Alcatel·Lucent

Newsvendor Networks and Assemble-to-Order Inventory Systems

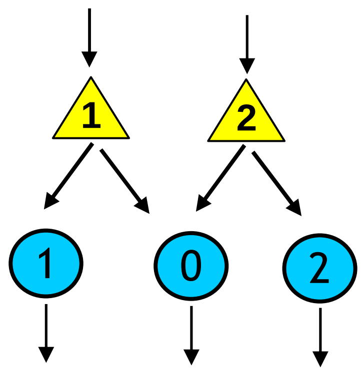

Marty Reiman
Alcatel-Lucent Bell Labs
(Based on joint work with Mustafa Dogru and Qiong Wang)

Talk Outline

- The Assemble-to-Order (ATO) inventory system
- The inventory control problem
- A related stochastic program (SP): The ‘Newsvendor Network’
- A relaxed SP
- SP solution for W model & translation into control policy
- Long lead time limits
- SP solution for M model & translation into control policy

The Assemble-to-Order (ATO) Inventory System

- *m* products assembled from *n* components
- Product i requires a_{ij} units of component j
- Stochastic demand for products: lead time demand D_i
- Identical component procurement lead time *L* *(suppliers are uncapacitated)*
- Assembly time is negligible
- Only component inventories are kept
- Backlogging: unit backlog cost b_i
- Unit inventory holding cost h_j
- Continuous or periodic review (compound Poisson or i.i.d. demand)

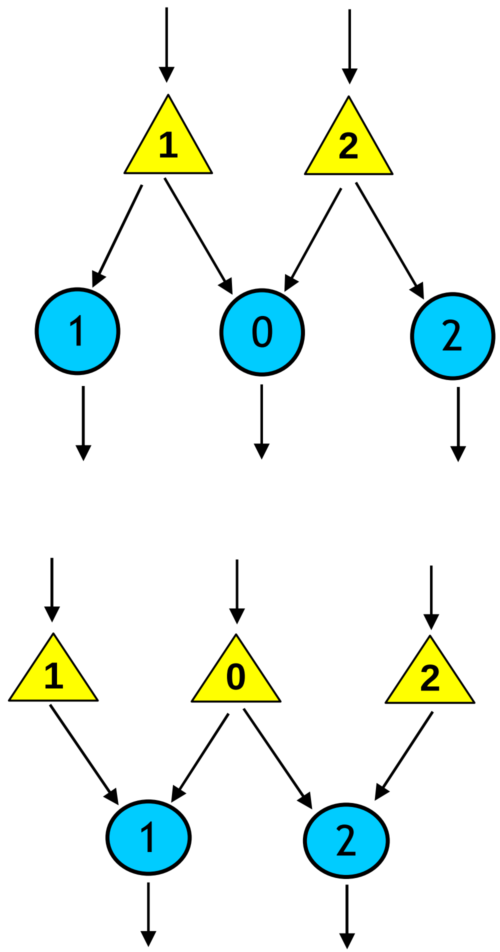

Goal: Find replenishment & allocation policies to minimize the long run average expected cost

The (Continuous Review) Demand Process

Let

$$\mathcal{D}_i(t) = \text{demand for product } i \text{ during } [0, t],\ \ 1 \leq i \leq m,$$

and $\mathcal{D}(t) \equiv (\mathcal{D}_1(t), \ldots, \mathcal{D}_m(t))$.

Assume that $\{\mathcal{D}(t), t \geq 0\}$ is a compound Poisson process.

Let

$$\Delta_i \equiv \mathsf{E}[(\mathcal{D}_i(1)],\ \ 1 \leq i \leq m, \quad \Delta \equiv (\Delta_1, \ldots, \Delta_m),$$

and

$$\sigma_{ij}^2 \equiv \mathsf{E}[(\mathcal{D}_i(1) - \Delta_i)(\mathcal{D}_j(1) - \Delta_j)].$$

Assume that

$$0 < \Delta_i < \infty \text{ and } 0 < \sigma_{ii}^2 < \infty, 1 \leq i \leq m.$$

Alcatel·Lucent

The (Continuous Review) Inventory Control Problem

- We want to choose a replenishment policy γ and an allocation policy p to minimize

$$C^{\gamma,p} \equiv \limsup_{T\to\infty} \frac{1}{T} \mathsf{E}\left[\int_0^T \left\{\sum_{i=1}^m b_i B_i(t) + \sum_{j=1}^n h_j I_j(t)\right\} dt\right]$$

where $\mathbf{B}_i(t)$ is the product i backlog at time t and $\mathbf{I}_j(t)$ is the component j inventory at time t.

- Feasible policy:

1. $B_i(t) \geq 0$
2. $I_j(t) \geq 0$
3. Decisions cannot be based on future demand information

The Stochastic Programming Based Approach

1. Introduce a (two stage) stochastic (linear) program (with complete recourse) whose solution provides a lower bound on the achievable cost in the inventory control system
2. Solve the stochastic program (SP)
3. Translate the SP solution into a control policy for the inventory system
4. Prove asymptotic optimality

This is similar to the approach introduced in Harrison (1988) and used in many papers since then

A Related Stochastic Program (The 'Newsvendor Network')

Let $y = (y_1, \ldots, y_n)$.
Choose $y \geq 0$ to minimize $C_s(y) \equiv \mathsf{E}[\varphi(y;D)] + \sum_{j=1}^{n} h_j y_j$, where

$$\begin{aligned}\varphi(y;D) &= \min_{z\geq 0}\left\{\sum_{i=1}^{m} b_i(D_i - z_i) - \sum_{j=1}^{n} h_j \sum_{i=1}^{m} a_{ij} z_i \,\middle|\, z_i \leq D_i, 1 \leq 1 \leq m, \sum_{i=1}^{m} a_{ij} z_i \leq y_j, 1 \leq j \leq n\right\} \\ &= b \cdot D - \max_{z\geq 0}\left\{c \cdot z \,\middle|\, z \leq D,\ zA \leq y\right\},\end{aligned}$$

$c_i = b_i + \sum_{j=1}^{n} a_{ij} h_j$, and D_i represents the demand for product i over a lead time (so that $D_i \overset{d}{=} \mathcal{D}_i(L)$).

Let $C_s^* = \min_{y\geq 0} C_s(y)$.

- This stochastic program 'relaxes' the allocation decisions (z): They are made at the end of the lead time, as opposed to each demand/replenishment arrival time. But it is not a relaxation of the inventory control problem because it assumes initial conditions B=0 and I=0.

References on Newsvendor Networks

- 1 Period ATO models
  - Baker, Magazine and Nuttle (1986)
  - Gerchak, Magazine and Gamble (1988)
  - Song and Zipkin (2003)
- Capacity Investment/Production/Inventory
  - Harrison & Van Mieghem (1999)
  - Van Mieghem and Rudi (2002)
- Call Center Staffing/Routing
  - Harrison and Zeevi (2005)
  - Bassamboo, Harrison and Zeevi (2006)

The Relaxed Stochastic Program

Let $\alpha = (\alpha_1, \ldots, \alpha_m)$.
Choose $y \geq 0$ and $\alpha \geq 0$ to minimize $\underline{C}_s(y, \alpha) \equiv \mathsf{E}[\varphi(y, \alpha; D)] + h \cdot y + b \cdot \alpha$, where

$$\varphi(y, \alpha; D) = b \cdot D - \max_{z \geq 0} \left\{ c \cdot z \middle| z \leq D + \alpha, zA \leq y \right\}.$$

Let

$$\underline{C}_s^* \equiv \inf_{y \geq 0, \alpha \geq 0} \underline{C}_s(y, \alpha).$$

Theorem (Dogru, Reiman, Wang, 2009). Let (γ, p) be any feasible policy, and $C^{\gamma,p}$ be the corresponding cost. If the demand is an i.i.d. sequence for the periodic review case, or a compound Poisson process in the continuous review case then

$$\underline{C}_s^* \leq C^{\gamma,p}.$$

Alcatel·Lucent

Solution of SP for W Model and Translation into Control Policy

- The newsvendor network for the W Model can be solved exactly. (There is no need for sampling.)
- The Relaxed SP can be transformed and solved exactly. (For all parameter values in our experiments, the solutions of the newsvendor network and relaxed SPs for the W model coincide.)
- For replenishment, follow a base-stock policy, where the base-stock levels are (y_0^*, y_1^*, y_2^*), the solution of the newsvendor network.
- For allocation, follow a priority policy, which is motivated by the solution of the recourse LP.

Base-Stock Policies

Let $\mathcal{R}(t) =$ total replenishment orders for component j placed up to time t
and $R_j(t) = \mathcal{R}(t) - \mathcal{R}(t-L)$: lead time replenishment orders.
The 'inventory position' of component j is defined as

$$Y_j(t) = I_j(t) - \sum_{i=1}^{m} a_{ij} B_i(t) + R_j(t), \;\; 1 \le j \le n.$$

A base-stock policy with base-stock levels $y = (y_1, \dots, y_n)$ places replenishment orders to keep

$$Y_j(t) = y_j, \;\; 1 \le j \le n.$$

After an initial 'start-up' order, this is 'order-for-order replenishment', so that

$$R_j(t) = \sum_{i=1}^{m} a_{ij} [\mathcal{D}_i(t) - \mathcal{D}_i(t-L)].$$

The Solution of the Recourse LP for the W System

Here $c_1 = b_1 + h_0 + h_1$ and $c_2 = b_2 + h_0 + h_2$.

Given $y \geq 0$ the recourse LP for the W system is

$$\max_{z \geq 0} \left\{ c_1 z_1 + c_2 z_2 \middle| z_i \leq D_i,\ z_i \leq y_i, i = 1, 2,\ z_1 + z_2 \leq y_0 \right\}.$$

Assume (without loss of generality) that $c_1 \geq c_2$.

The *recourse* LP has solution

$$z_1^* = D_1 \wedge y_1, \quad z_2^* = D_2 \wedge y_2 \wedge \left(y_0 - z_1^*\right).$$

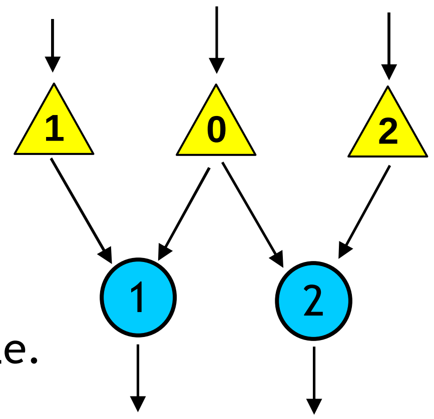

This motivates priority to product 1 in allocation:

- A demand is served as long as all required components are available.
- If both products have backlogs due to the lack of the common component, when a replenishment arrives, then all product 1 backlogs are cleared first (as long as there are enough unique components) before serving product 2 demand

Alcatel·Lucent

The SP-Based Solution is Sometimes Optimal for W Model

- Symmetric cost case: $c_1=h_0+h_1+b_1=c_2=h_0+h_2+b_2$
  - All ‘myopic’ policies (maintaining $I_i * I_0 * B_i=0$, i=1,2) achieve the same total allocation (z_1+z_2) as the associated recourse LP.
  - With $c_1=c_2$, $C_s^* = \underline{C}_s^*$ and all allocations with the same total have the same cost.
- SP solution has $y_0^* = y_1^* + y_2^*$
  - In this case the supply chain decomposes into 2 separate supply chains: There is no advantage to component commonality here, even if the common component costs the same as the components that it is replacing.
  - Numerical examples show that this is not a rare occurrence, especially if h_1 and/or h_2 are large relative to h_0

Asymptotic Optimality?

For $0 < L < \infty$, let $\underline{C}_s^*(L)$ denote the optimal objective of the relaxed SP with demand $\mathcal{D}(L)$.

Let $C_{SP}^{(L)}$ denote the expected long-run average cost of the inventory system using the SP based policy in the W system with lead time L.

Conjecture. $$\frac{C_{SP}^{(L)}}{\underline{C}_s^*(L)} \to 1 \quad \text{as } L \to \infty.$$

This is borne out numerically.

Alcatel·Lucent

The Component Shortage Process

Let

$$Q_j(t) = \sum_{i=1}^{m} a_{ij} B_i(t) - I_j(t) : \text{ component j 'shortage'}$$

Recall that, under a base-stock policy,

$$I_j(t) - \sum_{i=1}^{m} a_{ij} B_i(t) + R_j(t) = y_j, \; 1 \leq j \leq n, \; t \geq 0,$$

and $R_j = \sum_{i=1}^{m} a_{ij}[\mathcal{D}_i(t) - \mathcal{D}_i(t-L)]$.
Thus

$$Q_j(t) = \sum_{i=1}^{m} a_{ij}[\mathcal{D}_i(t) - \mathcal{D}_i(t-L)] - y_j.$$

This process is not controllable.

Target Backlog Levels

Given $Q(t) = Q$, we can look for feasible values of $B(t)$ and $I(t)$ that minimize the total inventory/backlog cost:

$$\min_{B,I\geq 0}\{\sum_{i=1}^{m} b_i B_i + \sum_{j=1}^{n} h_j I_j | \sum_{i=1}^{m} a_{ij} B_i - I_j = Q_j\}$$

$$= \min_{B\geq 0}\{\sum_{i=1}^{m} c_i B_i | \sum_{i=1}^{m} a_{ij} B_i \geq Q_j\} - \sum_{j=1}^{n} h_j Q_j$$

This is a translation of the recourse LP of the relaxed SP.
Let $B^*(Q)$ denote an optimal solution.
If for all $t \geq 0$ we can control the suystem so that $B(t) = B^*(Q(t))$ then the inventory system will achieve the lower bound.

Functional Central Limit Theorem

Let

$$\hat{\mathcal{D}}^{(L)}(t) = \frac{\mathcal{D}(Lt) - LT\Delta}{\sqrt{L}}, \; L > 0, \; t \geq 0, \; \text{and } \hat{Q}^{(L)}(t) = \frac{Q(Lt)}{\sqrt{L}}.$$

By Donsker's (Functional Central Limit) Theorem,

$$\hat{\mathcal{D}}^{(L)} \xrightarrow{d} \hat{\mathcal{D}},$$

where $\hat{\mathcal{D}}$ is a driftless Brownian motion with covariance matrix $\{\sigma_{ij}^2, 1 \leq i, j \leq m\}$. This motivates representing $y(L)$ and $z(L)$ as

$$y(L) = L\Delta A + \sqrt{L}\hat{y}(L) \quad \text{and} \quad z(L) = L\Delta + \sqrt{L}\hat{z}(L).$$

If $\hat{y}(L) \to \hat{y}$ as $L \to \infty$, then

$$\hat{Q}^{(L)} \xrightarrow{d} \hat{Q} \text{ as } L \to \infty,$$

where

$$\hat{Q}_j(t) = \sum_{j=1}^{n} [\hat{\mathcal{D}}_j(t) - \hat{\mathcal{D}}_j(t-1)] - \hat{y}_j.$$

State Space Collapse for W

Let

$$\hat{B}^{(L)}(t) = \frac{B^{(L)}(Lt)}{\sqrt{L}} \text{ and } \hat{I}^{(L)}(t) = \frac{I^{(L)}(Lt)}{\sqrt{L}}.$$

For the W system,

$$\hat{Q}_0^{(L)}(t) = \hat{B}_1^{(L)}(t) + \hat{B}_2^{(L)}(t) - \hat{I}_0^{(L)}(t),$$

$$\hat{Q}_1^{(L)}(t) = \hat{B}_1^{(L)}(t) - \hat{I}_1^{(L)}(t), \quad \hat{Q}_2^{(L)}(t) = \hat{B}_2^{(L)}(t) - \hat{I}_2^{(L)}(t).$$

The target backlog levels, given $\hat{Q}^{(L)}(t) = \tilde{Q}$, are

$$\tilde{B}_1 = \tilde{Q}_1^+ \text{ and } \tilde{B}_2 = \tilde{Q}_2^+ \vee (\tilde{Q}_0 - \tilde{Q}_1^+)^+.$$

Suppose that $\hat{B}_1^{(L)}(t) > \hat{Q}_1^+(t)$. Then $\hat{I}_1^{(L)}(t) > 0$.
With product 1 receiving priority in using component 0,

$$\hat{B}_1^{(L)}(t + \frac{s}{\sqrt{L}}) - \hat{B}_1^{(L)}(t) \approx -\Delta_2 s.$$

So $\hat{B}_1^{(L)}$ hits $\hat{Q}_1^+$ in $O(\frac{1}{\sqrt{L}})$ time. By Bramson(1998), this implies state-space collapse.

Solution of SP for M Model and Translation into Control Policy

- The newsvendor network for the M Model can be solved exactly. (There is no need for sampling.)
- The Relaxed SP for the M model can be transformed and solved exactly. (The solutions of the newsvendor network and relaxed SPs for the M model are sometimes different.)
- For replenishment, follow a base-stock policy, where the base-stock levels are (y_1^*, y_2^*), the solution of the newsvendor network.
- For allocation, follow a priority policy, possibly with reservation.
  - The priority is motivated by the solution of the recourse LP
  - The need for reservation follows from a 'fluid analysis'

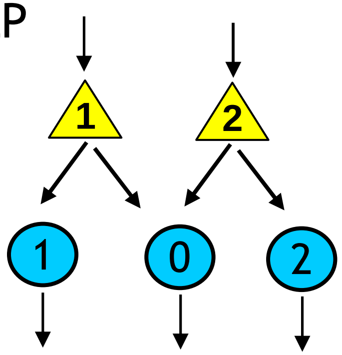

Alcatel·Lucent

Parameter Regions for M System

Here $c_0 = b_0 + h_1 + h_2, c_1 = b_1 + h_1, c_2 = b_2 + h_2$.

Assuming (without loss of generality) that $c_1 \geq c_2$, there are 4 cost parameter regions:

(A) $c_1 + c_2 \leq c_0$ (B) $c_1 \leq c_0 \leq c_1 + c_2$

(C) $c_2 \leq c_0 \leq c_1$ (D) $c_0 \leq c_2$

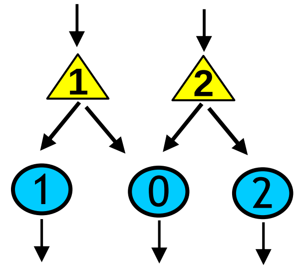

- The recourse solution in region A motivates priority to product 0 in allocation
  - Priority alone does not yield asymptotically optimal performance
  - Minimal reservation: hold 1 unit of component 2 in reserve for product 0. Fluid analysis (and numerical results) suggest state-space collapse
- The recourse solution in region B motivates a state-dependent policy:
  - If all products have backlogs, clear product 1 and 2 backlogs before product 0
  - If only products 0 and 1 have backlogs, clear product 0 backlog first
  - If only products 0 and 2 have backlogs, clear product 0 backlog first

Alcatel·Lucent