Column Generation II : Application in Distribution Network Design

Teo Chung-Piaw (NUS)

27 Feb 2003, Singapore

1 Supply Chain Challenges

1.1 Introduction

Slide 1

Network of facilities: *procurement* of materials, *transformation* of these materials into intermediate and finished products, and the *distribution* of these finished products to customers.

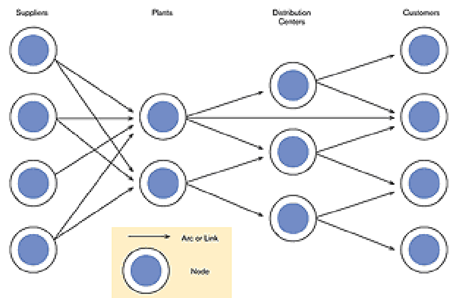

1.1.1 Distribution Network

Slide 2

Distribution Network Design looks at strategies to distribute finished products to customers efficiently at lowest cost - a difficult but important problem N1

- Factory storage with drop shipping
- Warehouse storage with package carrier delivery
- Local storage with last mile delivery
- Factory or Warehouse storage with consumer pickup
- Local storage with consumer pickup

Note 1 — ***Delivery Options***

- Drop shipping refers to the practice of shipping directly from the manufacturer to the end consumer, bypassing the warehouse who takes the order and initiates the delivery request. Delivery costs tend to be high with this option, but it eliminates the need to setup and manage warehouse delivery operations.
- In warehouse storage with package carrier delivery, the customers request are met from a few centralized warehouse strategically located in the region. Transportation cost tends to be lower than the previous option, but inventory holding and warehouse management costs increase.

- In the local storage with last mile delivery option, more warehouses need to be setup and spread across the region (leading to higher warehouse and inventory cost), but transportation cost will be lower.
- In the factory/Warehouse storage with consumer pickup approach, inventory is stored at the factory or warehouse but customers place their orders online and then go to designate pickup points to collect their orders. Inventory and warehousing cost will be low due to demand aggregation at the factory. however, new pickup sites have to built, leading to higher facility costs. A significant information infrastructure is also needed to provide visibility of the order until the customer picks it up.
- In the last option, inventory is stored locally at a warehouse (or retail outlet) nearby. Customers either walk into the warehouse/retail outlet or place an order online or on the phone and pickup at the warehouse/retail store. Local storage increases inventory costs but transprotation cost reduces. This option is more suitable for fast moving goods so that there is only marginal increase in inventory even with local storage.

Slide 3

How to set up the distribution network to move products from the factory to the customers?

2 Network Design

2.1 Issues

2.1.1 Location

Slide 4

Distribution Centers (DC): An intermediate storage location for stocking of finished products for subsequent shipment to customers.

- Economy of scale in bulk shipping from factory in DC

- Faster response time to ship from DC, rather than to ship from factory direct.

Location choices for the distribution centers?

- Different sites may have different facility fixed cost structures

2.1.2 Transportation

SLIDE 5

- Transportation Cost: **depends on which DC is being used to serve the customer**.
- Outbound Transportation cost
- Inbound Transportation cost

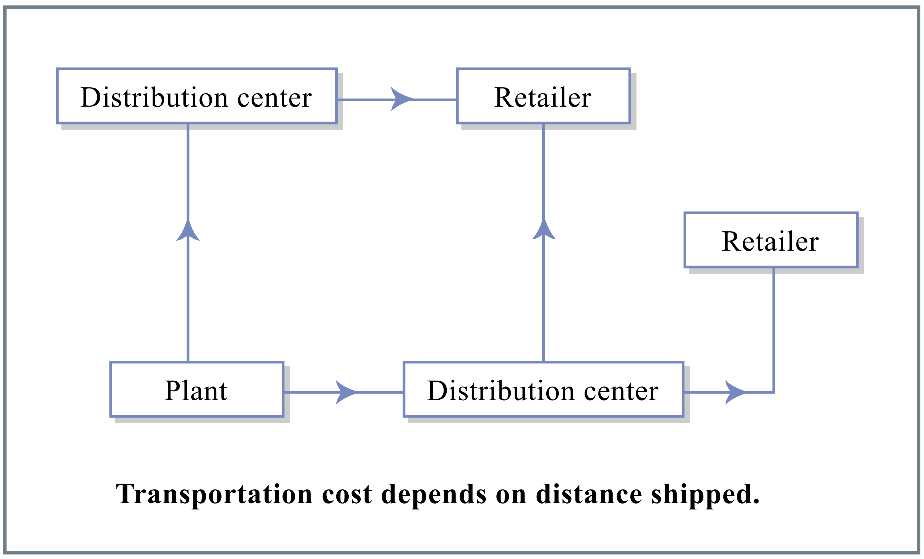

Figure by MIT OCW.

2.1.3 Storage and Material Handling

SLIDE 6

Warehousing and Material handling Cost: packaging, order-picking, replenishment. **Depends on the volume of business served**.

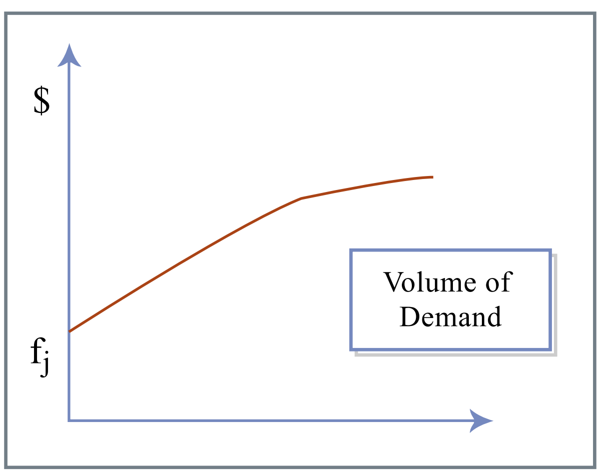

Figure by MIT OCW.

2.1.4 Safety Stock Inventory

SLIDE 7

Inventory Management refer to means by which inventories are managed. Inventories exist at every stage of the supply chain as

- raw materials,
- semi-finished,
- *finished goods*,
- in-transit between locations.

SLIDE 8

- Replenishment Lead Time: Time from order is being placed to time when shipment is received.
- To protect against stock out, company normally keeps **excess** inventory (called **Safety Stock**) to protect itself against stockout during the replenishment lead time.

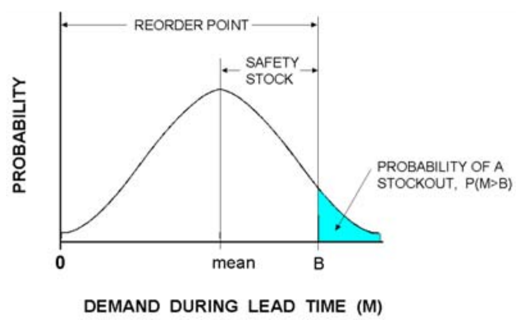

SLIDE 9

- **Safety Stock**: Excess above the expected lead time demand. Their primary purpose is to buffer against any uncertainty that might exist in the supply chain. But they cost $$$.
- Safety stock (SS) level depends on the **standard deviation** (SD) of the lead time demand (denoted by σ). Eg. SS = 3σ indicate a close to 99.7% safety protection against stockout.

Goal: To maintain high service level with the least safety stock investment

2.2 Comparison

SLIDE 10

No. of DCs	Few	Many
Fixed Cost		
Transportation Cost		
Handling Cost		
Safety Stock Cost		

2.3 The Problem

Slide 11

- Find
  - Number and location of DCs
  - Assignment of customers to DCs
- To minimize
  - Transportation costs: From factory to DC; DC to customers
  - Inventory (i.e., safety stock) costs at DCs
  - Warehousing and material handling cost.

2.4 Inputs and Parameters

Slide 12

- W: set of **potential** DC locations
- I: set of customer locations
- μ_i: mean (yearly) demand at customer i
- σ_i^2: variance of (daily) demand at customer i
- Operating Cost
  - f_j: fixed (annual) cost of locating a distribution center at location j
  - d_{ij}: cost per unit to ship from DC j to customer i
  - a_j: per unit shipment cost from factory to distribution center j
  - L_j: Replenishment lead time for DC j
  - h: inventory holding cost per unit of product per year

2.5 Cost Model

Slide 13

Decision Variables

- $X_j = 1$, if customer j is selected as a distribution center location, and 0, otherwise
- $Y_{ij} = 1$, if customer i is in served by a distribution center j, and 0, otherwise

Cost: Suppose customers in set S is served by DC j.

- Transportation cost for this set of customers

$$\sum_{i \in S} (\mu_i \times d_{j,i} + \mu_i \times a_j) = \sum_i Y_{ij} \Big((\mu_i \times d_{j,i} + \mu_i \times a_j) \Big)$$

- Warehousing/Material handling cost

$$f_j X_j + \Gamma_j (\sum_i \mu_i Y_{ij}).$$

- Safety Stock Inventory Cost

$$3h\sqrt{L_j}\sqrt{\sum_i \sigma_i^2 Y_{ij}}.$$

2.6 Complete Model

SLIDE 14

$$\min \sum_{j\in W}\left(f_j X_j + \beta\sum_{i\in I}\hat{d}_{ij}Y_{ij} + \theta\Gamma_j\left(\sum_{i\in I}\mu_i Y_{ij}\right) + \theta q_j\sqrt{\sum_{i\in I}\sigma_i^2 Y_{ij}}\right)$$

$$\text{subject to}\quad \begin{array}{ll} \sum_{j\in I} Y_{ij} = 1, & Y_{ij} - X_j \le 0 \\ Y_{ij}\in\{0,1\}, & X_j\in\{0,1\}\end{array}$$

- β and θ are weightage factors that we will adjust for computational testing.
- $\hat{d}_{ij} = \beta\mu_i(d_{j,i} + a_j)$
- $q_j = 3h\sqrt{L_j}$
- $\Gamma_j(\cdot)$ is concave and non-decreasing

Not surprisingly, this problem is NP-hard.

2.7 Set-Covering Model

SLIDE 15

- Decision variables: Let $z_{j,R} = 1$ if DC j is used to serve the customers in the set R.
- Constraints: Each customer must be served by one DC. Hence we want to find a partition of the entire customer set into $(R_1, \dots, R_k)$ and the corresponding DC assignment $(j_1, \dots, j_k)$. DC j_l will be used to serve the customers in R_l.
- Cost of using DC j to serve customers in R:

$$c_{j,R} = f_j + \beta\sum_{i\in R}\hat{d}_{ij} + \theta\Gamma_j(\sum_{i\in R}\mu_i) + \theta q_j\sqrt{\sum_{i\in R}\sigma_i^2})$$

SLIDE 16

Let $\mathcal{R}$ denote the set of all subsets of customers.

$$\begin{array}{ll} \min & \sum_{j\in W}\sum_{R\in\mathcal{R}} c_{j,R}z_{j,R} \\ \text{subject to} & \sum_{j\in I}\sum_{R\in\mathcal{R}:i\in R} z_{j,R} \ge 1, \quad \forall i\in I, \\ & z_{j,R}\in\{0,1\}, \quad \forall\ \ j,R. \end{array}$$

The optimal solution actually gives rise to a partition of the customers. N2

Note 2 ***Covering=Partition***

The set covering model proposed above can be shown to be equivalent to the following set-partitioning model:

$$\begin{array}{ll} \min & \sum_j \sum_{R\in\mathcal{R}} c_{j,R} z_{j,R} \\ \text{subject to} & \sum_{j\in I} \sum_{R\in\mathcal{R}: i\in R} z_{j,R} = 1, \quad \forall i \in I, \\ & z_{j,R} \in \{0,1\}, \quad \forall \ \ j, R. \end{array}$$

This follows from monotinicity of the cost function $c_{j,R}$, i.e., if $S \subset T$, $c_{j,S} \leq c_{j,T}$. Hence all optimal 0-1 solution to the set covering model must be a solution to the set partitioning model.

Slide 17

Example: $W = \{a, b\}, I = \{1, 2\}$.

$$\begin{array}{rcccccccccccccc} \min & c_{a,1}z_{a,1} & + & c_{b,1}z_{b,1} & + & c_{a,2}z_{a,2} & + & c_{b,2}z_{b,2} & + & c_{a,12}z_{a,12} & + & c_{b,12}z_{b,12} & & \\ \text{s.t.} & z_{a,1} & + & z_{b,1} & + & & & & & z_{a,12} & + & z_{b,12} & \geq & 1 \\ & & & & & z_{a,2} & + & z_{b,2} & + & z_{a,12} & + & z_{b,12} & \geq & 1 \\ & z_{a,1}, & & z_{b,1}, & & z_{a,2}, & & z_{b,2}, & & z_{a,12}, & & z_{b,12} & \geq & 0 \end{array}$$

2.7.1 Column Generation

Slide 18

- Number of variables in the formulation $= |W| \times 2^{|I|}$. Huge!
- We solve the LP relaxation of the set covering model using column generation
- Given a set of dual prices $\bar{\pi}$ (obtained by solving the restricted master problem), to find a column with negative reduced cost, we need to solve the following pricing problem:
  - Is there any S with $c_{j,S} - \sum_{i\in S} \pi_i < 0$?

Slide 19

To answer this question, we solve:

$$\min_S (c_{j,S} - \sum_{i\in S} \pi_i),$$

i.e.,

$$\min_S \left(f_j + \sum_{i\in S} (\beta \hat{d}_{ij} - \pi_i) + \theta \Gamma_j (\sum_{i\in S} \mu_i) + \theta q_j \sqrt{\sum_{i\in S} \sigma_i^2} \right)$$

2.7.2 Pricing Problem

Slide 20

For designated distribution center $j \in I$, define

$$\min_{S\subset I} g_j(S) \equiv \sum_{i\in S} a_i + \theta \Gamma_j (\sum_{i\in S} b_i) + \theta q_j \sqrt{\sum_{i\in S} c_i}.$$

where

$$\begin{array}{rcl} a_i & \equiv & \hat{d}_{ij} - \pi_i, \\ b_i & \equiv & \mu_i \geq 0, \\ c_i & \equiv & \sigma_i^2 \geq 0. \end{array}$$

WLOG: $a_i < 0$ for all i. Otherwise, i will not be in the optimal S. SLIDE 21

- There are $2^{|I|}$ possible choices of S.
- Impossible to enumerate over all possible S.

General form of the pricing problem:

$$\min_S h_1(\sum_{i\in S} a_i) + h_2(\sum_{i\in S} b_i) + h_3(\sum_{i\in S} c_i)$$

where h_1, h_2, h_3 are concave, non-decreasing functions. SLIDE 22

Proposition:

- Consider the convex hull $\mathcal{C}$ in 3D formed by the points $(\sum_{i\in S} a_i, \sum_{i\in S} b_i, \sum_{i\in S} c_i)$, for all subset S.
- Let S^* be an optimal solution to our pricing problem

$$\min_{S\subset I} g_j(S) \equiv \sum_{i\in S} a_i + \theta\Gamma_j(\sum_{i\in S} b_i) + \theta q_j \sqrt{\sum_{i\in S} c_i}.$$

- $(\sum_{i\in S^*} a_i, \sum_{i\in S^*} b_i, \sum_{i\in S^*} c_i)$ is a corner point in $\mathcal{C}$.

SLIDE 23

- S^*: a corner point in the convex hull $\mathcal{C}$.
- There exists (α, β, γ) such that

$$\begin{array}{rl} & \alpha a(S^*) + \beta b(S^*) + \gamma c(S^*) \\ = & \min_{(x,y,z)\in\mathcal{C}}(\alpha x + \beta y + \gamma z) \\ = & \min_S(\alpha a(S) + \beta b(S) + \gamma c(S)) \\ = & \min_S \sum_{i\in S}(\alpha a_i + \beta b_i + \gamma c_i). \end{array}$$

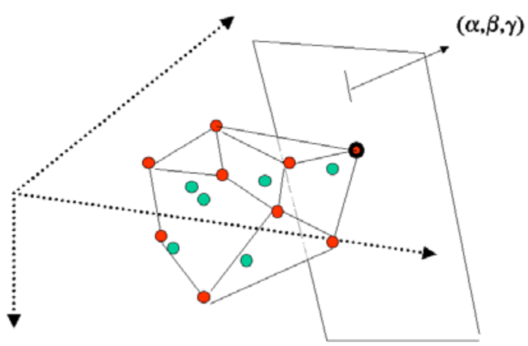

SLIDE 24

Observation:

- $i \in S^*$ iff $\alpha a_i + \beta b_i + \gamma c_i < 0$. Since $a_i < 0$ for all i, this is equivalent to

$$\text{CHECK:}\alpha + \beta \frac{b_i}{a_i} + \gamma \frac{c_i}{a_i} > 0?$$

- **To recover S^*, search through all possible α, β and γ!.**
- Issue: How to find α, β, γ?

SLIDE 25

Consider the set of points $(\frac{b_i}{a_i}, \frac{c_i}{a_i})$, $i = 1, ..., |I|$.

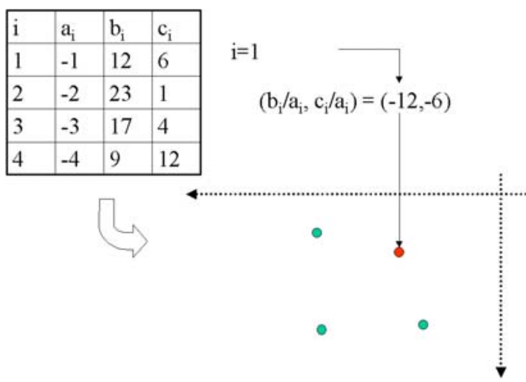

i	a_i	b_i	c_i
1	-1	12	6
2	-2	23	1
3	-3	17	4
4	-4	9	12

SLIDE 26

The set S^* contains points that lie on one side of the line $\alpha + \beta \frac{b_i}{a_i} + \gamma \frac{c_i}{a_i} = 0$

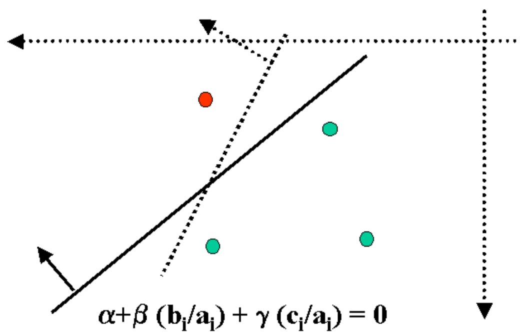

3 Computational Results

3.1 Test Instances

SLIDE 27

Transportation cost is proportional to the euclidian distance in the plane.

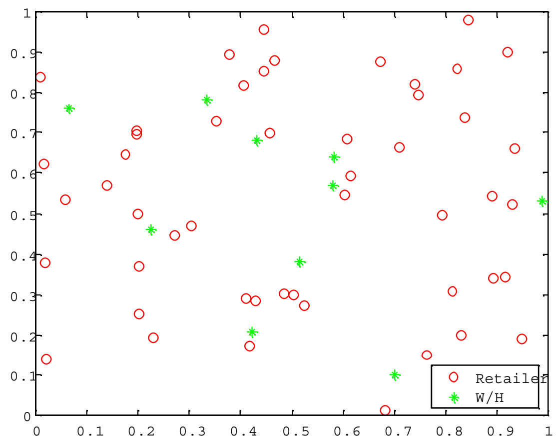

Figure 1: Warehouse-Retailer Locations

3.2 Solving the pricing problem

SLIDE 28

No. of Retailers	Total CPU Time(seconds)
10	0.01
20	0.07
40	0.73
80	6.91
160	60.3
320	554

3.3 40 customers and 40 warehouses

SLIDE 29

INPUT		OUTPUT			
β	θ	No. of DCs Opened	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	5	62	2661	1
0.002	0.1	7	16	961	1.001
0.003	0.1	8	7	543	1
0.004	0.1	10	6	466	1
0.005	0.1	14	5	364	1
0.001	0.1	5	62	2661	1
0.002	0.2	6	16	974	1.001
0.005	0.5	9	6	492	1
0.005	0.1	14	5	364	1
0.005	1	11	6	430	1.003
0.005	5	8	9	828	1
0.005	10	7	28	1643	1.001

3.4 80 customers and 80 warehouses

SLIDE 30

INPUT		OUTPUT			
β	θ	No. of DCs Opened	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	7	414	9984	1
0.002	0.1	9	201	6031	1
0.003	0.1	12	94	2548	1
0.004	0.1	21	47	1042	1
0.005	0.1	24	26	511	1
0.001	0.1	7	414	9984	1
0.002	0.2	13	84	2077	1
0.005	0.5	18	52	1196	1
0.005	0.1	24	26	511	1
0.005	1	10	102	3096	1
0.005	10	8	364	9072	1

3.5 120 customers and 120 warehouses

SLIDE 31

INPUT		OUTPUT			
β	θ	No. of DCs Opened	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.0001	0.01	11	742	21537	1
0.0002	0.01	15	516	11476	1
0.0003	0.01	24	213	4997	1.001
0.0004	0.01	28	103	2014	1
0.0005	0.01	33	57	1043	1
0.0001	0.01	11	742	21537	1
0.0002	0.02	23	215	5021	1
0.0005	0.05	29	99	1998	1
0.0005	0.01	33	57	1043	1
0.0005	0.1	15	522	11628	1
0.0005	0.5	12	635	17142	1.002

4 Column Generation

4.1 Speeding Up

SLIDE 32

- Bottleneck: For each j, a related pricing problem where the j is assumed to be the DC is solved.
- How to choose j wisely?

Approach: We use information on the primal and dual solution to “fix” variables, so that we can determine whether a DC will or will not be selected in an optimal solution early in the column generation routine.

4.2 Variable Fixing

SLIDE 33

Recall that the set covering model we are trying to solve is of the form:

$$\begin{array}{ll} \min & \sum_{j \in W} \sum_{R \in \mathcal{R}} c_{j,R} z_{j,R} \\ \text{subject to} & \sum_{R \in \mathcal{R}: i \in R} \left(\sum_{j \in R} z_{j,R} \right) \geq 1, \quad \forall i \in I, \\ & z_{j,R} \in \{0,1\}, \quad \forall R \in \mathcal{R}. \end{array}$$

At each stage of the column generation routine, we have: SLIDE 34

- A set of dual prices $\{\pi_j\}$.
- A set of primal feasible (fractional) solution $z_{j,S}$.

- After solving the pricing problem (one for each potential DC), we obtain the reduced cost $r_j \equiv \min_S\left(c_{j,S} - \sum_{k\in S}\pi_k\right)$. **Note that some of the r_j's may be non-negative**.

Let Z_{IP} and Z_{LP} denote the optimal integral and fractional solution to the set covering problem. SLIDE 35

Claim 1 . *$\sum_{j:r_j\le 0} r_j + \sum_j \pi_j$ is a lowerbound to Z_{LP}. Hence it is a lowerbound to Z_{IP} too.*

N3

Let j^* be a customer such that $r_j^* > 0$. Let UB be an upperbound for Z_{IP}.

Claim 2 . *If $\sum_{j:r_j\le 0} r_j + \sum_j \pi_j + r_{j^*} > UB$, then customer j^* will* **never be used** *as a DC in the optimal solution to the (integral) set covering problem.*

N4

Note 3 ***Proof***

Proof: Note that for the set partitioning model, by monotonicity of cost function $c_{R,j}$, we know that there is an optimal LP solution which satisfies each covering constraint at equality. Hence the optimal LP solution is actually a solution to a set partitioning problem. For each customer j, the constraint $\sum_{S:j\in S} z_{S,j} \le 1$ is thus a redundant constraint implicit in the formulation.
The Lagrangian dual of the LP relaxation is thus equivalent to

$$L(\lambda) = \sum_j \lambda_j + \min\left\{\left(\sum_j \sum_{S:j\in S}(c_{j,S} - \sum_{k\in S}\lambda_k) z_{j,S}\right) : 1 \ge z_{j,S} \ge 0, \sum_{S:j\in S} z_{j,S} \le 1 \,\forall\, j\right\}.$$

The problem decomposes for each customer j, and hence $Z_{LP} = \max_\lambda L(\lambda) \ge L(\pi) = \sum_j \pi_j + \sum_{j:r_j\le 0} r_j$.

Note 4 ***Proof***

Proof: To see this, suppose otherwise. Then Z_{IP} remains unchanged if we impose the additional condition: $\sum_{S:j^*\in S} z_{j^*,S} = 1$ to the set covering formulation for customer j^*. The Lagrangian dual, in this case, reduces to

$$L'(\lambda) = \sum_j \lambda_j + \min\left\{\left(\sum_j \sum_{S:j\in S}(c_{j,S} - \sum_{k\in S}\lambda_k) z_{j,S}\right) : 1 \ge z_{j,S} \ge 0, \sum_{S:j\in S} z_{j,S} \le 1 \,\forall\, j, \sum_{S:j^*\in S} z_{j^*,S} = 1\right\}.$$

Hence $Z_{IP} \ge \max_\lambda L'(\lambda) \ge L'(\pi) = \sum_j \pi_j + \sum_{j:r_j\le 0} r_j + r_{j^*}$. On the other hand, we have $Z_{IP} \le UB$. This gives rise to a contradiction.

4.2.1 Heuristic

SLIDE 36

The variable fixing method depends largely on the quality of the upperbound UB.

- We generate an upperbound for the IP by generating a feasible (integral) solution using some heuristic.
  - Let z^* be the optimal LP solution obtained by solving the problem using a **partial** set of columns.
  - Order the customers accoring to non-decreasing value of μ_i.
  - Starting from the first customer (say $i = 1$) on the list,
    - For some S and j, $i \in S$ and $z^*_{j,S} = 1$: customer i is served by DC j.
    - Otherwise, if there exists S, T, both containing i, and j, k are the designated DCs for set S and T respectively: such that $z^*_{j,S} > 0$, $z^*_{T,k} > 0$. We serve i using the DC that will lead to the least total cost.
  - Proceed to the next customer, until all customers have been assigned.

4.3 Computational Results

4.3.1 40 Customers and Warehouses

SLIDE 37

INPUT		OUTPUT				
β	θ	No. of DCs Opened	No. of DCs OUT	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	5	34	6	117	1
0.002	0.1	7	31	3	73	1
0.003	0.1	8	31	2	66	1
0.004	0.1	10	29	1	44	1.001
0.005	0.1	13	27	1	32	1
0.001	0.1	5	34	6	117	1
0.002	0.2	6	34	4	92	1.001
0.005	0.5	8	32	2	56	1
0.005	0.1	13	27	1	32	1
0.005	1	10	29	1	40	1.002
0.005	5	7	32	4	96	1
0.005	10	5	34	6	122	1.001

4.3.2 80 Customers and Warehouses

SLIDE 38

INPUT		OUTPUT				
β	θ	No. of DCs Opened	No. of DCs OUT	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	6	74	31	348	1
0.002	0.1	8	72	20	202	1
0.003	0.1	12	67	14	147	1
0.004	0.1	21	58	9	84	1
0.005	0.1	24	56	7	66	1
0.001	0.1	6	74	31	348	1
0.002	0.2	13	66	14	142	1
0.005	0.5	18	62	10	113	1
0.005	0.1	24	56	7	66	1
0.005	1	10	69	15	162	1
0.005	10	7	72	27	303	1

4.3.3 120 Customers and Warehouses

SLIDE 39

INPUT		OUTPUT				
β	θ	No. of DCs Opened	No. of DCs OUT	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.0001	0.01	10	109	94	804	1
0.0002	0.01	15	105	61	480	1
0.0003	0.01	24	94	38	296	1
0.0004	0.01	28	90	27	184	1
0.0005	0.01	33	87	16	97	1
0.0001	0.01	10	109	94	804	1
0.0002	0.02	23	96	38	309	1
0.0005	0.05	29	90	26	166	1
0.0005	0.01	33	87	16	97	1
0.0005	0.1	15	104	62	496	1
0.0005	0.5	11	109	90	743	1.002

4.3.4 500 Customers and Warehouses

Slide 40

INPUT		OUTPUT				
β	θ	No. of DCs Opened	No. of DCs OUT	CPU Time	No. of Columns Generated	Z_H/Z_{LP}
0.0001	0.01	42	458	512	3742	1
0.0002	0.01	57	442	426	2819	1.001
0.0003	0.01	95	404	248	1405	1
0.0004	0.01	114	386	146	717	1
0.0005	0.01	146	354	86	446	1
0.0001	0.01	42	458	512	3742	1
0.0002	0.02	90	409	314	1833	1
0.0005	0.05	132	368	117	586	1
0.0005	0.01	146	354	86	446	1
0.0005	0.1	61	439	404	2633	1
0.0005	0.5	44	455	503	3572	1

5 Applications and Extensions

Slide 41

The model captures the three most important concerns in distribution network design: Transportation cost, Material and Warehousing Cost and Safety stock inventory cost

How about:

- Service level, say defined in terms of response time?
- Robustness issues, say a set of scenarios (concerning cost and input parameters) are given?
- Capacity issues, say a DC can only handle up to a fixed amount of demand ?