{l-1} (l = 1, \ldots, L)$.
- We have already stated that the transportation cost for a given l is a linear function of # routes, # delivery stops, # items carried and the total distance.
- Clearly, the combined cost for all l must also be a function of these four descriptors. Because vehicles travel full, three of these (the **total number of items** $D(t_{\max})N$, **the number of vehicle tours** $D(t_{\max})N/v_{\max}$, and **the total number of delivery stops** NL) are fixed; they do not depend on when or how much is shipped at each t_l.

► Recall that *the cost for n shipments* $\approx c_s(1 + n_s)n + c_s nd + c_s' V$, c_s the stop cost; c_d vehicle cost for each mile traveled; c_s' added cost of carrying an extra item

Total combined transportation cost

- For a given L, the total combined distance for all dispatches is also independent of the t_l.
- As indicated by the VRP formula, it is the sum of a local distance term proportional to the total number of stops made NL, $kLNE(\delta^{-1/2})$, and a line-haul component which is proportional to the (fixed) number of vehicle tours: $2E(r) \times \#\text{tours} = 2E(r)D(t_{\max})N/v_{\max}$. Note that the line-haul component is independent of L
- With the cost coefficients, the **total transportation cost combined** between $t = 0$ and $t = t_{\max}$ is approximately:

$$c_s N \left\{ \frac{D(t_{\max})}{v_{\max}} + L \right\} + c_d kLNE(\delta^{-1/2}) + c_d 2E(r) \frac{D(t_{\max})N}{v_{\max}} + c_s' D(t_{\max})N,$$

which only depends on one decision variable, L

- The last three terms are obvious. Recall $C = \lfloor \frac{v_{\max}}{v} \rfloor$ is the number of stops that a vehicle visits and $D(t_{\max}) = vL$. We have the 'stop' cost to be $c_s L \lceil \frac{N}{C} \rceil (1 + C) \approx c_s N \left\{ \frac{D(t_{\max})}{v_{\max}} + L \right\}$

- An expression based on few tour formula instead of the many tour one would be quite similar, and also independent of the $\{t_l\}$.

Total combined transportation cost (cont.)

$$c_s N \left\{ \frac{D(t_{\max})}{v_{\max}} + L \right\} + c_d k L N E(\delta^{-1/2}) + c_d 2 E(r) \frac{D(t_{\max}) N}{v_{\max}} + c_s' D(t_{\max}) N,$$

- 随着 L 不同，每个配送时间 l 各个顾客点产生的需求量不同。尽管每辆车都会用完容量 $v_{\max}$，但是所访问的顾客点数不同。这导致 local distance 的不同。
- The formula holds regardless of how many items are included in each shipping period l—even if customer lot sizes are greater than $v_{\max}$.
- It holds in particular if one decides to ship larger quantities than necessary in anticipation of future increases in the demand curves. This has a profound implication for inventory control. Given a number of shipments L to be received by a customer, their sizes and timing can be chosen to minimize holding cost without affecting the transportation cost.

Very cheap items: $c_i \ll c_r$

库存/等待成本远小于租赁成本时的成本分析

- We examine first a case where items are so cheap (c_i is small) that most of the holding cost arises because of the rent paid to hold the items, $c_h \approx c_r$.
- In future lectures, with more expensive items and different customer types, the CA approach will be used to solve this problem. This is not possible now because, since the rent cost is a function of the *maximum* inventory held, said cost cannot be prorated to (small) time intervals based only on the inventories held at those times.
- Fortunately, for a given L the transportation cost is fixed, and the headways only influence the rent cost. Clearly, the headway selection problem is analogous to that examined in the 1-to-1 distribution problem. -> 给定送货次数时，运输成本为固定值，因此仅有租赁成本受到影响。

---

c_r rent cost per item-year; c_i wait cost per item-unit time

Very cheap items: $c_i \ll c_r$ (cont.)

- We saw in the lot size problem with variable demand that holding cost is minimized if **all shipments are just large enough to run out before the next delivery**
- If rent costs were the dominant holding costs (so that the rent cost was proportional to the maximum lot size), then one should choose the dispatching times so as to minimize the maximum lot size →All the lot sizes should be equal, and given by $D(t_{\max})/L$.
- The same occurs here. The minimum holding cost (for L dispatching periods) is thus:

$$\text{Combined holding cost} = N[\frac{D(t_{\max})}{L}]c_r t_{\max}$$

Optimal # of dispatching times

- The total combined logistic cost consists

$$\underbrace{N[\frac{D(t_{\max})}{L}]c_r t_{\max}}_{\text{combined holding cost}} + \underbrace{c_s N\left\{\frac{D(t_{\max})}{v_{\max}} + L\right\} + c_d KLNE(\delta^{-1/2}) + c_d 2E(r)\frac{D(t_{\max})}{v_{\max}} + c_s' D(t_{\max})N}_{\text{combined transportation cost}}$$

- The optimal number of dispatching times L should be chosen by minimizing such a sum. Only the first and second terms of the transportation cost capture the local stop cost and the local distance cost and depend on L. The other terms, corresponding to the line-haul travel and the loading/handling cost do not.
- Thus, the optimal L^* is the solution of an integer constrained EOQ equation that balances the local transportation cost and the rent cost; the solution is close to:

$$L^* \approx \left[\frac{c_r t_{\max} D(t_{\max})}{c_s + c_d kE(\delta^{-1/2})}\right]^{1/2}, \text{ if } L \text{ is greater than 1.}$$

注意书本 107 页公式 (4.9) 有误

Optimal total combined cost/item

The total combined cost per item is approximated by:

$$\frac{c_s + 2c_d E(r)}{v_{\max}} + c_s' + 2[c_r \frac{c_s + c_d k E(\delta^{-1/2})}{\bar{D}'}]^{1/2}$$

where we use $\bar{D}'$ for the average demand rate per customer, $D(t_{\max})/t_{\max}$. Remarkably, the optimal cost does not depend on the shape of $D(t)$. Not many details are needed to provide a reasonable estimate of operating cost.

More expensive items: $c_i \gg c_r$

- We now discuss the problems for items so expensive per unit volume that most of **the holding cost is inventory cost**. Our lectures on lot size problem showed how a CA approach could be used to locate points on the time line (the delivery times) in order to minimize approximately the sum of the holding and motion costs
- The latter was modeled by a constant c_f that represented the added cost of each dispatch. Reasonable for the one-to-one problem examined at the time, this simple formulation also applies now
- From the equation of the combined transportation cost, we notice that with each additional dispatch, the transportation cost still increases by a constant amount (对 L 求导)

$$c_f \approx [c_s + c_d k E(\delta^{-1/2})]N.$$

  This constant represents the local transportation cost induced by the N additional customer visits resulting from the extra dispatch. The line-haul cost remains unchanged.

More expensive items: $c_i \gg c_r$ (cont.)

- Consequently, the results and methods of the lot size problem for the EOQ with variable demand also apply here if one defines $c_f \approx [c_s + c_d k E(\delta^{-1/2})]N$ and replaces $D(t)$ by $ND(t)$. The CA formulation for 1-to-1 problems can then be used to estimate cost. Don't forget to add the (large) fixed components of combined transportation cost that do not depend on L
- Once the dispatch times $\{t_l\}$ and the corresponding delivery lot sizes $\{v_l\}$ have been determined, the vehicle routes can be designed as described in the non-detailed VRP, recognizing that the number of stops per vehicle ($C = n_s^l \approx v_{\max}/v_l$) changes with l.

More expensive items: $c_i \gg c_r$ (cont.)

- For the special case with *uniform density and constant demand*, the cost formula reduces to a form analogous to formula for cheap goods, with c_i, D' and $(|\boldsymbol{R}|/N)^{1/2}$ substituted for c_r, $\bar{D}'$ and $E(\delta^{-1/2})$*.
- This approach has been used to streamline General Motors' finished product distribution procedures. The results have been compared with those of (less efficient) direct shipping strategies†.

*Burns et al. 1985

†Gallego and Simchi-Levy, 1988

Inventory at the origin

- The theory we have described focused on the holding cost at the destination and used cost expressions as if there were an equivalent cost at the origin.
- This assumption is reasonable for the 1-to-1 problems and is now shown also to be reasonable if the one-to-many system is operated as we described.
- However, a modification to the operating procedure can drastically reduce the origin holding costs.

灵活生产策略下的租赁成本

- With our dispatching strategy, where all the destinations are served with each l, the number of items accumulated at the origin reaches a maximum immediately before a dispatch, and at the destinations immediately after a reception.
- If production is flexible, one will produce by dispatch l only those items that must be sent by time t_l (and no more) ; thus, the maximum accumulation at the origin is the size of the largest shipment received by any customer, times N.
- Because shipments arrive as supplies run out, this is also the maximum accumulation for all the customers. It is thus reasonable to represent rent cost by the product of a constant, c_r, and the maximum accumulation, as we have done.

灵活生产策略下的库存成本

- Inventory costs are slightly different. If one could produce the items as fast as desired, one would produce item during a short time interval prior to t_l for each combined shipment l; and would therefore avoid inventory costs at the origin. This is not likely to happen, however.
- Although the production rate can change with time to satisfy a slow varying demand $D(t)$, items are normally produced at a roughly uniform rate during each inter-dispatch interval, since most production processes benefit from a smooth production curve.
- Thus, inventory costs should not be reduced in this manner. If some destinations request more expensive items than others, then inventory cost may be reduced without altering the production rate, simply by changing the order of production. One might want to produce the cheap items at the beginning of the inter-dispatch interval and the most expensive at the end.

灵活生产策略下的库存成本 (cont.)

- In most cases, however, only a fraction of the inventory cost at the origin could be saved by exploiting these differences.
- Thus, the waiting cost at the origin should be comparable to the waiting cost at the destinations, and a strategy which assumes that both holding costs are equal should yield costs close to one which recognizes the inventory cost at the origin more accurately
  - Remember that an error in a cost parameter by a factor of 2 only increases the resulting EOQ cost by about 10%.

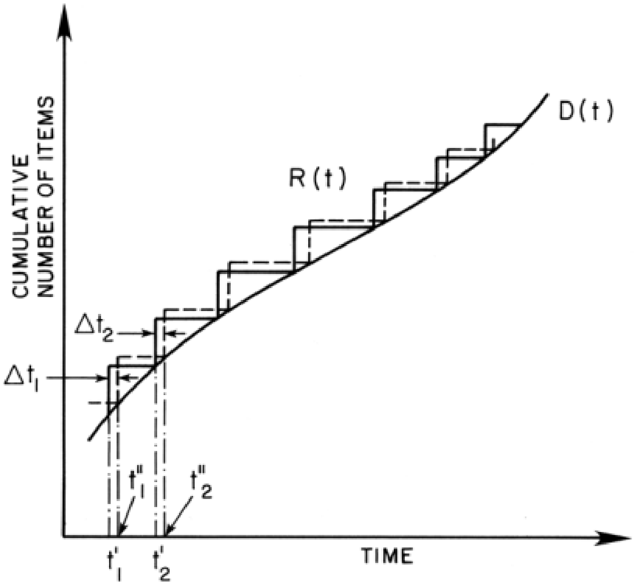
CUMULATIVE NUMBER OF ITEMS
D(t)
R(t)
Δt2
Δt1
t''1
t''2
t'1
t'2
TIME

Staggering production for delivery regions (交错生产策略)

- With our operating strategy, all the points in the region $\boldsymbol{R}$ are visited at each instant l.
- However, if instead of waiting for time t_l, vehicles are dispatched just as soon as their last item is produced, both the storage room and the inventory cost at the origin may be reduced. As shown below, this reduction is largest if one can produce all the items for each one of the delivery districts, in sequence.

交错生产策略

- If the delivery times to any customer are shifted by a time Δt_l smaller than one headway (i.e., the new delivery times are $t_l' = t_l - \Delta t_l > t_{l-1}$), and if Δt_l changes slowly with l so that the new headways are close to the old, then the total holding cost does not change appreciably.
- With a slow varying $D(t)$, the maximum accumulation remains virtually unchanged, and so does the total number of items-hours; see the difference between the solid and dotted $R(t)$ curves.
- This is consistent with the CA solution; the cost is sensitive to the delivery headways used as a function of time but much less so to the specific dispatching times.

交错生产策略

- Suppose that we label the tours used for the l-th shipment: $j = 1, 2, 3$, etc.
- Assume that items for destinations in tour $j = 1$ are produced first, items for destinations in $j = 2$ second, etc; and assume as well that every tour is started as soon as the orders for its customers have been completed.
- If the delivery districts do not change with every l, it would be possible to label them consistently so that all destinations would have the same label in successive dispatches. This would ensure that the l-th delivery headway to every customer is close to $(t_{l+1} - t_l)$, and that as a result the holding cost at all the destinations would remain essentially unchanged.
- The ordered production schedule, though, would cut the maximum and average inventory at the origin by a factor equal to the number of tours used for the l-th shipment, drastically reducing holding costs at the origin.

交错生产策略

Unless the demand is constant, $D(t) = \lambda t +$ constant, it is not reasonable to assume that all the delivery districts remain the same; in that case a less ambitious version of our staggered production schedule can be employed.

- The service region can be partitioned into production subregions $\boldsymbol{P}_1, \boldsymbol{P}_2, \ldots, \boldsymbol{P}_P$, where P is a number small compared with the number of tours in any l, but significantly larger than 1 (so that it can make a difference.)
- Each production subregion should contain the same number of customers (i.e., the same total demand) and require at least several tours to be covered. Under such conditions, the distance for covering $\boldsymbol{R}$ with a VRP is not much different from the collective distance of separate VRP's to cover $\boldsymbol{P}_1, \boldsymbol{P}_2$ etc.
- This is true because, like the TSP, the VRP exhibits a partitioning property. (This should be obvious, since: (i) the cost in each subregion is the sum of the costs prorated to each of its points, and (ii) the cost per point is independent of the partition).

交错生产策略

The following strategy cuts inventories at the origin by a factor P , while preserving virtually unchanged the motion and holding costs at the destination:

1. produce the items for any shipment in order of production subregion: $\boldsymbol{P}_1$ first, then $\boldsymbol{P}_2$, etc
2. On completing production for a subregion, $\boldsymbol{P}_P$, dispatch the vehicles to the subregion on VRP routes constructed for the subregion alone

交错生产策略

- As a practical matter, P does not need to be very large; once it reaches a moderate value (say $P \approx 5$) additional increases yield decreasingly small benefits.
- In fact, even if the demand was perfectly constant, it is unlikely that one would choose a P much larger than 5 because larger P's imply shorter production runs within each $\boldsymbol{P}_p$, which hinders our ability to sequence the production to meet other objectives, such as operating with smoothing worker loads and materials requirements. P 的大小决定了工厂面向不同子区域生产计划的不同，如平滑的员工工作量或者材料要求。P 越大，生产计划变化地越频繁，反而不利于工厂生产。

交叉生产策略

- If production schedules are staggered as described, then the search for the optimal dispatching times should recognize that holding costs will be lower.
- The analysis could be repeated with a changed holding cost equation (e.g., combined holding cost $= N\frac{D(t_{\max})}{L}c_r t_{\max}$ for the case $c_r \gg c_i$) but this is unnecessary; a suitable (downward) adjustment to the holding cost coefficient, either c_i or c_r, has the same effect and also preserves our results.
- If holding costs at the origin can be neglected, the coefficient should be halved; of course, there is no need to pinpoint its value very precisely, since the solution to our problem is robust to errors in the cost coefficients.

目录 |Outline

1 Review

2 Identical Customers and Fixed Vehicle Loads

3 Identical Customers and Vehicle Loads Not Given
- Limits to Route Length
- Accounting for Pipeline Inventory Cost

4 Implementation Considerations

车辆载重未定带来的问题

- In every case discussed so far, the total cost expression decreases with the vehicle load carried $v_{\max}$ $\leftarrow$ the larger $v_{\max}$ the smaller the total number of vehicles that need to be dispatched. $\rightarrow$ In any practical, one would be well advised to use vehicles as large as the (highway, railway ...) network would allow.
- However, the analysis ignored *pipeline inventory cost* and did not consider possible *route length restrictions*. With either one of these complications, it may not always be desirable (or possible) to dispatch full vehicles all the time; vehicle load size becomes a decision variable.

内容构成

- We will discuss route length restrictions first, and will then incorporate pipeline inventory into the models.
- It will be shown that pipeline inventory cost can be ignored for freight that is neither perishable nor extremely valuable, and that it cannot be ignored for passengers.
- Were it not for this complication, the results for fixed vehicle load problems could be used for 1-to-N passenger logistics (e.g., to design a commuter rail network serving a CBD).
- We concludes with a discussion of restrictions on the delivery lot size.

Limits to Route Length

- If the optimization of the identical customer problem results in **very small delivery lot sizes**, each vehicle may have to make an **unreasonably large number of stops**.
- Very long routes may not be feasible if there are restrictions to the duration of a vehicle tour. For example, due to **labor regulations**
- We may explore the consequences of such restrictions

Limits to route length (cont.)

- Tour duration limitations essentially impose a **location-dependent limit** on the number of stops.
- Presumably, locations distant from the depot will need to be served with fewer stops than those which are nearer since more time is needed to reach their general vicinity.
- To recognize this dependence, we use $C_{\max}(\boldsymbol{x})$ for the maximum number of stops around $\boldsymbol{x}$; we assume that $C_{\max}(\boldsymbol{x})$ varies slowly with $\boldsymbol{x}$

Limits to route length (cont.)

- Assume first that N is large, so that most delivery districts do not reach all the way to the depot. Then, to minimize distance one should still attempt to design delivery districts of **width** $[6/\delta(\boldsymbol{x})]^{1/2}$, while making them long enough to include a desired number of stops at (or near) coordinate $\boldsymbol{x}$, $n_s(\boldsymbol{x}) < C_{\max}(\boldsymbol{x})$. This yields: **length** $= n_s(\boldsymbol{x})/[6\delta(\boldsymbol{x})]^{1/2}$. The total distance is then given by expressions

$$\begin{aligned}\text{Total distance} &\approx \sum_i \left[\frac{2r_i}{n_{s,i}} + k\delta^{-1/2}(\boldsymbol{x}_i)\right] \\ &\approx 2NE\left(\frac{r}{n_s}\right) + kNE(\delta^{-1/2})\end{aligned}$$

  where $n_{s,i}$ denotes the number of stops per tour used for tours near $\boldsymbol{x}_i$;
- if $n_s(\boldsymbol{x}) \equiv C$, it coincides precisely with previous expressions. Although the line-haul distance component (the first term) is somewhat different if $n_s(\boldsymbol{x})$ varies with $\boldsymbol{x}$, *the local component remains unchanged*.

Total distance with restrictions on the route length (cont.)

$$\text{Total distance} = 2NE(\frac{r}{n_s}) + kNE(\delta^{-1/2})$$

This expression decreases with $n_{s,i}$ → # stops per tour should be made as large as practicable.
For our problem, # stops used near location $\boldsymbol{x}$ on the l-th dispatch, $n_s^l(\boldsymbol{x})$, should satisfy:

$$n_s^l(\boldsymbol{x}) = \min\{C_{\max}(\boldsymbol{x}); v_{\max}/v_l\},$$

where v_l denotes the delivery lot size used for period l. The expression indicates that the vehicle either reaches its **route length constraint**, or else is **filled to capacity**.
运输批量的两个上界：点 $\boldsymbol{x}$ 处允许的最大值，容量

Dependence on the specific headways?

- With this restriction some of the tours may carry less than a full load. As a result, it may appear that neither the total number of vehicle tours nor the line-haul transportation cost (长途运输的距离是 $2NE(\frac{r}{n_s})$) are fixed.
- We shows that, while not fixed, **the number of tours** (and thus the sum of the line-haul and stop costs) **can sometimes be approximated by an expression that only depends on the number of headways** L; then, the scheduling and routing decisions can still be decomposed.

Approximation for the number of tours

- Assume that $\boldsymbol{R}$ can be partitioned into just a few subregions, $\boldsymbol{P}_p$, with the same limitation on the number of stops: $n_s(\boldsymbol{x}) < C_{\max}(\boldsymbol{x}) \approx C_p$. Characterize each subregion by the number of destinations N_p, and their average distance to the depot $E(r_p)$.
- We will show that the number of tours in each subregion only depends on L. As a result, an expression for the total number of tours is developed.
- The number of tours in period l for subregion p is:

$$\{\#\text{ tours}; l, p\} = \max\left\{\frac{N_p v_l}{v_{\max}}; \frac{N_p}{C_p}\right\},$$

and for all periods:

$$\{\#\text{ tours}; p\} = N_p \sum_{l=1}^{L} \max\left\{\frac{v_l}{v_{\max}}; \frac{1}{C_p}\right\} \geqslant N_p \max\left\{\sum_{l=1}^{L} \frac{v_l}{v_{\max}}; \frac{L}{C_p}\right\}$$

需求接近恒定

- This inequality is a good approximation for the number of tours if rent costs dominate, as then the delivery lot size should be independent of l.
- The approximation will also be good, for the same reason, if the demand is nearly stationary. Then, we can write:

$$\{\#\ \text{tours};p\} \cong N_p \max\left[\frac{D(t_{\max})}{v_{\max}};\frac{L}{C_p}\right] = N_p\left\{\frac{D(t_{\max})}{v_{\max}} + \max\left[0,\frac{L-L_p}{C_p}\right]\right\}$$

where $L_p = C_p D(t_{\max})/v_{\max}$.

- L_p represents a critical number of dispatching periods for subregion p. If $L > L_p$, then the lot sizes are so small that the vehicle cannot be filled in subregion p; the number of stops constraint is binding. L 越大，表示每两次配送时间间隔产生的需求量越小；如果车辆满载，则访问的顾客数会非常多，此时行程数量受最大距离限制。

† L_p 表示车辆在 p 区域所能服务的配送次数临界值。

Total transportation cost

- If the equation is a good approximation for # tours used in P_p, then the sum of **the origin stop cost plus the line-haul cost** for all tours is:

$$\sum_{p=1}^{P}\{\#\text{ tours}; p\}[c_s + 2c_d E(r_p)]$$
$$= \frac{D(t_{\max})}{v_{\max}} N[c_s + 2c_d E(r)] + \sum_{p=1}^{P}[c_s + 2c_d E(r_p)]\left[N_p \max\left(0, \frac{L - L_p}{C_p}\right)\right]$$

  which only depends on the dispatching times through L.

- For small L the expression is constant, and matches the sum of the 1st and 3rd term of the total combined transportation cost. But once L exceeds some of the L_p (some tours hit the length constraint and are only partially filled), it increases with L at an increasing rate.

† Recall the expression for combined transportation cost: $c_s N\{\frac{D(t_{\max})}{v_{\max}} + L\} + c_d kLNE(\delta^{-1/2}) + c_d 2E(r)\frac{D(t_{\max})N}{v_{\max}} + c_s' D(t_{\max})N$。该公式对应于总运输成本的第一项(除去 c_sNL 部分)和第三项

Find the optimal dispatching time L

- The optimal L can be found still as a trade-off between inventory cost* and transportation cost†, with the first and third terms revised.
- Because the revised combined transportation cost equation is piecewise linear and convex, the sum of the inventory cost and transportation cost has only one local/global minima. The revised derivative of total combined transportation cost with respect to L is now a step function:

$$\left[c_s + c_d k E(\delta^{-1/2})\right] N + \sum_{L_p < L} \frac{N_p}{C_p} \left[2 c_d E(r_p) + c_s\right],$$

where the summation only includes p's for which $L_p < L$. The second term represents **the cost increase for the extra tours that need to be sent because (some) vehicles cannot be filled to capacity**. The first term keeps unchanged.

*Recall the expression for combined holding cost: $N[\frac{D(t_{\max})}{L}]c_r t_{\max}$

†Recall the expression for combined transportation cost: $c_s N\{\frac{D(t_{\max})}{v_{\max}} + L\} + c_d k L N E(\delta^{-1/2}) + c_d 2E(r)\frac{D(t_{\max})N}{v_{\max}} + c_s' D(t_{\max})N$

Find the optimal dispatching time L (cont.)

$$\left[c_s + c_d k E(\delta^{-1/2})\right] N + \sum_{L_p < L} \frac{N_p}{C_p} \left[2 c_d E(r_p) + c_s\right],$$

- In the special case where C_p is the same ($C_{\max}$) for all points, there is only one subregion, with $L_1 = C_{\max} D(t_{\max})/v_{\max}$ and $N_1 = N$. Therefore, the second term is zero if $L \leqslant C_{\max} D(t_{\max})/v_{\max}$, and equals $(N/C_{\max})(2c_d E(r) + c_s)$ otherwise.
- The optimal L can be found as follows: If there is a value of L for which the sum of this equation and the derivative of combined holding cost equals zero, then that value is optimal; otherwise, the optimal value is the L_p for which the sum changes sign.
- Because the derivative is larger than before, the optimal L will tend to be smaller and the resulting cost greater. This is intuitive; with limits to route length it may be advisable to increase the lot sizes (by reducing L) to make sure that most of the vehicles travel full.

C_p 随着子区域变化显著的情形

- Our results assume that all customers share the same L and v_l. Although this simplification facilitates production scheduling, it may also **increase logistics costs** when C_p changes significantly across subregions.
- If a different L can be used for different subregions, then **fewer dispatching intervals and larger delivery lot sizes** can be used for subregions with a low C_p; all the vehicles can be filled as a result. A strategy (a set of dispatching times and delivery districts) can then be tailored to each one of the subregions independently of the others. 对于 C_p 不均匀的区域，可为各个子区域设计不同的配送策略；C_p 越小，说明子区域 p 能访问的最大点数越小，可以通过降低配送频率和增大运输批量使得服务该区的车辆被装满。
- We will explore this point — the determination of routing/dispatching strategies that vary in time and space –more thoroughly in the following talks.

Route length restriction for few vehicle tours

- To conclude our discussion on route length restrictions, we must consider the case with few vehicle tours, $N \ll C^2$
- Very simple. The transportation cost is insensitive to # stops per vehicle for this case.
- Hence, route length restrictions do not influence either the optimal dispatching strategy or the final cost.

Accounting for Pipeline Inventory Cost

- In all the optimization problems described so far we have found a solution which minimizes the sum of the motion cost, the holding (rent) cost and the stationary inventory cost. We did not consider the pipeline inventory cost of the items in the vehicles.
- Recall that the pipeline inventory cost/item was $c_i t_m$, where t_m is the average time an item spends inside a vehicle.
- On average an item spends in a vehicle a time approximately equal to one-half of the duration of the tour. If the vehicle travels at a speed s, and takes t_s time units per stop, the duration of a tour with n_s stops and d distance units long is $d/s + (n_s + 1)t_s$; thus:

$$t_m \approx \frac{1}{2}\left[\underbrace{d/s}_{\text{途中时间}} + \underbrace{(n_s+1)t_s}_{\text{停靠时间}}\right], \text{ and } c_i t_m \approx \frac{1}{2}c_i \times \frac{d}{s} + \frac{1}{2}c_i t_s (n_s + 1)$$

Compositions of the pipeline inventory cost

- Added for all items for all L shipping periods, the pipeline inventory cost becomes approximately a simple function of the total number of (item-miles), (items) and (item-stops):

$$\frac{c_i}{s} \times \# \text{ item-miles} + \frac{c_i t_s}{2} \times \# \text{ items} + c_i t_s \times \# \text{ item-stops}$$

- The total number of items is $D(t_{\max})N$. The total number of item-miles and item-stops can be obtained easily if there are no route length restrictions.
- In that case vehicles travel full (from the depot) and every stop delays on average $v_{\max}/2$ items; therefore, the total number of item-stops is $NLv_{\max}/2$. Similarly, each vehicle carries on average $v_{\max}/2$ items and the item-miles equal the product of the vehicle-miles and $(v_{\max}/2)$.

Compositions of the pipeline inventory cost (cont.)

- We have already seen that the total vehicle-miles are (recall the derivation for the total combined distance):

$$\frac{2E(r)D(t_{\max})N}{v_{\max}} + kNLE(\delta^{-1/2})$$

- Thus, the pipeline inventory cost can also be expressed as a function of the decision variables through L alone:

$$\left[\frac{c_i E(r) D(t_{\max}) N}{s}\right] + \left[\frac{c_i KN}{2s} E(\delta^{-1/2}) v_{\max}\right] L + \frac{c_i t_s}{2}[D(t_{\max})N + N v_{\max} L]$$

这里仅仅计算总旅行长度，并代入 'item-mile' 取值。

†注意课本 116 页 'total combined distance' 公式有误。

渠道存货成本 vs 运输成本

$$\frac{c_i E(r) D(t_{\max}) N}{s} + [\frac{c_i K N}{2s} E(\delta^{-1/2}) v_{\max}] L + \frac{c_i t_s}{2} [D(t_{\max}) N + N v_{\max} L]$$

- As a function of L, this expression is similar to the equation for transportation cost*, but it increases much more slowly: at a rate $N[c_i v_{\max}/2][t_s + kE(\delta^{-1/2})/s]$ as opposed to $N[c_s + c_d k E(\delta^{-1/2})]$
- Normally, the quantity $c_i v_{\max} t_s$ represents the cost of delay to the items in a full vehicle during a stop. It should be several orders of magnitude smaller than c_s (the truck cost and driver wages during the stop).
- Likewise, the quantity $c_i v_{\max}/s$ represents the inventory cost of a full truck per unit distance. It should be much smaller than c_d (the vehicle operating cost per unit distance, including driver wages).
- Thus, if pipeline inventory costs had been considered from the beginning, the results would not have changed.

*Recall the expression for combined transportation cost: $c_s N\{\frac{D(t_{\max})}{v_{\max}} + L\} + c_d k L N E(\delta^{-1/2}) + c_d 2E(r)\frac{D(t_{\max})N}{v_{\max}} + c_s' D(t_{\max}) N$

渠道存货成本不能忽略的情形

- If the items are so expensive that the pipeline inventory component cannot be neglected, then **the pipeline inventory cost, unlike the transportation cost, increases with** $v_{\max}$.
- One could thus imagine a situation where a $v_{\max}$ smaller than the maximum possible might be advantageous; the vehicle loads cannot be assumed to be known. **The transportation of people is a case** in point, where the inventory cost of the items carried (the passengers) vastly exceeds the operating cost.
- That is why airport limousine services do not distribute people from an airport to the hotels in the outlying suburbs in large buses; this would result in unacceptably large routes, with some passengers spending too much time in the vehicle*. 机场礼车服务为什么不使用大型车将顾客从机场拉到远郊的各个宾馆？因为这会使行程长的不可接受，部分旅客在车上花的时间太多。

*See Banks, et al. 1982, for a discussion

Re-examine the total transportation cost

- Let us now see how to select the routes and schedules for a system carrying items so valuable that vehicle loads are not necessarily maximal.
- Without an exogenous vehicle load, the total transportation cost no longer can be expressed as a function of L alone; the total vehicle-miles and the number of tours depend on the specific vehicle-loads used, and this has to be recognized in the optimization.
- To cope with this complication, we will consider a set of strategies more general than the ones just examined, but will analyze them less accurately.

Headway

- We will now **allow different parts of $\boldsymbol{R}$ to be served with different delivery headways at the same time**. To do this, we define the smooth and slow varying function $H(t, \boldsymbol{x})$, which represents the headways one would like to use for destinations near $\boldsymbol{x}$ at times close to t.
- Until now we had assumed that the headways were only a function of t: $H(t, \boldsymbol{x}) = H(t)$. As a result, the optimal dispatching times $\{t_l\}$ could be found with the exact numerical techniques; or if $D(t)$ was slow varying, with the CA approach.

stops per tour

- For the present analysis we also seek a function $n_s(t, \boldsymbol{x})$ which indicates the number of stops made by tours near $\boldsymbol{x}$ at a time close to t. Of course, this number cannot be so great that the vehicle capacity is exceeded; the following must be satisfied:
$$\{n_s(t, \boldsymbol{x})D'(t)\}H(t, \boldsymbol{x}) \leqslant v_{\max},$$
- The quantity in braces represents the combined demand rate at the n_s destinations visited by a tour, and the left side of the inequality the load size carried by the vehicle.
- The approach we had used assumed that this equation was a pure equality, so that n_s was only a function of t, $n_s(t) = v_{\max}/[H(t)D'(t)]$, implicitly given by $H(t)$.
- Like $H(t, \boldsymbol{x})$, the function $n_s(t, \boldsymbol{x})$ will be allowed to be continuous and slow-varying during the optimization.

Decision variables

- Once $H(t, \boldsymbol{x})$ and $n_s(t, \boldsymbol{x})$ have been identified, a set of delivery districts and dispatching times consistent with these functions must be found. This will be illustrated after the optimization has been described.
- Let us write the total logistics cost per item that items at time-space point $(t, \boldsymbol{x})$ would have to pay if the parameters of the problem were the same at all other times and locations, i.e., $D'(t) = D', \delta(\boldsymbol{x}) = \delta,$ and $r(\boldsymbol{x}) = r$.
- The decision variables H and n_s that minimize such an objective function will become the sought solution, varying continuously with t and $\boldsymbol{x}$ ($H(t, \boldsymbol{x})$ and $n_s(t, \boldsymbol{x})$).

Total motion cost per item

- The minimum value of the objective function for these coordinates $z(t, \boldsymbol{x})$, is the CA cost estimate.
- Noticing that a vehicle load consists of $D'n_sH$ items and a delivery lot of $D'H$ items, we can express the total motion cost per item as:

$$z_m = \frac{2rc_d}{D'n_sH} + c_dk\delta^{-1/2}\frac{1}{D'H} + c_s\frac{1}{D'H} + c_s\frac{1}{D'n_sH} + c_s'.$$

- Recall the expression for combined transportation (motion) cost: $c_sN[\frac{D(t_{\max})}{v_{\max}} + L] + c_dkLNE(\delta^{-1/2}) + c_d2E(r)\frac{D(t_{\max})N}{v_{\max}} + c_s'D(t_{\max})N$. We may obtain z_m by simply putting $v_{\max} = D'n_sH, D'H = D(t_{\max})/L$ in the expression and dividing it by $D(t_{\max})N$.

Physical interpretation

$$z_m = \frac{2rc_d}{D' n_s H} + c_d k \delta^{-1/2} \frac{1}{D' H} + c_s \frac{1}{D' H} + c_s \frac{1}{D' n_s H} + c_s'.$$

This expression has an intuitive physical interpretation.

- Each tour incurs a cost $(2rc_d + c_s)$ for overcoming the line-haul distance and stopping at the origin, which prorated to all the items in the vehicle yields the first and fourth terms.
- The tour also incurs a cost $(c_d k \delta^{-1/2} + c_s)$ for each local stop and detour, which prorated to the items in a delivery lot, yields the second and third terms of the expression.
- The last term is the (constant) cost of handling each item.

Physical interpretation (cont.)

$$z_m = \frac{2rc_d}{D' n_s H} + c_d k \delta^{-1/2} \frac{1}{D' H} + c_s \frac{1}{D' H} + c_s \frac{1}{D' n_s H} + c_s'.$$

Thus, the first two terms are the cost of overcoming line-haul and local distance (assuming that many tours are needed); the third term is the cost of stopping at the destinations; the fourth the cost of stopping at the origin, and the last one the handling/loading cost.

Holding costs

- The holding costs can be expressed in a similar manner. For the pipeline inventory cost per item, we have the following expression when $C = n_s$:

$$z_p = c_i \frac{r}{s} + c_i k \delta^{-1/2} \frac{n_s}{2s} + c_i \frac{t_s}{2} n_s + \frac{1}{2} c_i t_s.$$

- As with the expression for the total motion cost, the four terms correspond to times spent in **line-haul travel, local travel, destination stops, and at the origin**. The stationary inventory cost per item averages $z_s = c_i H$ if we count it both at the origin and the destination.
- The rent cost can be ignored because if items are expensive compared to transportation costs, they will certainly satisfy $c_i \gg c_r$; thus $c_h = (c_i + c_r) \approx c_i$, and we can write $z_s = c_h H$. .

*Inclusion of rent costs would pose a problem because rent does not depend only on local characteristics such as H and n_s. An exception arises if the demand is stationary in time, $D'(t) = D'$, because then the optimal solution is also stationary; i.e., $H(t, \mathbf{x})$ is independent of t, and the rent cost is $c_r H$

Total logistics cost

If instead of H (and as is often done in the literature) we use the delivery lot size $v = D'H$ as a decision variable, keeping ns as the other variable, then the sum of costs can be expressed as:

$$z = \alpha_0 + \alpha_1 \frac{1}{n_s v} + \alpha_2 \frac{1}{v} + \alpha_3 n_s + \alpha_4 v.$$

where the $\alpha_0, \ldots, \alpha_4$ are the following interpretable cost constants, which will be used from now on:

$\alpha_0 =$ $(c_s' + c_i r/s + c_i t_s/2)$; handling and fixed pipeline inventory cost per item,
$\alpha_1 =$ $(2rc_d + c_s)$; transportation cost per dispatch,
$\alpha_2 =$ $(c_d k\delta^{-1/2} + c_s)$; transportation cost added by a customer detour,
$\alpha_3 =$ $1/2c_i(k\delta^{-1/2}/s + t_s)$; pipeline inventory cost per item caused by a customer detour and the ensuing stop,
$\alpha_4 =$ c_h/D'; stationary holding cost of holding one item during the time $(1/D')$ between demands.

Total logistics cost (cont.)

$$z = \alpha_0 + \alpha_1 \frac{1}{n_s v} + \alpha_2 \frac{1}{v} + \alpha_3 n_s + \alpha_4 v.$$

- z is a "logistics cost function" (LCF) that relates the cost per item distributed to the decision variables of our problem.
- With the new notation, we have: $n_s v \leqslant v_{\max}$. In addition, we require $n_s \geqslant 1$. Clearly, the LCF is constrained by these inequalities.
- We will see that the determination of a realistic LCF is perhaps the most important step in the design of a logistics system with the CA approach. In the present case, the minimum of the total cost subject to these inequalities is the solution to our problem.

Total logistic cost (cont.)

- Note that α_0 can (and often will) be omitted for optimization purposes. Note as well that, with a small modification to the expressions for α_1, α_2, and α_3, This expression also applies to the VRP case with few tours*; k should be replaced by k' and the term $2rc_d$ should be omitted.
- We will assume for the remainder of this section that the α_1, α_2, and α_3 for large N are used in the optimization; if the resulting n_s found in an application is inconsistent with these values, then the α's should be changed to recognize that N is "small". Our qualitative discussion also applies to this case, which is very similar.

*Recall N is small compared with n_s^2

The full vehicle condition

- We now identify a condition under which the pipeline inventory term ($\alpha_3 n_s$) can be neglected, and show that in that case $n_s v = v_{\max}$.
- For any integer n_s, a feasible solution to the LCF is $v = v_{\max}/n_s$, which (ignoring α_0) yields:

$$z(n_s) = \frac{\alpha_1}{v_{\max}} + \frac{\alpha_2 n_s}{v_{\max}} + \alpha_3 n_s + \alpha_4 \frac{v_{\max}}{n_s}$$

- An upper bound, z^u, to the minimum of the LCF, z^*, is obtained from $z(n_s)$, using $n_s \approx \max\{1, v_{\max}(\alpha_4/\alpha_2)^{1/2}\}$ that is:

$$z^* \leqslant z^u \approx \begin{cases} \frac{\alpha_1}{v_{\max}} + 2(\alpha_2\alpha_4)^{1/2} + \alpha_3 v_{\max}(\alpha_4/\alpha_2)^{1/2*}, & \text{if } v_{\max}(\alpha_4/\alpha_2)^{1/2} \geqslant 1 \\ \frac{\alpha_1}{v_{\max}} + \frac{\alpha_2}{v_{\max}} + \alpha_3 + \alpha_4 v_{\max}, & \text{otherwise} \end{cases}$$

*注意课本 121 页笔误

Lower bound

- A lower bound to the optimal cost is obtained by neglecting the pipeline inventory term $\alpha_3 n_s$ of the LCF, and optimizing the problem. We see at a glance that LCF decreases with n_s for any v; thus, one will always choose the largest n_s satisfying constraints: $n_s \leqslant v_{\max}/v$. (Note that if $v < v_{\max}$, then $n_s \geqslant 1$ holds.)
- If this value is substituted for n_s in the LCF, without its first and fourth terms, we obtain a function

$$z(v) = \frac{\alpha_1}{v_{\max}} + \frac{\alpha_2}{v} + \alpha_4 v$$

whose minimum (subject to $v < v_{\max}$) is a lower bound, z^l. Its expression is:

$$z^* \geqslant z^l \approx \begin{cases} \frac{\alpha_1}{v_{\max}} + 2(\alpha_2\alpha_4)^{1/2}, & \text{if } v_{\max}(\alpha_4/\alpha_2)^{1/2*} \geqslant 1 \\ \frac{\alpha_1}{v_{\max}} + \frac{\alpha_2}{v_{\max}} + \alpha_4 v_{\max}, & \text{otherwise} \end{cases}$$

*注意课本 121 页公式 4.21 错误

Gap between z^u and z^l

- Notice that the expressions for z^u and z^l are almost identical: $z^u - z^l = \alpha_3 v_{\max}(\alpha_4/\alpha_2)^{1/2}$ if $v_{\max}(\alpha_4/\alpha_2)^{1/2} \geqslant 1$, and $z^u - z^l = \alpha_3$, otherwise.
- The relative difference between any two of z^u, z^* and z^l should be lower than $\epsilon = \frac{\alpha_3 v_{\max}}{2\alpha_2}$, the ratio of the maximum value of $(z^u - z^l)$ to $2(\alpha_2\alpha_4)^{1/2}$, which is the second term of z^l when $v_{\max}(\alpha_4/\alpha_2)^{1/2} \geqslant 1$. It bounds z^l from below.
- The numerator of this constant, $\alpha_3 v_{\max}$, is the pipeline inventory cost accruing to a full vehicle for one delivery detour; the denominator is double the vehicle motion cost per detour. For most commodities this ratio is orders of magnitude smaller than 1, so that the lower and upper bounds will nearly coincide.
- In summary, if $\epsilon \ll 1$, then filling the vehicles (as done with the strategy leading to z^u) is near optimal; the resulting cost is close to the lower bound, obtained without pipeline inventory costs.

Problems with large gaps

- The incentive to fill vehicles, used so far, does not apply if $\epsilon = \alpha_3 v_{\max}/(2\alpha_2)$ is large compared with 1. The LCF minimization problem then yields a strict inequality for $n_s v \leqslant v_{\max}$. We now examine the solution to this minimization problem with varying conditions in time-space.
- The unconstrained minimum of LCF can be obtained numerically, and it can also be expressed analytically as a function of one single parameter β. To see this, let n_s be close to the unconstrained minimum of LCF: $n_s \approx (\alpha_1/\alpha_3 v)^{1/2}$; then $z^*(v) = 2(\alpha_1\alpha_3/v)^{1/2} + \alpha_2/v + \alpha_4 v$. This expression reflects an achievable cost if $n_s > 1$.

Problems with large gaps (cont.)

- Because $z^*(v)$ is convex, its minimum is the root of $\mathrm{d}z^*(v)/\mathrm{d}v = 0$. Using $v' = (\alpha_1 \alpha_3 v)^{1/2}/\alpha_2$, we can express this equation in terms of v' as follows:

$$\beta \times (v')^4 = 1 + v'; \text{ i.e., } \beta = (v')^{-4} + (v')^{-3}$$

  where $\beta = \alpha_4 \alpha_2^3/(\alpha_1 \alpha_3)^2$
- When v' is small compared with 1 the second term in the last expression can be neglected; in this case the solution is: $v' \approx \beta^{-1/4} \ll 1$ for $\beta \gg 1$.
- Conversely, if v' is large compared with 1, i.e., $\beta \ll 1$, the first term can be neglected and the solution becomes $v' \approx \beta^{-1/3}$.

Problems with large gaps (cont.)

- The largest of the two extreme solutions can be used as a rough approximation when $\beta \approx 1$.
- The optimal vehicle load is $n_s v = v' \alpha_2/\alpha_3$, and $n_s = \alpha_1/\alpha_2 v'$. If the vehicle load is smaller than $v_{\max}$ and $n_s > 1$, then the solution can be accepted. (This happens if $\alpha_2 v' < \alpha_1$ and $\alpha_3 v_{\max}$). The optimal H and z can also be expressed as a function of v', and thus of β.

Ignoring the pipeline inventory cost

- Without pipeline inventory, the solution z^* is as the derivation for the lower bounder. $z^* \approx z^l \approx \begin{cases} \frac{\alpha_1}{v_{\max}} + 2(\alpha_2\alpha_4)^{1/2}, & \text{if } v_{\max}(\alpha_4/\alpha_2)^{1/2} \geqslant 1 \\ \frac{\alpha_1}{v_{\max}} + \frac{\alpha_2}{v_{\max}} + \alpha_4 v_{\max}, & \text{otherwise} \end{cases}$ increases linearly with α_1*
- Because there is an intercept, both z^* and the total cost/ unit time $ND'z^*$ increase "less-than-proportionately" with r; the ratio of cost to distance decreases.
- We also see that z^* decreases with the demand rate/customer D'†, but increases with the spatial density of customers δ‡ if their aggregate demand rate ND' (i.e., $\delta D'$) is constant. However, the total cost/unit time $ND'z^*$ is non-decreasing with D'.
- While not so obvious, these scale economies are also shared by the solution to LCF minimization problem as just described. While $ND'z^*$ increases with D', z^* decreases; the optimal cost also increases less than proportionately with distance from the depot.

*which also increases linearly with the distance from the depot r since $\alpha_1 = 2rc_d + c_s$

† $\alpha_4 = c_h/D'$

‡ $\alpha_2 = c_d k \delta^{-1/2} + c_s$

Extensions

- To estimate cost for a problem with varying $D'(t), \delta(\boldsymbol{x})$ and $r(\boldsymbol{x})$, one would need to **average the analytical solution over t and $\boldsymbol{x}$**. Although it may be possible to do this in closed form using statistical approximation formulas for expectations (these indicate that cost increases with variable conditions), a few numerical calculations should suffice.
- One could calculate z^* for all the $D(t_{\max})N$ items demanded, using their respective t and $\boldsymbol{x}$, but this would be too laborious. Instead, one can partition the time axis into $m = 1, \ldots, M$ intervals and $\boldsymbol{R}$ into $p = 1, \ldots P$ subregions so that each (m, p) combination includes roughly the same amount of demand. We use any interior point $(t, \boldsymbol{x})$ of each combination to calculate both the parameters of the optimization and the resulting cost, z^{mp}. The estimated cost is then the arithmetic average of the z^{mp}.

目录 |Outline

1. Review
2. Identical Customers and Fixed Vehicle Loads
3. Identical Customers and Vehicle Loads Not Given
4. Implementation Considerations
   - Clarens and Hurdle's Case Study
   - Fine-Tuning Possibilities

- We now describe how specific solutions can be designed from the optimization results in prior sections. It also discusses systematic ways for fine-tuning the designs.
- We already know that changes in the input parameters of an EOQ optimization have a dampened effect on the decision variables; this is also true for the objective function now at hand.
- Thus, if $D(t)$ and $\delta(\boldsymbol{x})$ change slowly, the decision variables H (or v) and n_s will change even more sluggishly over t and $\boldsymbol{R}$. Because, as with the EOQ optimization, the decision variables themselves do not need to be set very precisely, it should be possible to identify large regions of the time-space domain where the decision variables can be set constant without a serious penalty.

- For our problem with identical customers, the partition is easily developed: (i) divide the time axis into $m = 1, 2, \ldots, M$ periods with nearly constant demand rates; and (ii) partition $\boldsymbol{R}$ into $p = 1, 2, \ldots P$ subregions with similar customer density and distance to the depot.
- The subregions and time periods should be large enough to include respectively several delivery districts and several headways. This ensures that the number of stops in each district can be close to ideal, and that the theoretical headway $H(t, \boldsymbol{x})$ can be approximated with an integer number of dispatches.
- We anticipate now that, by designing a different spatial partition for every time period, this method can be extended to situations with different customers and time varying customer densities.

- An application of the technique for a very similar problem has been reported by Clarens and Hurdle (1975).
- These authors explored the best way of laying out transit routes from a CBD to its outlying suburbs. They assumed that the demand was stationary and changed with position.

- They describe the solution in terms of slightly different variables and inputs, but the differences are only superficial. They define the vehicle operating cost as a function of time (and not distance), c_t, and do not explicitly account for the number of stops; instead they assume that one knows from empirical observations the time that it takes for a bus to cover one unit area — a constant, $\tau(\boldsymbol{x})$, that can vary with position.
- They define the demand as a density per unit area and unit time, $\lambda(\boldsymbol{x})$, which changes with position. Instead of a distance from the CBD, $r(\boldsymbol{x})$, they define an express (line-haul) travel time, $T(\boldsymbol{x})$, and as a decision variable they use the area of a bus service zone, $A(\boldsymbol{x})$, instead of $n_s(\boldsymbol{x})$. Thus, they work with the following logistic cost function, which is equivalent to LCF:

$$z = \frac{2c_t T}{A\lambda H} + \frac{\tau c_t}{\lambda H} + c_i\{T + \tau A/2\} + c_h H/2$$

where the bus load, $A\lambda H$ is restricted to be below $v_{\max} = 45$ passengers. Note that the constraints are also similar.

Demand distribution for a transit line design problem

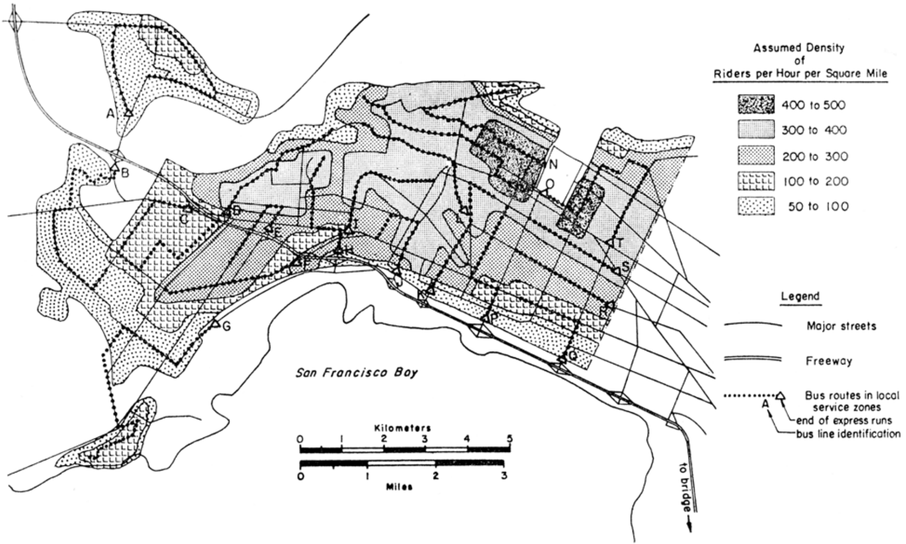

Worksheet for a transit design problem

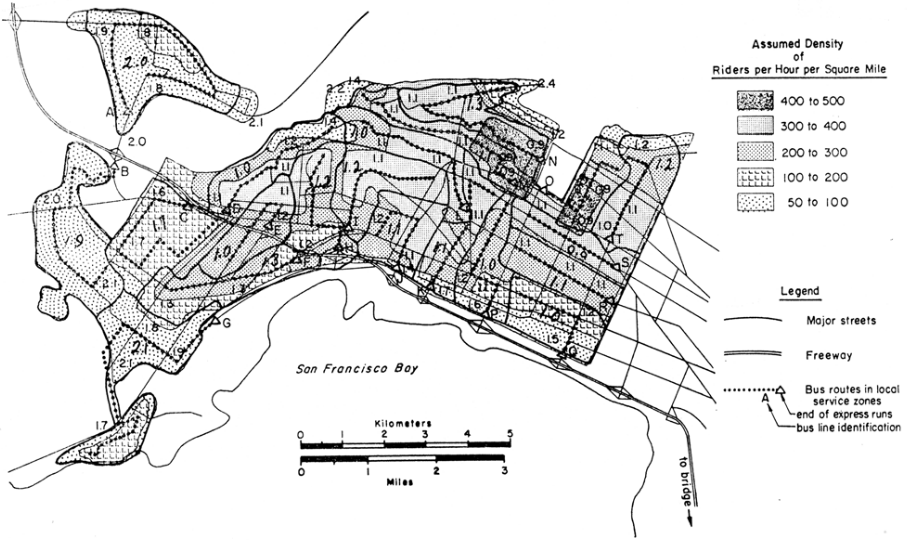

Comparison of the actual and ideal zone size

Results of the transit line design process
(Source: Clarens and Hurdle, 1975)

Zone	Area (square miles)		Headway (minutes)	Average Load On Bus (persons)	Load Factor	$T(x,y)$ (minutes)	$r(x,y)$ (min./sq mi)
	Actual	$A^*(x,y)$					
A	2.0	1.9	13	35	78	27	9
B	1.9	1.9	14	31	69	27	10
C	1.7	1.3	10	43	96	25	10
D	1.0	1.1	9	39	87	26	11
E	1.0	1.2	11	40	89	24	9
F	1.3	1.4	14	36	80	26	10
G	2.1	1.9	8	27	60	21	9
H	1.2	1.5	7.3	38	84	22	8
I [a]	1.2	1.2	7	45	Full	26	8
J	1.1	1.2	7	36	80	20	8
K	1.1	1.1	9.0	38	84	19	10
L	1.0	1.0	6.7	45	Full	24	13
M [a]	1.1	0.9	6.7	45	Full	26	13
N [a]	1.3	0.8	5.8	45	Full	29	16
O [a]	0.9	1.0	8	43	96	25	11
P	1.0	1.0	9	30	67	17	11
Q	1.0	1.3	10	23	51	15	8
R	1.1	1.3	8	35	38	20	8
S	0.9	1.0	7	40	89	21	10
T	1.2	1.5	7	39	87	22	7

[a] Zones where $A^*(x,y) = A_e(x,y)$.

- Given the close agreement between these two columns of figures and the robustness of the CA solution to small departures from the recommended settings, one would expect to have a cost that is very close to the minimum.
- The Clarens-Hurdle case study was an published example where the CA guidelines have been translated into a proposed design for a two-dimensional problem.
- On reviewing the procedure, it becomes clear that a great deal of human intuition is required to complete a design. Furthermore, careful efforts notwithstanding, the designer may miss opportunities for small improvements at the margin that depend on specific details (e.g., stop locations, street intersections, etc.) of the particular problem. It might be worthwhile to use fine-tuning software to find these possible improvements if any exist.

- The rest of this section describes the results of some experiments where fine-tuning software was used to improve detailed VRP solutions developed quickly from the guidelines of TSP/VRP.
- These authors tested simulated annealing (SA) as a technique that is well suited for fine-tuning purposes. The brief discussion of simulated annealing provided in this reference is included as Appendix B. The technique is attractive because:
  - A prototype computer program can be developed quickly for most problems since the SA logic is very simple. (These authors developed software for the VRP, from scratch, in about three mandays.)
  - The optimization can be controlled by means of input variables (called initial "temperature" and "cooling rate" or "annealing schedule") which determine how much the algorithm is allowed to increase (worsen) the objective function at different stages of the process in the hope of finding larger reductions later.

- Simulated annealing is known to converge in probability to the global optimum of combinatorial optimization problems, such as those arising when designing in detail logistics systems.
- Unfortunately, convergence is slow. To be guaranteed, the initial temperature has to be very large and the cooling rate very slow; the computer time required rapidly becomes prohibitively long with increasing problem size. However, with an overall idea of the system's structure, and a near optimal initial solution as would be obtained with nondetailed methods, the scope of the annealing search can be restricted. As demonstrated in Robusté et al. (1990), a low initial temperature achieves that.
- It prevents the search from wandering away from the initial solution, while systematically testing variations that exploit the details (specific locations of customers, for example.)

- One of the examples in this reference considers a VRP problem with $N = 500$ points (randomly located according to a uniform density in a 6-inch by 10-inch rectangle), $C = 45$ stops per tour and a centrally located depot; distances are Euclidean.
- For this test the VRP formula with $k \approx 0.57$, predicts a total distance averaging 179 inches. With a high initial temperature, the SA approach yielded tours that were very long in reasonable times; after one day of computation it obtained a set of tours 180.4 inches long. This was reasonable, but longer than the hand constructed tours using the VRP guidelines presented earlier.
- When the hand constructed tours were used to initiate SA with a low initial temperature, the SA algorithm found enough modifications to reduce the total length by about four percent — to 173.6 inches.

SA solution

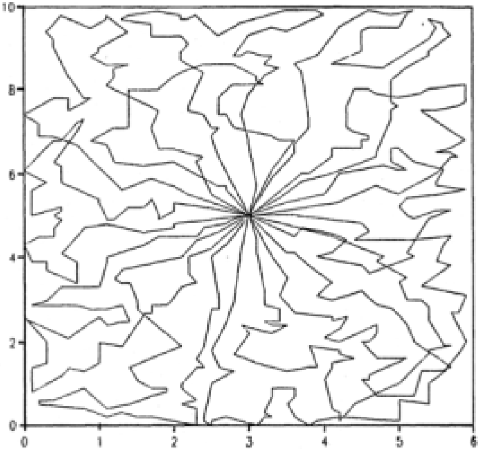

Figure: 500 point VRP. C= 45. 12 tours with total length = 180.4inches.

Manual solution

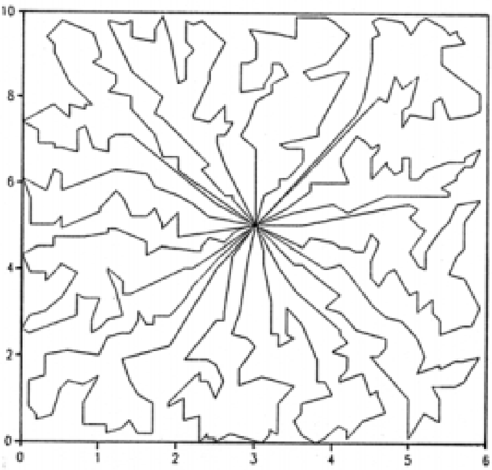

Figure: 500 Point VRP. C=45. 12 tours with total length = 179.8 inches.

Other tests performed in this reference show that the non-detailed approach, fine-tuned with SA, can obtain solutions with objective functions as low as those currently believed to be optimal. The efficiency of the twostep approach has also been demonstrated in practice — the (non-detailed) results in Burns and Daganzo (1987) were used in conjunction with SA to schedule the assembly lines in some GM plants

- These observations are in agreement with our philosophical conclusions. Like the evolution processes in nature, to design a complex logistic system it seems best to develop a preliminary design based on the overall characteristics of the problem, and use the details later to fine-tune the preliminary design. This view has been adopted in the recent works of Langevin and StMleux (1992) and Hall et. al. (1994).
- Although the CA approach and the SA algorithm seem to be ideal companions for this twostep approach, other methods may also be useful. The critical thing is not the specific approach for each step, but the fact that the first step disregards details in searching over all possible solutions, and the second step—restricted to a small subset of possible solutions—incorporates all the details.
- Perhaps other computer fine-tuning methods will improve on SA (Neural Networks and Tabu Searches...etc.). But the improvement should not be measured only on computation grounds; the ability to develop the software quickly is just as important.

Any questions?

Readings

- Daganzo. Logistics System Analysis. Ch.4. Page 105-132.