Management of International Business and Economic Systems Conference 2009 (MIBES '09), Florina, Greece, 18-20 September 2009.

The Economic Lot Sizing Problem for Continuous Multi-grade Production with Stochastic demands

Olympia Hatzikonstantinou
Department of Mechanical Engineering
University of Thessaly, Pedion Areos, Volos 38334, Greece
ohatzikon@mie.uth.gr

George Liberopoulos
Department of Mechanical Engineering
University of Thessaly, Pedion Areos, Volos 38334, Greece
glib@mie.uth.gr

Dimitrios Pandelis
Department of Mechanical Engineering
University of Thessaly, Pedion Areos, Volos 38334, Greece
d pandelis@mie.uth.gr

Abstract
We study a variant of the Stochastic Economic Lot Scheduling Problem (SELSP) in which a single production facility must produce several grades to meet random stationary demand for each grade from a common finished-goods (FG) inventory buffer with limited storage capacity. Demand that can not be satisfied directly from inventory is lost. Raw material is always available, and the production facility produces at a constant rate. When the facility is set up to produce a particular grade, the only allowable changeovers are from that grade to next lower or higher grade. All changeover times are deterministic and equal to each other. There is a changeover cost per changeover occasion, a spill-over cost per unit of product in excess, whenever there is not enough space in the FG buffer to store the produced grade, and a lost-sales cost per unit short, whenever there is not enough FG inventory to satisfy demand. We model the SELSP as a discrete-time Markov Decision Process (MDP), where in each time period we must decide whether to initiate a changeover to a neighboring grade or keep the setup of the production facility unchanged, based on the current state of the system, which is determined by the current setup of the facility and the FG inventory levels of all the grades. The goal is to minimize the infinite-horizon long-run average cost. For 2 and 3-grade problems we can numerically solve the resulting MDP problem using successive approximation. For problems with more than three grades, we develop a heuristic solution which is based on approximating the original multi-grade problem into many 3-grade sub-problems and numerically solving each sub-problem using successive approximation. We present and discuss numerical results for incidences with 2, 4 and 5 grades, using both the exact numerical and the heuristic solution procedure.

Keywords: Stochastic economic lot sizing problem, Dynamic scheduling, Process industry, Markov decision process

JEL classifications: Industry Studies, Manufacturing, L60, L65.

Introduction

Scheduling production of multiple products, each with random demand, on a single facility with limited production capacity and significant changeover costs and times between products is a classic problem in production planning research that is often referred to as the Stochastic Lot Scheduling Problem (SLSP). Sox et al. (1999) distinguishes between two versions of the SLSP: the Stochastic Economic Lot Scheduling Problem (SELSP) and the Stochastic Capacitated Lot Sizing Problem (SCLSP), for consistency with the deterministic demand literature. The SELSP assumes an infinite planning horizon and stationary demand, whereas the SCLSP assumes a finite planning horizon and allows for non-stationary demand. The SELSP is better suited for continuous-processing manufacturing, whereas the SCLSP is more appropriate for discrete-parts manufacturing. Discrete-parts manufacturing is characterized by individual parts that are clearly distinguishable and is often encountered in the industries of computer and electronic products, electrical equipment and appliances, transport equipment, machinery, fabricated metal, wood, furniture products, etc. Process industries, on the other hand, operate on material that is continually flowing, as is the case with petroleum and coal products, metallurgical products, non-metallic mineral products, food and beverage products, paper products, etc. Generally, process industries are capital intensive and focus on high-volume, low-variety production. In a typical process industry, the production facility operates continuously, and the different products are really variants of the same family that differ in one or more attributes, such as grade, quality, size, thickness, etc. Often, the different grades are related in such a way that the only allowable changeovers are from one grade to the next higher or lower grade in the chain. For example, if the facility produces three grades, A, B, and C – A being the lowest and C being the highest – the allowable changeovers are between A and B, between B and C, but not directly between A and C.
The deterministic version of the SELSP, the so-called ELSP has received considerable attention (e.g., see the surveys of Elmaghraby, 1978 and Salomon, 1991). Both analytical and heuristic solutions for the ELSP derive rigid cyclic production plans, which in many multi-grade plants take the form of rigid product slates or wheels, whereby all grades are produced sequentially in a cycle, starting from the lowest grade, going up all the way to the highest grade, and returning down to the lowest grade. In the previous example with the three grades, a complete product grade slate would be A-B-C-B-A. Unfortunately, cyclic plans do not work well for the stochastic problem, for two reasons. They focus on lot-sizing and not on dynamic capacity allocation and the inventories of finished products serve not only to reduce the number of changeovers but also to hedge against stock-outs. In the stochastic problem, both lot-sizing and capacity allocation have to be considered simultaneously, and the dynamics have to be included in the plan (Graves, 1980).

In this paper, we study a variant of the SELSP in which a single production facility must produce several grades to meet random stationary demand for each grade from a common finished-goods (FG) inventory buffer with limited storage capacity. Demand that can not be satisfied directly from stock is lost. Raw material is always available, and the production facility produces at a constant rate all the time. When the facility is set up to produce a particular grade, the only allowable changeovers are from that grade to next lower or higher grade. In many industries, it is customary to divide the intermediate grade produced during a changeover, say from grade A to grade B, into two halves, and classify the first half as A and the second half as B, although in reality the grade of the product coming out of the production facility may gradually change from grade A to grade B. In this paper, for simplicity, we assume that the grade produced during a changeover from A to B is classified as A, and the grade produced during the reverse changeover, from B to A, is classified as B. Under this assumption, the amounts of grades A and B that will be produced in the long run will be the same as those that would have been produced had we divided the produced grade during a changeover into two halves. We also assume that all changeover times are deterministic and equal to each other. The cost structure of our model includes a changeover cost per changeover occasion, a spill-over cost per unit of product in excess, whenever there is not enough space in the FG buffer to store the produced grade, and a lost-sales cost per unit short, whenever there is not enough FG inventory to satisfy demand. The assumptions presented above are realistic and are based on a real dynamic scheduling problem of a PET processing plant, presented in Liberopoulos et al. (2009).

We model the SELSP problem described above as a discrete-time Markov Decision Process (MDP), where in each time period the decision is whether to initiate a changeover to a neighboring grade or keep the setup of the facility unchanged, based on the current state of the system, which is determined by the current setup and the FG inventory levels of all the grades. The goal is to minimize the infinite-horizon long-run average cost.

Because of its theoretical and practical importance, the SELSP problem has received considerable attention in the literature. A comprehensive review of related works can be found in Sox et al. (1999) and Winands et al. (2005). From these reviews, it is apparent that there have been two approaches to the SELSP. One approach is to develop a cyclic schedule using a deterministic approximation of the stochastic problem and develop a control rule for the stochastic problem to pursue this schedule. The other approach, which we follow in this paper, is to develop a dynamic schedule that determines which product to produce based on the current state of the system.

One of the first papers that looked at the SELSP as a discrete-time stochastic control problem with dynamic sequencing is Graves (1980). Graves first solves a one-product problem with inventory-backorder costs and changeover costs, but no changeover times, where the decision in each period is to produce or idle the facility. He then uses the solution of the one-product problem as the basis for a heuristic procedure to solve the multi-product problem. In that heuristic, scheduling conflicts among different products are solved by comparing the value functions derived for each individual and “composite” product from the one-product analysis. The composite product is a concept that Graves introduces to help anticipate possible scheduling conflicts in the multi-product problem. The idea is that the composite inventory of several products should indicate the need for current production, in case the individual product inventories are deemed just adequate when viewed separately.

Qiu and Loulou (1995) look at a problem with Poisson demand, deterministic processing and changeover times, and changeover and inventory-backlog costs. They model the problem as a semi-Markov decision process, where the objective is to decide in each review epoch which product, if any, to set up the facility to produce, in order to minimize the infinite-horizon, discounted cost. The review epochs are those points in time when either the production facility is idle and some demand arrives, or when a part has just been processed and the production facility is free. They use successive approximation to generate near-optimal control policies by solving the problem on a truncated inventory space, and compute error bounds caused by the truncation. They present numerical results for two-product problems, and state that systems with more than two products are limited by the curse of dimensionality.

Leachman and Gascon (1988) develop a dynamic, periodic review control policy that determines which products to produce and how much, based on solutions of deterministic ELSP that account for non stationary demand. This solution is modified if two or more products are close to being stocked out or are backordered.

Finally, Sox and Muckstadt (1997) and Karmarkar and Yoo (1994) develop finite-horizon stochastic mathematical programming models for the SELSP, that can also be classified as SCLSP, with deterministic production and changeover times, and use Lagrangian relaxation for finding optimal or near-optimal solutions for problems of small sizes.
Our work in this paper follows the stream of papers that view the SELSP as a discrete-time periodic-review control problem with dynamic production sequencing and global lot sizing, and is most closely related to Graves (1980) and Qiu and Loulou (1995). It differs from previous works in that it considers a new variant of the SELSP, where the only allowable changeovers are from one grade to the next lower of higher grade. The latter feature renders problems with a large number of grades amenable to heuristic solution procedures that are based on approximating the original problem by many smaller (i.e., with fewer grades) sub-problems that are computationally easier to solve. Thus, for two-grade and three-grade problems we are able to numerically solve the resulting MDP problem using successive approximation, and obtain insight into the optimal control policy. For problems with N grades, where N > 3, we develop a heuristic solution which is based on decomposing the original N-grade problem into (N – 2) 3-grade sub-problems and numerically solving each sub-problem using successive approximation. Each 3-grade sub-problem is an approximation of the original N-grade problem, where the middle grade in the sub-problem corresponds to one of the grades in the original problem, the low (left) grade in the sub-problem is the composite of all grades in the original problem that are lower than the middle grade, and the high (right) grade is the composite of all grades that are higher than the middle grade. For example, if the original problem consists of five grades, A-B-C-D-E, we formulate the following 3-grade sub-problems: A-B-(C+D+E), (A+B)-C-(D+E), and (A+B+C)-D-E, where the notation "(A+B)" indicates the composite grade formed by grades A and B. After solving all the sub-problems, the heuristic control policy for the original N-grade problem is obtained by combining parts of the optimal policies of the sub-problems.

The rest of this paper is organized as follows. In Section 0, we present the stochastic dynamic programming formulation and solution of the MDP model of the original N-grade problem. The heuristic procedure for solving problems with more than three grades is outlined in Section 0. Finally, numerical results for problem incidences with 2, 4 and 5 grades, using both the exact numerical and the heuristic solution procedure are presented in Section 0, and conclusions are drawn in Section 0.

Problem Formulation and Dynamic Programming Solution

We consider a discrete-time model of a production facility that can produce N different grades, one at a time. Grade changeovers are only allowed between neighboring grades, n and n + 1, n = 1, …, N − 1. The changeover time is one period. In each time period, the production facility produces P units of the grade that is was set up for at the beginning of the period. The quantity produced is stored in a common FG buffer with a finite storage capacity of X units; any excess amount that does not fit in the buffer is spilled over, incurring a spill-over cost of CS per unit of excess product. The FG buffer is flexible in that it can contain any quantity of any grade at the same time, as long as the total amount does not exceed X. After the quantity produced by the facility has been added to the FG buffer, a vector of random demands $\mathbf{D} \equiv (D_1, \ldots, D_N)$ must be met from FG inventory. The demand for grade n, D_n, n = 1, …, N, is a discrete random variable with known stationary joint probability distribution. For each grade n, the part of the demand that can not be satisfied from FG inventory, if any, is lost, incurring a lost-sales cost of CL_n per unit of unsatisfied demand. In many real problems, P is not considered as a control variable for scheduling purposes, because changing P may cause instabilities in the production process. In this paper, we assume that P is fixed and equal to (or close to) the total expected demand for all grades.

We formulate the dynamic scheduling problem of the production facility as a discrete-time MDP, where the state of the system at the beginning of each period is defined by the vector $\mathbf{y} \equiv (s, x_1, \ldots, x_N)$, where s is the grade that the facility is set up for during that period and x_n, n = 1, …, N, is the FG inventory level of grade n at the beginning of the period. Note that $s \in \{1, \ldots, N\}$, and the set of allowable inventory levels is determined by all integers x_n, n = 1, …, N, such that $0 \le \Sigma_n x_n \le X$. Thus, the size of the state space is $½ \cdot N \cdot X^N$. The decision, u, to be made at the beginning of each period is whether to initiate a changeover to a neighboring grade or leave the facility setup unchanged. Thus, if the current setup is s, the allowable decisions are given by the set U(s), where U(1) = {1, 2}, U(N) = {N − 1, N}, and U(s) = {s − 1, s, s + 1}, s = 2, …, N − 1. If the decision is to initiate a changeover, then this changeover will be in effect at the beginning of the next period. A decision to initiate a changeover at the beginning of a period incurs a changeover cost CC in that period. Suppose that the state of the system at the beginning of a period is $\mathbf{y}$, decision u is taken, and demand $\mathbf{D}$ is realized. Let $g(\mathbf{y},u,\mathbf{D})$ be the cost incurred during that period and let $\mathbf{y}' \equiv (s', x_1', \ldots, x_N') = f(\mathbf{y},u,\mathbf{D})$ be the state of the system at the beginning of the next period. From the above discussion, it is clear that $s' = u$ and $x_n' = [x_n + p(\mathbf{y}) \cdot I_{n=s} - D_n]^+$, n = 1, …, N, where $p(\mathbf{y})$ is the amount added to the FG buffer after the facility produces P units and before the demand is satisfied and is given by $p(\mathbf{y}) \equiv \min\{P, X - \Sigma_n x_n\}$, I_a is the indicator function which takes the value of 1 if a is true, and 0 otherwise, and $[x]^+ \equiv \max\{0, x\}$. Moreover, $g(\mathbf{y},u,\mathbf{D}) = CC \cdot I_{u \neq s} + CS \cdot (P - p(\mathbf{y})) + \Sigma_n CL_n \cdot [D_n - x_n - p(\mathbf{y}) \cdot I_{n=s}]^+$.

The objective is to find a state dependent policy $u = \mu(\mathbf{y})$ that minimizes the long-run (infinite-horizon) expected average cost per period, denoted by J. To find such a policy we need to solve the so-called Bellman equation, which can be written as $J + V(\mathbf{y}) = \min_{u\in U(s)} T_u(V(\mathbf{y}))$, where $V(\mathbf{y})$ is the differential cost starting from state $\mathbf{y}$, and the operator $T_u(\cdot)$ is defined as $T_u(V(\mathbf{y})) \equiv E_D\{ g(\mathbf{y},u,\mathbf{D}) + V(\mathbf{y}')\}$. The minimizer in the Bellman equation determines the optimal policy when the system is in state $\mathbf{y}$, denoted by $\mu^*(\mathbf{y})$.

We solve the Bellman equation by the method of successive approximations. We denote by $V_k(\mathbf{y})$ the value of the differential cost function at the kth iteration. Initially, we set $V_0(\mathbf{y}) = 0 \; \forall \; \mathbf{y}$. The values at the (k + 1)th iteration are obtained from the previous iteration by the recursion $V_{k+1}(\mathbf{y}) = T(V_k(\mathbf{y})) - T(V_k(\hat{\mathbf{y}}))$, where $T(V_k(\mathbf{y})) = \min_{u\in U(s)} T_u(V_k(\mathbf{y}))$ and $\hat{\mathbf{y}}$ is an arbitrarily chosen special state. Note that in each iteration the differential cost for the special state is reset to zero. Assuming that the iteration scheme converges to some values $V(\mathbf{y})$, then from the recursion equation, these values must satisfy $T(V(\hat{\mathbf{y}})) + V(\mathbf{y}) = T(V(\mathbf{y}))$. A comparison of this equation and the Bellman equation reveals that $J = T(V(\hat{\mathbf{y}}))$.

To implement the successive approximation method, at each iteration k = 1, 2, … we compute the maximum and minimum differences, $V_k^U = \max_{\mathbf{y}}\{V_k(\mathbf{y}) - V_{k-1}(\mathbf{y})\}$ and $V_k^L = \min_{\mathbf{y}}\{V_k(\mathbf{y}) - V_{k-1}(\mathbf{y})\}$. The procedure is terminated when $V_k^U - V_k^L < \varepsilon \cdot T(V_k(\hat{\mathbf{y}}))$, where ε is some small positive scalar.

Heuristic Solution

Although the exact method presented in the preceding section can in principle determine the optimal policy for any number of grades, it becomes computationally intractable for more than three grades. In this section, we present a heuristic procedure that approximates any N-grade problem, N > 3, by several 3-grade sub-problems and then uses the sub-problem solutions (determined by the exact numerical method) to construct a heuristic policy for the original problem. More specifically, the heuristic procedure works as follows. Let S denote the original N-grade problem. Then, for each grade n, n = 2, …, N − 1, we formulate a 3-grade sub-problem, denoted by S_n, in which the middle grade is grade n, the low grade is the composite of all grades that are lower than n, i.e., grades 1, …, n − 1, and the high grade is the composite of all grades that are higher than n, i.e., grades n + 1, …, N; hence S_n is an approximation of the original problem S. For each sub-problem S_n, we define the state of the system by the vector $\mathbf{y}_n = (s_n, w_n, x_n, z_n)$, where $s_n \in \{1, 2, 3\}$ and w_n and z_n are the total inventory levels of the low and high composite grades, respectively, and are given by $w_n \equiv x_1 + \ldots + x_{n-1}$ and $z_n \equiv x_{n+1} + \ldots + x_N$. In each sub-problem S_n, the demand distribution of the middle grade is the same as the demand distribution of grade n in the original problem, the demand distribution of the low grade is the convolution of the demand distributions of grades 1, …, n − 1 in the original problem, and the demand distribution of the high grade is the convolution of the demand distributions of grades n + 1, …, N in the original problem. We use the exact method to obtain the optimal policy $\mu_n^*(\mathbf{y}_n)$ of sub-problem S_n. The heuristic policy for the N-grade problem, denoted by $\mu(\mathbf{y})$, is then constructed by using parts of the optimal policies of the sub-problems, as follows: $\mu(1, x_1, \ldots, x_N) = \mu_2^*(1, \hat{w}_2, x_2, \check{z}_2)$, $\mu(N, x_1, \ldots,$

$x_N) = \mu_{N-1}^*(3, \hat{w}_{N-1}, x_{N-1}, \check{z}_{N-1})$, and $\mu(n, x_1, \ldots, x_N) = \mu_n^*(2, \hat{w}_n, x_n, \check{z}_n)$, $n = 2, \ldots, N-1$, where $\hat{w}_n$ and $\check{z}_n$ are the "aggregate" inventory levels of the low and high composite grades, respectively, which represent in some aggregate way their individual components and hence are given by $\hat{w}_n = h(x_1, \ldots, x_{n-1})$ and $\check{z}_n = h(x_{n+1}, \ldots, x_n)$ for some appropriate function h, which will be defined next.

First, note that $\hat{w}_2 = x_1$ and $\check{z}_{N-1} = x_N$, because in these cases the composite grade corresponds to a single grade. We now focus on $\hat{w}_n$, $n > 2$, as $\check{z}_n$ is obtained in exactly the same way. An obvious choice for the aggregate inventory level of the composite grade made up of grades 1, …, n − 1 is the sum of the inventory levels of the individual grades, i.e., $\hat{w}_n = w_n$. This is a reasonable choice, especially with respect to estimating potential spill-over costs, but may underestimate the possibility of lost sales when one or more of the individual components of the composite grade is very low compared to the others. To illustrate this, suppose that the facility is currently set up to produce grade 4, and that the inventory levels of grades 1-4 are $x_1 = x_2 = 15$, $x_3 = 0$, and $x_4 = 6$. Then, in sub-problem S_4, the inventory level of the middle grade would be $x_4 = 6$, and the total inventory level of the low composite grade would be $w_4 = x_1 + x_2 + x_3 = 30$. In this case, the optimal policy obtained by solving S_4 might indicate that it is optimal for the facility not to changeover to the low composite grade, because there is enough of it (30 units) in storage compared to the inventory level of the middle grade 4, which is much lower (6 units). What the heuristic fails to see in this case is that although the sum of the inventory levels that make up the composite grade is relatively high, one of its components, namely x_3, is zero, and unless the facility changes over to grade 3, a heavy stock-out penalty is likely to be incurred in the current and in the following period.

To take into account such a situation, we seek an aggregate inventory level, $\hat{w}_n$, for the composite grade made up of grades 1, …, n − 1 that would result in the same value of the expected lost sales for that composite grade as that computed by summing the expected lost sales of the individual component grades of the composite. The sum of the expected lost sales of the individual grades is given by $LS = E[D_1 - x_1]^+ + \ldots + E[D_{n-1} - x_{n-1}]^+$. The expected lost sales for a given inventory level w of the composite grade is equal to $E[(D_1 + \ldots + D_{n-1}) - w]^+$; therefore, $\hat{w}_n$ is the value of w that makes the latter expression as close as possible to LS. To compute this expression we first need to derive the probability distribution of the aggregate grade demand by convolution of the probability distributions of individual grade demands. In case this is not computationally convenient we propose the following faster alternative.

We approximate the sum of the expected lost sales of the individual grades by $LS = [E(D_1) - x_1]^+ + \ldots + [E(D_{n-1}) - x_{n-1}]^+$. If all inventory levels x_i are large enough so that LS = 0, we set $\hat{w}_n = w_n$. Otherwise, we define $\hat{e}_n = [E(D_1) + \ldots + E(D_{n-1})] - LS$, and we set $\hat{w}_n$ to be a linear combination of $\hat{e}_n$ and w_n, i.e., $\hat{w}_n = \alpha\hat{e}_n + (1 - \alpha)w_n$, rounded to the nearest integer, for some $0 \le \alpha \le 1$.

Numerical Results

In this section, we present numerical results for problem examples with 2, 4 and 5 grades, using both the exact numerical and the heuristic solution procedure presented in the previous sections. First, we look at a 2-grade example (N = 2), where P = 5 and the demand distribution for the two grades is given in Table **1**.

Table 1: Demand distribution of the 2-grade example

	$\Pr(D_n=i)$							
$n \setminus i$	0	1	2	3	4	5	6	$E(D_n)$
1	0.1	0.15	0.15	0.2	0.15	0.15	0.1	3
2	0.15	0.15	0.4	0.15	0.15	0	0	2

Table **2** shows the number of iterations of the successive approximation procedure until convergence, denoted by k_c, for convergence tolerance criterion $\varepsilon = 0.001$, and the resulting optimal long-run average cost, J, for various combinations of space capacity, X, and cost parameters, where it is assumed that both grades have the same lost-sales cost rate, i.e., $CL_1 = CL_2 = CL$. From the results, is can be seen that as X increases, k_c increases and J decreases, as is expected. J also increases as the cost parameters increase.

Table 2: Results for the 2-grade example

				$X = 40$		$X = 60$		$X = 80$		$X = 100$	
Case	*CC*	*CS*	*CL*	k_c	*J*	k_c	*J*	k_c	*J*	k_c	*J*
1	1	5	5	186	0.9824	474	0.618	895	0.4503	1447	0.354
2	1	10	10	188	1.7454	472	1.0965	891	0.7985	1444	0.6277
3	2	5	5	179	1.1640	448	0.7342	844	0.5354	1367	0.421
4	5	10	1	181	1.6842	437	1.0682	806	0.7806	1286	0.6146
5	5	1	10	211	1.6933	515	1.074	956	0.7848	1538	0.6178
6	2	10	10	186	1.9648	474	1.2361	895	0.9006	1447	0.7079
7	10	1	1	340	1.1409	369	0.7536	408	0.5587	588	0.4445
8	10	5	10	168	2.7141	411	1.7277	761	1.2644	533	0.9962
9	1	10	5	225	1.3610	555	0.855	1032	0.6228	1659	0.4896
10	1	5	10	253	1.3679	632	0.8593	1184	0.626	1908	0.4921

Figure 1 shows the optimal changeover policy as a function of inventories x_1 and x_2, for cases 1 and 3 of Table 2, for X = 40, and is representative of the other cases as well.

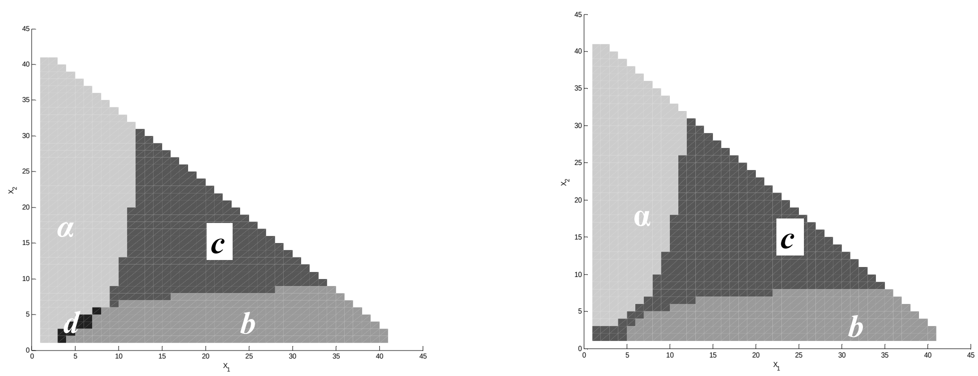

Figure 1: Optimal changeover policy for cases 1 (left) and 3 (right) of Table 2, for X = 40

In both cases 1 and 3, the optimal policy partitions the inventory space in several regions, each characterized by a different optimal changeover action. If we let $\mu^*(s, R)$ denote the optimal policy when the facility is set up to produce grade s and the inventory levels are in region R, then $\mu^*(1, a) = \mu^*(2, a) = 1$, $\mu^*(1, b) = \mu^*(2, b) = 2$, $\mu^*(1, c) = 1$, $\mu^*(2, c) = 2$, $\mu^*(1, d) = 2$, $\mu^*(2, d) = 1$. Thus, the optimal policy dictates the following actions: When the inventory levels are in region a, changeover to produce grade 1, when in b, changeover to produce grade 2, when in c, do not changeover, and when in d, changeover to the other grade. A typical production sequence when the inventory levels are in and around region d would be one where the facility changes over from one grade to the other in each period. When the inventory levels are in region c, the facility would be producing grade 1 in successive periods until the inventory levels cross the border between regions c and b and then changing over to grade 2 and producing that grade until the inventory levels cross the border between regions c and a. Notice that region c is wider in case 3 than in case 1, indicating that in case 3 it is optimal to produce longer runs of the two grades with less frequent changeovers, because the changeover cost in case 3 is twice as much as in case 1. In fact, the widening of region c in case 3 is big enough to eliminate region d. Another observation is that the inventory space partition is more or less symmetric for the two grades but with a slight displacement in favor of grade 1, because grade 1 has a higher demand than grade 2. Next, we look at a 4-grade (N = 4) and a 5-grade (N = 5) example. In each example, we assume that the demand for each grade is identically distributed to one of the random variables D_j, j = A, B, …, E, F, whose distributions are given in Table **3**.

Table 3: Demand distributions for the 4-grade and 5-grade examples

	$\Pr(D_j = i)$				
$J \setminus i$	0	1	2	3	$E(D_j)$
A	0.65	0.25	0.05	0.05	0.5
B	0.4	0.5	0.05	0.05	0.75
C	0.25	0.5	0.25	0	1
D	0.25	0.25	0.5	0	1.25
E	0.25	0.25	0.25	0.25	1.5
F	0.05	0.2	0.45	0.3	2

For each example, we consider four cases, each with a different demand pattern that captures a different way that total demand is distributed among the individual grades. In each case, the total expected demand is equal to the production rate. First, we solve each case optimally by dynamic programming, using $\varepsilon = 0.001$. Then, we solve each case by the heuristic. In the implementation of the heuristic we use the faster alternative to approximate the sum of the expected lost sales of the individual grades, described at the end of Section 0, for values of α ranging from 0 to 1 with a step of 0.1. In all cases we assume that CC = 1, $CS = CL_n = 1$, n = 1,…, 5, and P = 6. The results for the 4-grade example, for X = 30, are shown in Table 4. The notation "F,C,F,C" in column 2 is used to indicate that D_1 is distributed as D_F, D_2 is distributed as D_C, etc. The computational (CPU) times are in hours. For the heuristic, we show the total CPU time it takes to solve the 3-grade problems and generate the heuristic policy, but not the time it takes to evaluate the heuristic policy. The optimal value of α in the heuristic is denoted α^* and the corresponding long-run average cost is $J(\alpha^*)$. The last column shows the percentage cost increase between the heuristic and the optimal policy.

Table 4: Results for the 4-grade example

	Demand	Exact			Heuristic			% cost
Case	pattern	k_c	CPU	J	α^*	CPU	$J(\alpha^*)$	Difference
1	F,C,F,C	187	52.41	1.1835	0.1	0.024	1.3207	11.59
2	F,C,C,F	110	41.84	1.2881	0.1	0.054	1.3139	1.96
3	C,F,F,C	55	21.57	1.0034	0.7	0.024	1.2442	24.00
4	F,F,C,C	156	48.22	1.0927	0.5	0.038	1.2253	12.13

From the results, we observe that cases 1 and 2 have higher expected costs compared to cases 3 and 4. This is because in the latter two cases the grades with the highest demands are adjacent in the sequence of allowed changeover transitions, while in the first two cases any transition between those two grades has to go through other grades, thus incurring higher switching costs. In all cases, except case 3, the heuristic average cost is insensitive to parameter α. Case 3 tends to have lower cost for α between 0.5 and 0.8 and significantly higher cost for α between 0.9 and 1. The cost difference between the heuristic and the exact solution is 1.96% for case 2, where the end grades 1 and 4 have the highest demand, and 24% for case 3, where the middle grades 2 and 3 have the highest demand. The heuristic is between 700 and 2,000 times faster than the exact solution.

The results for the 5-grade example, for X = 20, are shown in Table **5**. Cases 2 and 3 have higher average costs because they require more product switches to move between products with the highest demands. A significant difference with the 4-grade example is that the heuristic average cost seems to be an increasing function of α, which means that the best heuristic policy is obtained when $\hat{w}_n = w_n$. The cost difference between the heuristic and the exact solution is between 10% and 20% in all cases, and the heuristic is between 3,000 and 120,000 faster than the exact solution.

Table 5: Results for the 5-grade example

	Demand	Exact			Heuristic			% cost
Case	pattern	k_c	CPU	J	α^*	CPU	$J(\alpha^*)$	Increase
1	C,C,F,C,C	48	32.27	2.944	0	0.010	3.414	15.96
2	E,D,A,D,E	87	142.77	4.076	0	0.014	4.918	20.66
3	E,B,E,B,E	65	125.51	3.851	0	0.023	4.293	11.48
4	B,D,F,D,B	35	38.05	2.652	0.1	0.008	3.036	14.48
5	F,D,D,B,B	71	78.70	3.002	0.1	0.004	3.451	14.96
6	F,D,B,B,D	129	369.40	3.492	0	0.003	3.876	11.05
7	F,B,D,B,D	129	140.30	3.657	0	0.003	3.935	7.59

Conclusions

We studied a variant of the SELSP in which a single production facility must produce several grades to meet random demand for each grade from a common FG inventory buffer with limited storage capacity. The only allowable changeover of the facility is from one grade to next lower or higher grade. All changeover times are deterministic. We modeled this problem as a discrete-time MDP, where in each time period it must be decided whether to initiate a changeover to a neighboring grade, based on the current state of the system. The goal is to minimize the infinite-horizon long-run average changeover, spill-over and lost-sales cost. For 2-grade and 3-grade problems we proposed to numerically solve the resulting MDP problem using successive approximation. For problems with more than three grades, we developed a heuristic solution which is based on approximating the original multi-grade problem into many 3-grade sub-problems and numerically solving each sub-problem using successive approximation. We presented numerical results for problem incidences with 2, 4 and 5 grades, using both the exact numerical and the heuristic solution procedure. For the 4 and 5-grade examples, the cost difference between heuristic and exact solution was as small as 1.96% and as high as 24%. The main advantage of the heuristic is that it was between 700 and 120,000 times faster than the exact solution.

Acknowledgements

This work was supported by grant "03ED913: Optimization of production planning and grade distribution of a PET resin chemical plant," within the Reinforcement Program of Human Research Manpower. It was co-financed by Greece's General Secretariat of Research and Technology (17%), the European Social Fund (68%), and Artenius Hellas S.A. PET Industry (15%).

References

Elmaghraby, *S.E., 1978, "The economic lot scheduling problem (ELSP): Review and extensions," Management Science, 24(6), 587-598.*

Graves, *S.C., 1980, "The multi-product production cycling problem," AIIE Transactions, 12(3), 233-240.*

Karmarkar, *U.S., J. Yoo., 1994, "The stochastic dynamic product cycling problem," European Journal of Operational Research, 73, 360-373.*

Liberopoulos, G., G. Kozanidis, O. Hatzikonstantinou, 2009, "Production scheduling of a multi-grade PET resin production plant," Computers and Chemical Engineering, (in press: doi:10.1016/j.compchemeng.2009.05.017).

Leachman, *R.C., A. Gascon., 1988, "A heuristic scheduling policy for multi-item, single-machine production systems with time-varying, stochastic demands," Management Science, 34(3), 377-390.*

Salomon, *M., 1991, "Deterministic lotsizing models for production planning," Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin.*

Sox, C.R., P.L. Jackson, A. Bowman, J.A. Muckstadt, 1999, "A review of the stochastic lot scheduling problem," International Journal of Production Economics, 62(3), 181-200.

Sox, *C.R., J.A. Muckstadt., 1997, "Optimization-based planning for the stochastic lot-sizing problem," IIE Transactions, 29(5), 349-357.*

Qiu, J., R. Loulou., 1995, "Multiproduct production/inventory control under random demands," IEEE Transactions on Automatic Control, 40(2), 350-356.

Winands, *E.M.M., I.J.B.F. Adan, G.J. van Houtum, 2005, "The stochastic economic lot scheduling problem: A survey," Working paper. Beta Research School for Operations Management and Logistics, Technical University of Eindhoven.*