物流系统分析
Logistics System Analysis

第 16 周　多到多配送问题（3）
Many-to-Many Distribution with Transshipments — Multi-Terminal Systems, Multiple Transshipments

葛乾

西南交通大学　系统所 & 交通工程系

目录 |Outline

- This lecture shows how further cost reductions can be achieved with additional transshipments. We first discuss two transshipments through break-bulk terminals (BBTs); as in the prior sections, it will be assumed that vehicles of maximum size, v_{max}, can reach the origins and destinations. It will also be assumed that the pipeline inventory cost can be neglected, relative to transportation costs. Systems with both BBTs and consolidation terminals (CTs) are examined later

目 录 |Outline

- With two transshipments, a non-hierarchical arrangement of terminals is no longer asymmetric with respect to collection and distribution; as we shall see, it requires few local stops at both ends of a trip. Thus, a hierarchy of terminal levels is not used to select an item's route through the terminals: each O-D pair is simply assigned to the least circuitous terminal pair, considering only terminals in the immediate neighborhood of the origin or the destination.

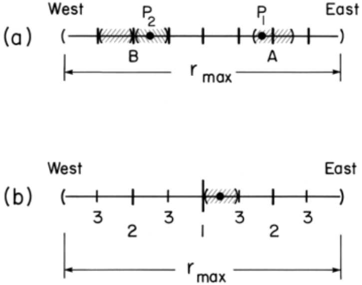

- In the figure, only the terminals on the four corners of the cell containing the origin (or the destination) would be potential candidates. Of the 16 possible combinations the pair adding the least distance should be chosen.
- A typical item would first travel to the origin terminal on a collection vehicle; it would then be sent to the destination terminal on an inter-terminal vehicle and finally, after a second transshipment, it would be delivered to its destination

We will assume here that:

1. all the vehicles arrive and leave the terminals full,
2. the system is operated on a clock with a common headway H,
3. every terminal pair is linked by a non-stop vehicle route along the shortest path; i.e., multiple stops at the terminals are not allowed.

Conditions (1) and (2), used with one transshipment systems, should be desirable here for the same reasons. Condition (3) ensures that the vehicle-miles of inter-terminal travel are minimum.

A more general set of conditions, e.g. with different headways for collection and distribution than for local travel, does not reduce cost (Daganzo, 1987c).

- With the system operated on a clock the holding cost per item is not increased by the second transshipment; it is still given by $c_h H$. The motion cost and terminal cost expressions are examined next.
- If we decompose collection and delivery vehicle-miles into line-haul and local components as before, and define inter-terminal travel as linehaul, then the total line-haul vehicle-miles still equal the total item-miles divided by $v_{\max}$.
- Consideration shows that for our routing scheme the distance added by the two transshipments is the same as the distance added with only one, $2|\boldsymbol{R}|^{1/2}/(3N_T)$. As a result, the line-haul circuity cost per item is still given by $\frac{c_d}{v_{\max}} \times \frac{2|\boldsymbol{R}|^{1/2}}{3N_T}$

- Local motion costs are still proportional to the number of stops, with the same proportionality factor. As before, the number of stops per customer equals the number of terminals serving the customer in both directions. But this number, given by $m^o + m^d = 3N_T^{1/2}$ for one transhipment, is now reduced to $4 + 4 = 8$. It is even less for customers along the boundary of the region.
- Consequently, the stop cost per item, assuming that $\delta^o \approx \delta^d$ as in $\frac{3\alpha_2\delta}{\lambda H|\boldsymbol{R}|}(N_T^{1/2})$ is instead about $\frac{3}{8}(N_T)^{1/2}$ times smaller: $8\alpha_2\delta/(\lambda H|\boldsymbol{R}|)$.
- For a given N_T and total demand, the fixed terminal costs don't change, but the terminal costs proportional to flow should just about double; after all, items are handled twice and spend twice the amount of time moving through terminals. Thus, the terminal cost per item should become: $2\alpha_5 + \alpha_6 N_T/|\boldsymbol{R}|$.

- The solutions of the strategic and tactical problems are now analogous to the previous lectures.

$$\text{stop cost/item} = \frac{3\alpha_2\delta}{\lambda H|\boldsymbol{R}|}N_T^{1/2}$$

$$\text{circuity cost/item} = \frac{c_d}{v_{\max}} \times \frac{2|\boldsymbol{R}|^{1/2}}{3N_T}$$

$$\text{terminal cost/item} = \alpha_5 + \alpha_6\frac{N_T}{|\boldsymbol{R}|}$$

$$\text{holding cost/item} = c_h H$$

- Now, however, we must introduce a flow conservation constraint. Because the system operates on a common schedule, vehicles are full, and multiple stops at terminals are not allowed, the number of collection tours arriving at a terminal (i.e. vehicle loads) must equal N_T.

- To see this note that the number of vehicle loads collected at a terminal for other terminals must equal the number departing for other terminals, and this number is $N_T - 1$. Because $1/N_T$ of the freight collected is local, the number of vehicle loads collected must be N_T. The same occurs for distribution, but with $\delta^o = \delta^d$ and spatially homogeneous demand the condition is redundant. Thus:

$$\frac{\lambda |\boldsymbol{R}|^2 H}{N_T v_{\max}} = N_T.$$

- If terminal costs are neglected, and we use $k = 0.5$ in the expression for α_2, the optimal number of terminals and number of stops can still be expressed as a function of N_o and K alone. For $N_o > 10^2 K$ the following closed form expressions are obtained:

$$N_T^* \approx (8N_o)^{1/4} \left(\frac{N_o}{K} \right)^{1/2}$$

$$\frac{n_s^*}{\sqrt{N_o}} \approx \frac{2^{1/2} K}{N_o}$$

- For smaller N_o, the results are given in Table II of Daganzo (1987c). The results in this reference use a slightly lower estimate for circuity cost than $\frac{c_d}{v_{\max}} \times \frac{2|\boldsymbol{R}|^{1/2}}{3N_T}$ but this has no noticeable numerical impact on the final result.

- The total cost per item also assumes a form similar to Eq. (6.10b). But now (for all N_o) the factor in braces is even smaller; it is $(2.8K/N_o)$. The difference between the two factors reaches a maximum, about 0.15, for N_o/K on the order of 10^1.
- But the actual difference is smaller because in deriving the one transshipment results we used two conservative simplifications, $2|\boldsymbol{R}|^{1/2}/(3N_T)$ and $m^o + m^d = 3N_T^{1/2}$, which lead to slightly higher cost estimates for low N_o/K. For N_o/K comparable with 101, the cost overestimation is on the order of 0.08.
- Thus, the maximum difference between 1-transhipment and 2-transshipment costs should be on the order of 7% (and not 15%) of the cost of driving an item across the service region in a full vehicle.

- Let the cost of transhipping a vehicle load including fixed delays and handling cost, $\alpha_5 v_{\max}$, be momentarily defined in terms of the cost of driving a vehicle a critical distance, $c_d r_{crit} = \alpha_5 v_{\max}$.
- Then, adding one transshipment to every item would have the equivalent effect of adding r_{crit} miles to the distance traveled by each item in a full vehicle. Clearly, two transshipments should not be considered if $r_{crit} > 7\%$ of the diameter of the service region. For $r_{crit} \approx 10^2$ miles (a value typical of trucking operations) only service regions as large as the largest countries in the world have the potential for benefiting from two BBT transshipments

目 录 |Outline

- The above statements do not imply that items (e.g. a letter) should not be handled more than twice between an origin and a destination; we are only stating that there is no need to have them pass through more than 1 or 2 *break-bulk terminals* — terminals serving multiple origins and multiple destinations where vehicles of similar characteristics swap their loads
- Consolidation terminals (connected with either a single origin or a single destination) should also be used to achieve the two main functions described in 1-to-N problems w/ transshipments: reducing the length of delivery and collection routes, and allowing small vehicles to reach the origins and destinations.

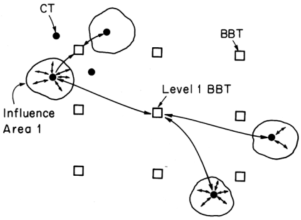

- A rationally designed many-to-many system might be organized as shown in the figure, using both consolidation (CT) and break-bulk (BBT) terminals. Each consolidation terminal would collect (and distribute) items from origins (and destinations) in an influence area around it; influence areas would form a partition of the service region to ensure that service is provided everywhere.

- Conceivably one could have smaller CT's within each of these influence areas, but this is unlikely. The upper level CT's, shown by dots in the figure, would then become the entry points in the many-to-many network of BBT's, shown by squares in the figure.
- The figure denotes by arrows the paths that items either originating or ending in influence area 1 would take on the network; a single BBT transshipment is assumed. Conceivably, the BBT's themselves could also be gates to the system, acting like upper level CT's with their own influence areas

- A conditional decomposition approach, combining the result of 1-to-N and N-to-N problems, can be used to develop desirable structures for an integrated logistic system such as the one in the figure. Conditional on the size, ICT, of the influence areas of the consolidation terminals, i.e. on the number of gates (e.g. post offices) to the break-bulk network, $N_o = N_d = |\boldsymbol{R}|/I_{CT}$, it is possible to determine the near-minimum cost per item on both portions of the system.
- On the consolidation portion of the system within the influence areas, one ould use the methods of 1-to-N w/ transshipments, and on the break-bulk portion those of the previous lecture. In addition to N_o, we may want to freeze N_T and the headway for the BBT network, H. In this way we can conveniently explore the economic merits of synchronizing the operations on both networks, and can invoke the results of previous lectures, depending on whether $N_T = 0, N_T = 1$, or $N_T \approx 1$. The values of N_o, N_T and H that minimize the sum of both costs should then be chosen. A more detailed design can then be developed as we have already learned, perhaps using fine tuning tools with detailed data. Problem 6.7 illustrates the approach.

- A recent application of these ideas, with some further development, is documented in Smilowitz (2001) and Smilowitz and Daganzo (2004), which describe an effort to design and evaluate large scale, integrated package delivery systems such as those of UPS and FedEx. These references examine the conditions under which it makes sense to integrate an air-express network into an existing ground network. The CA techniques proved to be practical and accurate cost predictors.
- They revealed that the larger a ground network, the more efficiently it can absorb a given air network. This helps explain why UPS has chosen to run an integrated air/ground network, but not FedEx. Of course, other factors can also contribute to such decisions. Labor issues are perhaps the most obvious, since existing contracts would likely have to be renegotiated after a structural change in operations.

目 录 |Outline

Variable demand

- We have not discussed in detail in this lecture how one should handle non-homogeneous origin and destination tables. It was assumed for the most part that origins and destinations were homogeneously distributed and that the flows between regions of comparable size was relatively independent of position. In practice, though, this is not likely to happen since population densities typically change over space.

- We have seen that if some customers are much larger than others it may be better to serve them without transshipments, but we did not explore how to deal with spatial variations in demand density and customer density (we treated λ, δ^o and δ^d as constants).
- In prior chapters we had used the continuous approximation method to deal with such variations, and this is also possible here. While it might now appear that we would have to specify an origin-destination flow table in detail, $\lambda(\boldsymbol{x}^o, \boldsymbol{x}^d)$, the solution is mostly sensitive to the generation and attraction rate densities: $\lambda \boldsymbol{x}^o$ and $\lambda \boldsymbol{x}^d$.

- With variable demand we would still try to locate the BBT terminals on a lattice, but would want to vary the number of consolidation terminals and their operation according to location. One may also wish to locate more BBT's in high density areas, but we will ignore this for the moment.
- The optimal solution, thus, is defined in terms of the H, N_T, $H_{CT}(\boldsymbol{x})$, and $\delta_{CT}(\boldsymbol{x})$, where $\delta_{CT}(\boldsymbol{x})$ is the spatial density of CT's in the vicinity of $\boldsymbol{x}$, and $H_{CT}(\boldsymbol{x})$ is the headway used at those CT's.
- We now show that, holding H and NT constant, the total cost decomposes locally in a manner that allows H_{CT} and δ_{CT} to be defined with the CA approach.

- The motion cost during consolidation is independent of H and N_T . For a given terminal, it only depends on $\lambda^o, \lambda^d, \delta_{CT}$ and H_{CT}, and can be prorated to small sub-regions of $\boldsymbol{R}$ as a function of $\boldsymbol{x}$ alone. If no transshipments take place in the consolidation area, the average motion cost per item for collection is given by the function $z_m^o(\lambda^o, \lambda^o, \delta_{CT}^{-1}, H_{CT})$, defined in connection with

$$z_m^i(\lambda, r, I, H^i) + z_m^o(\lambda, \delta, I, H^o) + (c_r + c_i)\max(H^o; H^i) + (\alpha_5 + \alpha_5 I^{-1})$$

.

- As stated, the arguments of z_m^o only depend on $\boldsymbol{x}$. Similarly, the motion distribution cost is: $z_m^o(\lambda^o, \lambda^d, \delta_{CT}^{-1}, H_{CT})$.

- The two components of BBT motion costs are also well behaved. Given N_T and H, the circuity (line-haul) cost per unit time over $\boldsymbol{R}$ is insensitive to $\delta_{CT}(\boldsymbol{x})$ and $H_{CT}(\boldsymbol{x})$. We note that flow will pulse through BBT's differently if these variable change, but the total item-miles should remain fixed.
- Therefore, this cost can be ignored for the minimization of the consolidation terminal variables. Important for the minimization of N_T and H, we also note that the average circuity cost per item should be rather insensitive to $\lambda(\boldsymbol{x}^o, \boldsymbol{x}^d)$ if O-D trips are comparable with the diameter of $\boldsymbol{R}$. $\frac{c_d}{v_{\max}} \times \left(\frac{2|\boldsymbol{R}|^{1/2}}{3N_T}\right)$ should then be a good approximation. The BBT (local) stop cost per unit time and unit area arising from visits to the CT's can be obtained simply. $m^o + m^d = 3N_T^{1/2}$ gives the number of stops at each CT every headway

- Since there are δ_{CT} CT's per unit area and each stop costs $\alpha_2' = c_s + c_d k \delta_{CT}^{-1/2}$, the BBT system=s local stop cost per unit time and unit area is:

$$z_s'(\delta_{CT}, N_T, H) \approx 3\frac{N^{1/2}}{H}(\alpha_2' \delta_{CT})$$

- We have argued in the 1-to-N problems with transshipment that the terminals' cost per unit time should be of the form α_5(# items/unit time) + α_6(# terminals). This total cost for the CT's can be prorated locally to small areas as a cost per unit time and unit area

$$z_T'(\lambda^o, \lambda^d, \delta_{CT}) = (\lambda^o + \lambda^d)\alpha_5 + \alpha_6 \delta_{CT}.$$

This expression yields the total cost when integrated over $\boldsymbol{R}$. It is again independent of the origin-destination flow details,$\lambda(\boldsymbol{x}^o, \boldsymbol{x}^d)$.

- Finally, we must account for holding cost. If the schedules of the breakbulk and consolidation vehicles are not synchronized, then the waiting cost for an item traveling from $\boldsymbol{x}^o$ to $\boldsymbol{x}^d$ is $c_h(H_{CT}(\boldsymbol{x}^o) + H + H_{CT}(\boldsymbol{x}^d))$. The total holding cost for all items can be prorated to a unit area per unit time so that it only depends on the location:

$$z'_h(\lambda^o, \lambda^d, H, H_{CT}) = (\lambda^o + \lambda^d)c_h[H_{CT}(\boldsymbol{x}^o) + H/2].$$

- If the two systems are synchronized and $H = H_{CT}$ for all locations, then the total holding cost can also be prorated locally:

$$z'_h(\lambda^o, \lambda^d, H) = (\lambda^o + \lambda^d)c_h H/2.$$

- In either case, z'_h is independent of $\lambda(\boldsymbol{x}^o, \boldsymbol{x}^d)$.

- Since all the cost components can be prorated to small sections of $\boldsymbol{R}$ as a function of $\boldsymbol{x}$ alone, given N_T and H, it is possible to obtain the best H_{CT} and δ_{CT} for any location, $\boldsymbol{x}$, by minimizing the sum of z'_h, z'_T, z'_s , the collection motion cost $\lambda^o z^o_m(\lambda^o, \delta^o, \delta^{-1}_{CT}, H_{CT})$, and the delivery motion cost $\lambda^d z^d_m(\lambda^d, \delta^d, \delta^{-1}_{CT}, H_{CT})$. The minimum of this sum is a function only of $\boldsymbol{x}, N_T$ and $H : z' * (\boldsymbol{x}, N_T, H)$. Integrated over $\boldsymbol{R}$, it yields the (approximate) total system cost per unit time for a given N_T and H, exclusive of the BBT circuity costs and BBT terminal costs.
- Reasonable values for N_T and H can now be found easily numerically by minimizing the sum of the integral, the BBT circuity costs $\frac{Dc_d}{v_{\max}} \times \frac{2|\boldsymbol{R}|^{1/2}}{3N_T}$ and the BBT terminal cost $2D\alpha_5 + \alpha_6 N_T$, where D is a constant, representing the total number of items traveling per unit time:

$$D = \int_{\boldsymbol{R}} \int_{\boldsymbol{R}} \lambda(\boldsymbol{x}^o, \boldsymbol{x}^d) \mathrm{d}\boldsymbol{x}^o \mathrm{d}\boldsymbol{x}^d.$$

- We have assumed that the N_T terminals would be homogeneously distributed over $\boldsymbol{R}$, but in practice one would try to locate them at the intersections of major flow corridors if these have been identified in order to reduce circuity costs. The location of the BBT's, however, does not affect any of the costs used in the above calculations, except perhaps for the linehaul circuity cost. By providing more BBTs in sections of $\boldsymbol{R}$, with heavy demand and higher concentrations of CT's it may be possible to reduce the extra line-haul distance traveled by an average item below $2|\boldsymbol{R}|^{1/2}/(3N_T)$.
- Any such adjustment, however, should change the distance considerably less than a small percentage increase in N_T. Thus, even if the circuity distance with the adjustment could be quantified by a more detailed expression, the resulting optimization would likely yield a similar value for N_T and H. Given what we know about the robustness of solutions, the solution we did obtain should result in costs that are not far from the ideal

- If desired, and once the locations of the CT's have been chosen, one may be able to formulate and solve approximately a detailed optimization program similar to the detailed solution, with the location of the N_T break-bulk terminals as decision variables in addition to the flow allocation variables, x_k^{ij}. We are optimistic about such endeavor when the desired number of BBT's is not large as would occur when the number of gates to the BBT network is itself moderate. If the number of BBT's is large, then circuity costs are small and minor reductions to it are of secondary importance.

Backhauls

- It was assumed throughout this chapter that vehicles could be usefully employed at the end of their trips. This effect was captured by adjusting the motion cost coefficient α_1.
- Although not exclusively, this assumption is reasonable if the origin-destination flows are balanced (i.e. $\lambda(\boldsymbol{x}^o, \boldsymbol{x}^d) \approx \lambda(\boldsymbol{x}^d, \boldsymbol{x}^o)$) ; then when a local vehicle finishes the last delivery it is automatically well positioned to start a collection run with little deadheading.
- Furthermore, with a symmetric O-D flow pattern the inbound and outbound flows at every BBT are equal, which obviates the need for interterminal empty vehicle travel

- If the demand is unbalanced, a more accurate accounting of vehiclemiles is necessary since partially empty vehicles will either be arriving or departing from the terminals. The problem is likely to be more severe for BBT's than for CT's, since vehicles travel longer distances between CT's and BBT's than between CT's and individual customers.

- Models and formulas have been developed to estimate CT vehicle mileage when vehicles and crews are based at an individual terminal and vehicles backhaul between the last delivery and the first pick-up (Daganzo and Hall, 1990).
- Such formulas would also apply to BBT routing with one transshipment.) An extensive algorithmic literature on the VRP with backhauls also exists (Casco et al., 1988). When the imbalance between inbound and outbound freight is significant, Daganzo and Hall (1990) shows that the distance traveled is just barely greater than the distance that would have to be traveled to collect (or distribute) the dominant direction of flow only, as if the other direction did not exist.
- We have already realized in the symmetric strategies that this would require doubling α_1 for the dominant direction and setting it equal to zero for the secondary. With the proper distance formula, it should not be difficult to duplicate the analysis

- If vehicles can visit a number of terminals, and some neighboring terminals have opposite imbalances, empty miles might be reduced further by backhauling from the last delivery of one terminal to a pickup of the neighboring one and balancing that trip by sending an empty vehicle from the second terminal to the first. The advantage of multi-terminal backhauling is particularly clear for inter-terminal vehicles in a two transshipment system. Now too, imbalances between pairs of BBT's result in some BBT's having an excess of inbound flow and others an excess of outbound flow.
- The solution to the Hitchcock problem of linear programming can be used to route empty BBT vehicles among terminals to minimize empty miles. The solution to the Hitchcock problem, however, may require some crews to visit several BBT's before returning to their home base; an outcome which is not desirable in practice. Other real life constraints also complicate the decision.

- Since carriers can greatly benefit from fewer empty miles, substantial research efforts have been made to improve backhauling decisions (see for example: Jordan, 1982, Powell et. al, 1984, and Dejax and Crainic, 1987). Most of these works, however, are algorithmic in nature, dealing with peculiarities such as real time control with imperfect information, and don't yield simple distance formulae as a function of few descriptors. This is indeed difficult for this problem. Jordan and Burns (1984) and Hall (1990) have sought estimates of empty miles for small networks, using strategies where each vehicle visits at most 2 BBT's before returning home; see also problem 6.8. This is pessimistic,

- We would like to estimate empty vehicle miles as a simple function of NT , which could then be used with Eqs. (6.8) to explore the various tradeoffs. A somewhat optimistic estimate of this quantity, good for large N_T, is given in Daganzo and Smilowitz (2004). This reference proved with a combination of dimensional arguments and mathematical analysis that the expected distance required to reposition an empty truck in a large homogeneous system operated with the Hitchcock recipe must be insensitive to the shape of the service region, and is given by

$$\delta_{BBT}^{1/2}(a + b\log_2 N_T)$$

where δ_{BBT} is the spatial density of BBT's, and a and b are dimensionless constants that depend on the metric.

- Simulations show that $a \approx 1$ and $b \approx 0.078$ in the Euclidean case. In practical terms, this means that the average distance traveled by an empty truck is roughly comparable with the separation between nearby terminals; i.e., that it is about twice as long as for the Euclidean TSP, for $N_T \approx 25 \rightarrow 210$. By multiplying the expected distance formula with the expected number of back hauls (easy to estimate if an underlying stochastic model is given), we can estimate the total expected empty vehicle miles. The result is optimistic because it is based on the Hitchcock problem, which slightly underestimates the distances of the real

- Fortunately, formulas for empty back hauls do not have to be very accurate, because in most cases empty miles should be a small proportion of the total, and do not depend heavily on NT.
- After all, if flow imbalances are serious, a carrier will normally take marketing steps to correct the imbalances since every extra item in the non-dominant direction can be carried without extra vehicle-miles. This can be done by pricing directions differently, or by other means. For example, rental car companies have drive-away programs to reposition their fleet and, because UPS's parcel flows tend to be heavier in the westward direction, that firm has considered using their trucks for carrying California produce toward the eastern U.S.

Large scale manufacturing systems

- The methods and ideas we have described can be extended to the organization of manufacturing systems (e.g., to the planning and design of supply chains). We saw in the 1-to-N problems w/ transshipments how the factory location problem was a special case of the terminal location problem in a one-to-many distribution network.
- This view was premised on the assumption that the inputs to the manufacturing process were ubiquitous; i.e. that changing the locations of a factory did not change the inbound logistics costs, which then could be ignored to define the system

- If some of the inputs are not ubiquitous, and must be obtained from fixed sources regardless of location, then the one-to-many model does not hold. But, we can view the *production* process as a many-to-many logistic process that conveys these raw materials from their sources to destination markets, in the form of a final product, passing through factories and terminals on their way.

- Factories can be viewed as special kinds of terminals which somehow change the nature of the items entering and leaving. In this lecture terminals satisfied a flow conservation equation ensuring that the number of items (e.g. tons) entering a terminal were in the long run equal to the number leaving. But this weight (or volume) conservation does not apply to factories. Burns (1986) likes to distinguish between factories that transform raw materials into parts, reducing weight and volume (production plants), from those that combine parts into bulkier final products (assembly plants).

- Transshipment points in a manufacturing network must be treated differently depending on whether they are bulk reducing, bulk conserving, or bulk increasing. Depending on the industry, each item may be produced and assembled at an integrated factory in a single location, or they may not.
- In an integrated system several factories may manufacture the items, but every item passes through only one factory. Suitably modified, the models for one transshipment* will apply. They would have to capture the different transportation needs of the inbound and outbound items. This has been preliminarily explored for the one-factory problem when all the vehicles arrive and leave the factory fully loaded in Bhaskharan and Daganzo (1987). This report shows that most of the logistics costs are independent of location; and that only the cost of overcoming distance (with full vehicles) depends on it.

*in a one-terminal system or multi-terminal system

As a result, the best location is the solution to a Weber problem (already discussed), *where origins and destinations have weights which reflect the ease of transport and the value of their items*. If inbound volume greatly exceeds outbound volume then there is an incentive to locate the factory close to the raw material sources; if the factory adds much value, then there is an incentive to locate the factory near the markets

- If there are clearly defined zones for raw materials and markets, and these are far apart, then there is an incentive to "dis-integrate" the system. Production plants could be located close to the raw materials and assembly plants close to the markets.
- In this manner transportation costs can be greatly reduced since raw materials make their way to the market in their most easily transportable form: parts, of which waste materials have been scrapped and burned away at the source but which have not yet been assembled into awkwardly shaped final products.

- For a large firm, able to operate multiple factories, it may be best to operate specialized parts plants and assembly plants. Parts with similar raw material needs would be produced in the same plant, located optimally with respect to the raw materials and recognizing the different cost of production (e.g. labor productivity and wages) at different locations.
- Parts would then be assembled into final products at assembly plants close to the markets, allowing production to take place where it is most efficient without incurring very large transportation costs. The practice is very prevalent in the automobile manufacturing industry, where parts are often shipped half way around the world for assembly in another country.

- It seems worthwhile to extend the non-detailed methods espoused in this book to aid in a more through understanding of large scale manufacturing systems in dynamic environments. The techniques seem ideally suited to that end.
- They have recently been used to unveil near-optimal designs and operating rules for some simple supply network scenarios, and to quantify the difference in performance between optimally-designed centralized and decentralized networks (Daganzo, 2002, 2004). This work, however, only begins to scratch the surface of possibilities.

Any questions?

Readings

- Daganzo. Logistics System Analysis. Ch.6. Page 249-266.