

Lecture Plan

1. The Economic Lot Scheduling Problem
2. The Joint Replenishment Problem

1 The Economic Lot Scheduling Problem

Economies of scale often dictate the choice of a single high speed machine capable of producing a given set of items over the choice of one dedicated machine for each item in the set. A scheduling problem arises when the high speed machine can only produce one item at a time. The problem is further complicated when the machine needs to be setup (often at a cost of both time and money) before a different item can be produced. The Economic Lot Scheduling Problem (ELSP) is that of scheduling the production of a set of items in a single machine to minimizing the long run average holding and set up cost under the assumptions of known constant demand and production rates.

Unfortunately, the ELSP belongs to a class of problems, known as NP-hard, see Gallego and Shaw [3], for which it is unlikely that an efficient solution procedure would ever be developed. What makes the ELSP hard is the coordination problem that needs to be resolved among the items competing for the machine capacity. Thus, much of the research effort on the ELSP (over 40 published papers) has dealt with the problem of finding low cost schedules via heuristics. Fortunately, good heuristics have been developed to obtain feasible production schedules whose average cost is often (but not guaranteed to be) close to the optimal. See Dobson [1]. Recently the research focus has shifted to deal with the more practical issues of real time scheduling, setup time reduction and the cost of offering variety.

We will start by presenting a lower bound on the long run average cost. This bound will be later used to assess the performance of the rotation schedule (RS) heuristic which calls for setting up and producing the items once per cycle in a given order. Our next goal is to discuss issues that arise in trying to implement feasible ELSP schedules. Our third goal is concerned with a brief discussion of the cost of offering variety and the effect of reducing setup times.

The data for the problem are:

- $i = 1, \dots, n$ index for the items,
- μ_i constant production rate of item i ,
- λ_i constant demand rate of item i ,
- h_i inventory holding cost of item i ,
- s_i set up time of item i ,
- K_i set up cost of item i .

For convenience let

$$H_i = 0.5h_i\lambda_i(1 - \rho_i),$$

where $\rho_i = \lambda_i/\mu_i$. With this notation, the average holding and ordering cost of using order interval T_i , for item i , is

$$C_i(T_i) = \frac{K_i}{T_i} + H_i T_i. \quad (1)$$

Clearly $T_i^* = \sqrt{K_i/H_i}$ minimizes (1) and $C_i(T_i^*) = 2\sqrt{K_i H_i}$. It is easy to verify that $C_i(T_i)$ can be written as

$$C_i(T_i) = C_i(T_i^*) \frac{1}{2} \left(\frac{T_i^*}{T_i} + \frac{T_i}{T_i^*} \right), \quad (2)$$

where (2) is often used to study the sensitivity of the cost function (1) to suboptimal choices of the order interval.

1.1 A Lower Bound on the Average Cost

A lower bound on the long run average cost is obtained by observing that feasible ELSP schedules must satisfy two constraints and then relaxing one of them. The *capacity* constraint arises because production rates are finite, while the *synchronization* constraint arises because the facility can only be used to setup or produce one item at a time. It is the synchronization constraint that makes the ELSP difficult. Our lower bound is obtained by finding the minimum cost schedule that satisfies the capacity constraint.

To derive the capacity constraint, notice that $s_i + \rho_i T_i$ units of time are needed to produce item i every T_i units of time. Consequently, $s_i/T_i + \rho_i$ is the proportion of time the machine is dedicated to setting up and producing item i , and

$$\sum_{i=1}^n \frac{s_i}{T_i} + \sum_{i=1}^n \rho_i$$

is the proportion of time that the machine is dedicated to setups and production of all the items. This proportion must be at most one, since otherwise the machine cannot keep up with demand.

Let $\kappa = 1 - \sum_{i=1}^n \rho_i$ denote the proportion of the machine time that is available for setups and idling. We assume that $\kappa > 0$. With this notation the lower bound is given by

$$\begin{aligned} C^{LB} &= \min \sum_{i=1}^n C_i(T_i) \\ \text{s.t.} \quad &\sum_{i=1}^n \frac{s_i}{T_i} \leq \kappa \\ &T_i \geq 0 \quad i = 1, \dots, n. \end{aligned} \tag{3}$$

Although the solution to the lower bound problem cannot be obtained in closed form when the capacity constraint is binding, it can be found numerically by a simple line search on the dual variable of the capacity constraint, or by the application of standard non-linear programming codes.

To be more precise, the solution to the lower bound problem can be shown to be of the form

$$T_i(a) = \sqrt{\frac{K_i + as_i}{H_i}}, \quad i = 1, \dots, n$$

for some $a \geq 0$. Here a is the economic value of the facility per unit time, e.g., its value per hour. You can think of $K_i + as_i$ as the total setup cost that consists of the out of pocket cost K_i plus the value consumed in setups as_i .

The problem reduces to finding the value of a . To do this you can try $a = 0$. If the capacity constrained is satisfied at $a = 0$ then this is the solution to the lower bound problem. If, however, the capacity constrained is violated, then we start increasing a until

$$\sum_{i=1}^n \frac{s_i}{T_i(a)} = \kappa.$$

Many commercial packages are available to solve non-linear programs. Excel, for example, has solver which can easily be invoked to solve the lower bound problem. See the excel file ELSP-Calculator.xls for more details.

Example:

Consider the following five item problem

The cost lower bound for this problem is \$2,140.62.

item	μ_i	λ_i	h_i	s_i (days)	K_i
1	10,000	900	2	1	25
2	5,000	1,000	1	5	25
3	6,000	1,000	5	10	5
4	4,000	500	3	2	10
5	7,500	1,500	4	5	50

Table 1: Data for the Economic Lot Scheduling Problem

1.2 Rotation Schedules

The Rotation Schedule (RS) heuristic forces all items to share a common order interval, say $T_i = T$, so each item is produced once per cycle in a given order resulting in a rotation schedule. Since items are produced one a time there is no risk of scheduling the setup or the production of two items at the same time. Consequently, any rotation schedule (RS) that satisfies the capacity constraint will also satisfy the synchronization constraint. The cost of an optimal rotation schedule can be obtained by simply adding the constraint $T_i = T$ for all $i = 1, \dots, n$. Equivalently, the cost of the best rotation schedule is given by

$$C^{RS} = \min \sum_{i=1}^n C_i(T) \quad (4)$$

$$s.t. \quad \frac{\sum_{i=1}^n s_i}{T} \leq \kappa \quad (5)$$

$$T \geq 0. \quad (6)$$

In the absence of the capacity constraint (5) the optimal choice of T is

$$\hat{T} = \sqrt{\frac{\sum_i K_i}{\sum_i H_i}}.$$

If \hat{T} satisfies (5), then it is the optimal common order interval. Otherwise, the optimal common order interval it is the smallest T satisfying (5), i.e.,

$$T_{\min} = \frac{\sum_i s_i}{\kappa}.$$

The overall optimal common order interval is thus

$$T^{RS} = \max(\hat{T}, T_{\min}).$$

The cost C^{RS} is obtained by evaluating (4) at T^{RS} .

Example: Using the data from Table 2 we find that the optimal rotation schedule is $T^{RS} = 0.293$ and that $C^{RS} = \$2,253.66$ which is 5.28% higher than the lower bound.

1.3 Assessing the Quality of the Rotation Schedule Heuristic

Ideally, we would like to be able to say something about the ratio C^{RS}/C^* where C^* is the unknown optimal long run average cost. Since $C^{RS}/C^* \leq C^{RS}/C^{LB}$, obtaining an upper bound on the later ratio yields a useful bound for the former. Here we will be satisfied by bounding C^{RS}/C^{LB} for cases where the capacity constraint is *not* binding, e.g., $T^{RS} = \hat{T}$. In order to obtain an *upper bound* on the cost ratio we order the items according to the ratios $\frac{K_i}{H_i}$, so that items with small ratios come first. Let

$$\gamma = \frac{K_n h_1 \lambda_1}{K_1 h_n \lambda_n}$$

then, it is possible to show that

$$\frac{C^{RS}}{C^*} \leq \frac{1}{2} (\sqrt[4]{\gamma} + 1/\sqrt[4]{\gamma}). \quad (7)$$

The bound (7) indicates that the RS heuristic performs well if the items are *similar* as measured by the ratio γ . Indeed, the cost ratio is at most 6% when $\gamma \leq 4$. Notice that while a small ratio C^{RS}/C^{LB} is prove positive that C^{RS}/C^* is small, a large ratio C^{RS}/C^{LB} does not necessarily imply that the rotation schedule is bad since it is possible that the lower bound is far from C^* .

1.4 Implementing a Schedule in Practice

Our second goal is to briefly touch on the problems of implementing a rotation schedule. To understand the implementation problem more clearly, assume that we have already determined T^{RS} and that γ is fairly small so that the C^{RS} is close to C^{LB} . Assume further that we schedule the production of the items as prescribed by the rotation schedule. That is, we setup the machine for item 1, produce item 1 for $\rho_1 T^{RS}$ units of time, set up for item 2, produce item 2 for $\rho_2 T^{RS}$ units of time,..., set up for product n and produce item n for $\rho_n T^{RS}$ units of time. Within the schedule we may need to insert some idle time (if $\sum_i s_i < \kappa T^{RS}$) until the elapsed time is a complete cycle of length T^{RS} . At the end of the first cycle, we will be on target only if

1. The initial inventories are on target
2. The facility is perfectly reliable
3. The setups actually consume a constant amount of time
4. The demand and the production rates are actually constant
5. The raw materials, tools, and fixtures are all available when required.

At best 1-5 represent an idealized situation that rarely holds in practice. So how can we effectively manage the production of the items? One effective way of doing so, when backorders are allowed, is to determine target, or produce-up to, levels for each of the items and to produce the items in the sequence prescribed by the rotation schedule and stop the production of an item when its inventory reaches its target level. An item's target level is determined by taking into account the nature of the schedule disruptions, i.e., random demands, production rates, setup times, etc. A target level is optimal for an item if and only if the time average probability of being out of stock is the ratio of the item's holding to the holding plus backorder cost. Thus, if the backorder cost for an item is $b = 19h$ then, the item's target, or base stock levels, should be such that the time average probability of being out of stock is $h/(b + h) = 1/20 = 5\%$. Alternatively, if backorder costs are not available, target service levels may be set by management from which target levels can be computed as indicated above.

Notice that the policy we have just described ignores the inventory levels of the items not being produced, and it prescribes to produce an item up to its target level even if other items are suffering severe shortages. This policy, of ignoring the inventories of the other items, works well in recovering a rotation (or any other target schedule) when the backorder costs are proportional to the processing rates, i.e., when the quantities $b_i \mu_i / \lambda_i$, $i = 1, \dots, n$ are all approximately the same. If the proportionality condition fails to hold, a more sophisticated control policy can be used where the production run time of the current item depends, in a linear way, on the inventories of all the other items in the group. See Gallego [2]

1.5 Related Issues

In this section, we address our final goal and briefly discuss a few related issues including the additional cost of offering variety and the effect of reducing setup times.

1.5.1 The Cost of Offering Variety

Studying the variety problem in all its generality is a difficult task. Here we will be satisfied by studying an important special case, where the setup costs are zero. This means that there are no out of pocket costs associated with the setup itself, other than those associated with the loss of capacity due to the time the machine spends on setups.

Under this conditions the lower bound problem is given by

$$\begin{aligned}
 C^{LB} &= \min \sum_{i=1}^n H_i T_i \\
 \text{s.t.} & \quad \sum_{i=1}^n \frac{s_i}{T_i} \leq \kappa \\
 & \quad T_i \geq 0 \quad i = 1, \dots, n.
 \end{aligned} \tag{8}$$

Fortunately, this problem can be solved in closed form, with

$$T_i^{LB} = \frac{1}{\kappa} \sqrt{\frac{s_i}{H_i}} \sum_j \sqrt{H_j s_j}$$

and

$$C^{LB} = \frac{1}{\kappa} \left(\sum_{i=1}^n \sqrt{H_i s_i} \right)^2. \tag{9}$$

Here C^{LB} is a lower bound on the average cost of producing n different items in a single machine.

We are interested in contrasting this lower bound cost with that of the average cost of producing each item in a dedicated machine (DM), where item $i = 1, \dots, n$ is assumed to be produced at rate μ_i . It is easy to see that in the absence of setup costs, is optimal to produce item i every $T_i = s_i/(1 - \rho_i)$ units of time, resulting in a total average cost of

$$C^{DM} = \sum_{i=1}^n \frac{H_i s_i}{1 - \rho_i}. \tag{10}$$

Now, let us compare the average cost C^{DM} of producing the items on dedicated machines, and the lower bound cost C^{LB} of producing the items in a single machine. It is easy to see that $C^{DM} < C^{LB}$. The following example illustrates the additional cost of offering variety.

Example Suppose that $H_1 = 1, \lambda_i = 1, s_1 = 1$, and $\mu_i = 10$ for $i = 1, \dots, n$. It is clear that at most nine items can be produced in a single facility, so we assume that $n \leq 9$. Under these conditions, $C^{LB} = 10n^2/(10 - n)$, $C^{DM} = 10n/9$, and

$$\frac{C^{LB}}{C^{DM}} = \frac{9n}{10 - n}.$$

Notice that the ratio indicates how much more we pay in terms of holding costs to offer a larger variety. This ratio is one when n is one, and it grows to eighty-one when n is nine.

1.5.2 Reducing Setup Times

An important concern in modern manufacturing is that of reducing setup times. Setup times are often reduced in pursuit of the benefits of just-in-time production, but setup time reductions can be justified directly in terms of reduced costs. Setup times are typically reduced either by externalizing setup operations, or by making direct investments to reduce the setup times of operations that are difficult to do off-line.

Out-of-pocket setup costs often *increase* when setup operations are externalize because, when done off-line, these operations require more time, additional or better trained workers, or more careful coordination by management. Nevertheless, externalizing setups can significantly reduce the long run average cost in facilities where lot sizes are mainly driven by the capacity constraint. These explains why some Japanese companies have been willing to spend more on setup costs to reduce internal setup times in order to reduce lot sizes and average cost. In effect they are trading setup times for setup costs. It can be shown that when out-of-pocket setup costs are zero, it is best to reduce the setup time of the item that is produced more frequently, i.e., the one with the smallest T_i^{LB} . The idea is that we enjoy the time savings every time the item is setup.

2 The Joint Replenishment Problem

The Joint Replenishment Problem (JRP) arises in a manufacturing setting when a machine requires a major setup to produce a set of products and a minor setup for each item included in the set. For example, the major setup may consist on placing the die into the machine and to adjust it to get good parts, while the minor setups may consist in opening and closing cavities in the die to produce different variants of the product. The JRP also arises in a distribution setting where multiple items can be shipped together, the cost per shipment is fixed, and there are in addition fixed, item dependent, costs for picking and processing. Notice that an order with a minor setup cost cannot be placed unless the major setup cost is also incurred. However, once the major setup cost is incurred then any item can be ordered by simply incurring its minor setup cost.

The relevant questions are: What is the optimal time between major setups? What is the optimal time between setups for each item?

When the major setup cost is zero the problem reduces to the case of independent items studied before. When the major setup cost is very large then all the items will be forced to order together. Thus, a manager that insists on ordering all items infrequently and together may be subconsciously charging the system a large major setup cost for his time. In most cases the major setup cost is large enough to force us to consolidate some, but not all, orders.

Let

- λ_i demand rate for item $i = 1, \dots, n$,
- h_i holding cost rate for item $i = 1, \dots, n$,
- K_i fixed minor setup cost for item $i = 1, \dots, n$,
- K_0 fixed major setup cost.

For convenience, we set $H_i = 0.5h_i\lambda_i$.

As stated, the solution can be exceedingly complicated because orders need to be coordinated to account for the major setup cost. Rather than studying this more general problem, we will restrict attention to a type of policy that seems to be quite constraining, but in fact is not. Specifically, we assume that there exists a base planning period β (day, shift, week, or month), where β is expressed in years, and that major setups occur as a non-negative integer power-of-two multiple of β . Finally, we assume that the policy that is followed is stationary. That is, the only solutions that will be considered assume that the time between setups, either major or minor, is always the same.

It helps to envision the solution before we delve into the details. There will be a set of items $C \subset \{1, \dots, n\}$ that will order together and will pay for the major setup cost. You can think of C as the set of items that want to order frequently. The other items will then order only when the set C orders. Thus, for example, if the set of items C orders every month, then other items will order either every month, every two months, every four months, etcetera.

The following mathematical program follows from minimizing average cost under the above assumptions.

$$\begin{aligned}
Z_{PT} &= \min \sum_{i=0}^n \left(H_i T_i + \frac{K_i}{T_i} \right) \\
\text{subject to} \quad & T_i = M_i \beta \\
& M_i \geq M_0 \\
& M_i \in \{2^l : l = 0, 1, \dots, \},
\end{aligned}$$

where $H_0 \equiv 0$, and T_i is the reorder interval of item $i = 1, \dots, n$, and T_0 is the interval between major setups.

Under formulation Z_{PT} the reorder intervals are powers-of-two multiples of a base planning period and every item has a reorder interval that is at least as large as the interval as that between consecutive payments of the major setup cost.

As stated the above problem is still difficult to solve because it is a non-linear program with integer constraints. Instead of solving Z_{PT} to optimality we will use a heuristic that will first relax the problem by ignoring the integer constraints and then round the solution to powers-of-two.

The first step of the heuristic is to relax the integrality constraint. This leads to the program

$$\begin{aligned}
Z_R &= \min \sum_{i=0}^n \left(H_i T_i + \frac{K_i}{T_i} \right) \\
\text{subject to} \quad & T_i \geq T_0 \geq 0.
\end{aligned}$$

Notice that $Z_R \leq Z_{PT}$ because any feasible solution to problem Z_{PT} is also feasible for problem Z_R . It can be shown that Z_R is a *lower bound* among all feasible policies (not only stationary policies) and that there exists a power-of-two solution whose cost is at most $1.06Z_R$, see [4].

2.1 Algorithm to Solve the Relaxed Problem

Step 1 Sort the items so $\frac{K_1}{H_1} \leq \dots \leq \frac{K_n}{H_n}$.

Step 2 Let $C = \{1, \dots, i^*\}$ where i^* is the largest index for which

$$\frac{\sum_{j=0}^i K_j}{\sum_{j=0}^i H_j} \geq \frac{K_i}{H_i},$$

and let $K(C) = \sum_{i=0}^{i^*} K_i$, $H(C) = \sum_{i=1}^{i^*} H_i$.

Step 3 Set $T_0^R = T_1^R = \dots = T_{i^*}^R = \sqrt{\frac{K(C)}{H(C)}} \equiv T^R(C)$

and $T_i^R = \sqrt{\frac{K_i}{H_i}}$ $i = i^* + 1, \dots, n$.

Intuitively the algorithm works as follows. After sorting the items, the natural order intervals $\sqrt{K_i/H_i}$ $i = 1, 2, \dots, n$ are non-decreasing. This means that in the absence of the major setup cost item 1 would like to order more frequently than item 2 etcetera. Set $C = \{1\}$ and suppose item 1 absorbs the major setup cost by itself. Then item 1's order interval is $T(C) = \sqrt{K(C)/H(C)} = \sqrt{(K_0 + K_1)/H_1}$. Now, if

$$T(C) > \sqrt{\frac{K_2}{H_2}},$$

then it is optimal to produce item 2 every time item 1 is produced and so it should share the burden of the major setup cost. We add item 2 to the set of items that have a common order interval, i.e., $C = \{1, 2\}$ and continue adding items to C until there is an item whose natural order interval is greater than $T(C)$.

2.2 Finding a Power-of-Two Solution from the Relaxation

Given $T^R = (T_1^R, \dots, T_n^R)$ with $T_i^R = T^R(C) = \sqrt{K(C)/H(C)}$ for $i \in C$ and $T_i^R = \sqrt{K_i/H_i}$ for $i \notin C$ we want to find a power-of-two policy based on a planning period β , i.e., we want to find order intervals of the form $2^k\beta$ where $k \in \{0, 1, 2, \dots\}$. We assume β is sufficiently small so that $\beta \leq T^R(C)$. We know that for the group of items in C any order interval in $[T^R(C)/\sqrt{2}, \sqrt{2} T^R(C)]$ will have a cost that is at most 1.06 times the cost $2\sqrt{K(C)H(C)}$ of using order interval $T^R(C)$. Therefore it is enough to find the smallest integer, say x_C such that $2^{x_C}\beta \in [T^R(C)/\sqrt{2}, \sqrt{2} T^R(C)]$. Similarly, for $i \notin C$ any order interval in $[T_i^R/\sqrt{2}, \sqrt{2} T_i^R]$ will have a cost that is at most 1.06 times the cost $2\sqrt{K_i H_i}$ of using the order interval T_i^R . Consequently, it is enough to find the smallest integer, say x_i , such that $2^{x_i}\beta \in [T_i^R/\sqrt{2}, \sqrt{2} T_i^R]$. Let

$$T^{PT}(C) = 2^{x_C}\beta \quad \text{for } i \in C$$

and

$$T_i^{PT} = 2^{x_i}\beta \quad \text{for } i \notin C.$$

Notice that by our choice of the x_i 's we have

$$\frac{1}{\sqrt{2}} \leq \frac{T_i^{PT}}{T_i^R} \leq \sqrt{2}$$

for all i , so the average cost based on the T^{PT} is at most 6% higher than the average cost based on T^R .

2.3 Numerical Example

Consider the data from the table below. We start by solving the relaxation.

Step 1 The items are already sorted as needed.

Step 2 Clearly $1 \in C$. Since $8/2 > 2$ we have $2 \in C$. Since $12/3 = 4$ we have $3 \in C$. Since $18/4 < 6$ we stop and $C = \{1, 2, 3\}$.

Step 3 Set $T_0^R = T_1^R = T_3^R = \sqrt{12/3} = 2$, $T_4^R = \sqrt{6}$ and $T_5^R = 4$.

We now round the solution to powers-of-two. Clearly $x_C = 7$ is the smallest integer such that $2^{x_C} \frac{1}{52} \geq T^R(C)/\sqrt{2} = \sqrt{2}$. Similarly, $x_4 = 7$ is the smallest integer such that $2^{x_4} \frac{1}{52} \geq T_4^R/\sqrt{2} = \sqrt{3}$, and $x_5 = 8$ is the smallest integer such that $2^{x_5} \frac{1}{52} \geq T_5^R/\sqrt{2} = 2\sqrt{2}$.

Thus, we have

$$T_i^{PT} = 2^7\beta \quad \text{for } i = 1, 2, 3, 4$$

and

$$T_i^{PT} = 2^8\beta \quad \text{for } i = 5.$$

In words, items 1, 2, 3 and 4 order every $2^7 = 128$ weeks while item 5 orders every $2^8 = 256$ weeks. The cost penalties are as follows: For items 1, 2, and 3 the cost penalty is 2.16%. For item 4 the cost penalty is less than 1%. For item 5 the cost penalty is 2.16%.

References

- [1] Dobson, G. (1987) "The Economic Lot Scheduling Problem: Achieving Feasibility Using Time-Varying Lot Sizes." *Operations Research*, 35, 764-771.
- [2] Gallego, G. (1990) "Scheduling the Production of Several Items with Random Demands in a Single Facility," *Management Science*, 36, 1579-1592.

$$\beta = 1/52$$

item	K_i	H_i	K_i/H_i
0	5	0	
1	1	1	1
2	2	1	2
3	4	1	4
4	6	1	6
5	16	1	16

Table 2: Data for the Joint Replenishment Problem

- [3] Gallego, G., and D. Shaw. 1997. On The Complexity of the Economic Lot Scheduling Problem with General Cyclic Schedules. *IIE Transactions*. **29**, 109-113.
- [4] Jackson, P., Maxwell, W., and Muckstadt J. (1985). *The Joint Replenishment Problem with a Powers-of-Two Restriction*, IIE-Transactions, 25-32.