This is a post-peer-review, pre-copyedit version of an article published in *Annals of Operations Research*. The final authenticated version is available online at: https://doi.org/10.1007/s10479-013-1323-6.

The Stochastic Economic Lot Sizing Problem for Non-Stop Multi-Grade Production with Sequence-Restricted Setup Changeovers

George Liberopoulos, Dimitrios G. Pandelis, Olympia Hatzikonstantinou

Department of Mechanical Engineering, University of Thessaly, Greece
glib@mie.uth.gr; d_pandelis@mie.uth.gr; ohatzikon@mie.uth.gr

January 2013

We study a variant of the *stochastic economic lot scheduling problem* (SELSP) encountered in process industries, in which a single production facility must produce several different grades of a family of products to meet random stationary demand for each grade from a common *finished-goods* (FG) inventory buffer that has limited storage capacity. When the facility is set up to produce a particular grade, the only allowable changeovers are from that grade to the next lower or higher grade. Raw material is always available, and the production facility produces continuously at a constant rate even during changeover transitions. All changeover times are constant and equal to each other, and demand that cannot be satisfied directly from inventory is lost. 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 the demand. We model the SELSP 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 set up of the production facility unchanged, based on the current state of the system, which is defined by the current set up of the facility and the FG inventory levels of all the grades. The goal is to minimize the (long-run) expected average cost per period. For problems with more than three grades, we develop a heuristic solution procedure which is based on decomposing the original multi-grade problem into several 3-grade MDP sub-problems, numerically solving each sub-problem using value iteration, and constructing the final policy for the original problem by combining parts of the optimal policies of the sub-problems. We present numerical results for problem examples with 2-5 grades. For the 2- and 3-grade examples, we numerically solve the exact MDP problem using value iteration to obtain insights into the structure of the optimal changeover policy. For the 4- and 5-grade examples, we compare the performance of the decomposition-based heuristic (DBH) solution procedure against that obtained by numerically solving the exact problem. We also compare the performance of the DBH method against the performance of three simpler parameterized heuristics. Finally, we compare the performance of the DBH and the exact solution procedures for the case where the FG inventory storage consists of a number of separate general-purpose silos capable of storing any grade as long as it is not mixed with any other grade.

Key words: Stochastic economic lot sizing problem; Dynamic scheduling; Process industry; Markov decision process

1. Introduction

Scheduling the 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 classical problem in production planning research that is often referred to as the *stochastic lot scheduling problem* (SLSP). Sox et al. (1999) distinguish between two versions of the SLSP, for consistency with the deterministic-demand literature: the *stochastic capacitated lot sizing problem* (SCLSP) and the *stochastic economic lot scheduling problem* (SELSP). The SCLSP assumes a finite planning horizon and allows for non-stationary demand, while the SELSP assumes an infinite planning horizon and stationary demand. The SCLSP is more appropriate for discrete-parts manufacturing, whereas the SELSP is better suited for continuous-process 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, furniture products, etc. Process industries, on the other hand, operate on material that is continuously flowing, as is the case with petroleum and coal products, metallurgical products, nonmetallic mineral products, basic chemicals, food and beverages, paper products, etc. Generally, process industries are capital intensive and focus on non-stop, high-volume, low-variety production.

This paper focuses on the SELSP. The deterministic version of the SELSP, the so-called ELSP, has received considerable attention in the literature over the past decades (e.g., see the surveys of Elmaghraby, 1978 and Salomon, 1991). Both analytical and heuristic solutions for the ELSP derive rigid cyclic production plans. Unfortunately, cyclic plans do not work well for the stochastic problem, for two reasons. Firstly, they focus on lot-sizing and not on dynamic capacity allocation, which is necessary to respond to random changes in demand. Secondly, in the stochastic problem, *finished-goods* (FG) inventories serve not only to reduce the number of changeovers, as is the case in the deterministic problem, 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 many process industries it is often the case that the products coming out of the production facility are variants of the same family that differ in one or more attributes, such as quality, color, consistency, weight, size, thickness, etc. These variants are frequently referred to as "grades". 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. This is because a grade changeover means a change in the physical or chemical properties of the product coming out of the facility (e.g., intrinsic viscosity) and results from the *gradual* change of the production conditions, i.e., the chemical reaction process variables, e.g., temperature, pressure, catalyst flow rate. 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 *A-B* and *B-C*, but not *A-C*. To indicate this ordering in the chain of allowable changeovers, we use the notation "*A-B-C*". Traditionally, multi-grade plants have been operated using cyclic production plans that take the form of rigid product *slates* or *wheels*, whereby all products are produced sequentially in a cycle of two phases: one phase with increasing order of grades, and the other in decreasing order. Nowadays, the product slate mode is often considered inadequate as it does not cope well with varying demands (Tousain and Bosgra 2006).

In many process industries it is also often the case that the production facility (usually a reactor) is never shut down, because of the tremendous cost of bringing it up or shutting it down, so the facility keeps producing 24 hours a day even during grade changeover transitions. The production rate of the facility may be fine-tuned once in a while in the long run so as to match the total long-run expected demand for all grades, in case the demand has seasonal or other long-run variations. For the purposes of short-term and medium-term scheduling, however, the production rate is considered to be constant and equal to (or very close to) the total medium-term expected demand for all grades. It has often been pointed out that in large-scale continuous processes, 100% utilization is often a policy objective because of the high capital cost of such processes and because variable costs are lowest at maximum throughput. In this case, it is the demand rate which is set equal to the (maximum) production rate, rather than the opposite. This can be accomplished by effectively managing the demand, e.g., finding customers and markets that are willing to absorb large fluctuations is sales volume in exchange for a low price. (Cooke and Rohleder, 2006). Whether the production rate is set equal to the demand rate or the reverse, the result is that on average, the FG inventory of the different grades remains unchanged. With this in mind, the main managerial concern is not to minimize the FG inventory holding cost, but to minimize the number of changeover transitions, while keeping the FG inventory within the bounds of the storage capacity. Avoiding grade changeovers is important, because during a changeover transition, the grade produced is off-specifications and the process is difficult to control.

The FG inventory buffer itself is finite and often consists of several separate silos, where each silo can store one type of grade at a time, so the FG inventories of the different grades are kept separately. In many cases, the number of silos is much larger than the number of grades. In other cases, the final grades are not kept in silos but are packed in small quantities (rolls, big bags, etc.) and stored in a common FG warehouse that can hold a large number of such quantities. In both situations, assuming that the FG inventory is common is realistic.

Our work in this paper was motivated by the need to find the optimal production schedule in a real continuous-process multi-grade *Polyethylene Terephthalate* (PET) resin plant. This need led to the development of two different mathematical models that address the production scheduling problem at two different levels: a short-term level and a medium-term level. The short-term scheduling problem, which is presented in Liberopoulos et al. (2009), is formulated as a deterministic, discrete-time, finite-horizon *mixed integer linear programming* (MILP) optimization model. It describes in great detail the real production scheduling problem in the short term (typically one week), where the orders for the different grades are known. Given that in real life, production and demand continue after the end of the scheduling horizon, the production schedule must be designed so as to hedge against the uncertainty of the unknown random demand after the end of the scheduling horizon. To accomplish this, it is necessary to develop a more macroscopic medium-term model that describes the system in less detail but takes into account the stochastic nature of demand. The goal of this paper is to develop and solve such a model.

More specifically, 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 FG inventory buffer with limited storage capacity. Demand that cannot be satisfied directly from stock is lost. Raw material is always available, and the production facility produces at a constant rate continuously even during changeover transitions. When the facility is set up to produce a particular grade, the only allowable changeovers are from that grade to the next lower or higher grade. In many industries, including the PET industry that motivated this work, 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 during the changeover transition is gradually changing from A to B. In this paper, we assume that the grade produced during a changeover from A to B is classified as A, and that the grade produced during the reverse changeover 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 the demand. As was explained earlier, on average, the FG inventory of the different grades remains unchanged, and the main managerial concern is to minimize the number of changeover transitions rather than FG inventory. For this reason, we chose not to include an inventory holding cost in our model, although we could have trivially included such a cost. In any case, excess inventory is penalized whenever the FG inventory buffer is full, in which case the grade coming out of the facility is spilled over. In many practical cases, the spilled-over material is discarded or is recycled into the raw material inventory. In other cases, it is driven into an evacuation buffer and from there it is disposed of at a discount price. In all cases, spilling over material is undesirable. If the long-run average percentage of spilled-over material is deemed high, then the management can take one or more of the following corrective actions: 1) lower the long-run production rate, 2) take measures to raise the long-run average demand, 3) take measures to lower the variability of demand, and 3) increase the FG buffer capacity. Actions 1 and 2 can help reduce the amount of spilled-over material, but may cause an increase in lost sales. Actions 3 and 4 can help reduce both the amount of spilled-over material and lost sales.

Our assumption that the FG inventory buffer is common is realistic, because as was explained earlier, in many practical situations each grade can be stored in extremely small storage units (e.g., big bags) compared to the total available space (e.g., warehouse). Nonetheless, in Section 6.2, we conduct further experiments after having redefined the inventory state space in our model by adjusting its outer face to accommodate the situation where the FG inventory storage consists of a number of separate general-purpose silos capable of storing any grade as long as it is not mixed with any other grade.

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 (long-run) expected average cost per period.

In theory, we can numerically solve the resulting MDP problem using the value iteration or any other appropriate method and obtain insight into the optimal changeover policy. We refer to this solution procedure as "exact", because it solves the exact problem. In practice, the exact solution procedure cannot solve problems with a large number of grades, because of the curse of dimensionality. For problems with N grades, $N > 3$, we develop a heuristic solution procedure that is based on decomposing the original N-grade problem into $(N - 2)$ 3-grade sub-problems and numerically solving each sub-problem using value iteration. 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 final changeover policy for the original N-grade problem is constructed by combining parts of the optimal changeover policies of the sub-problems. We call the heuristic outlined above *decomposition-based heuristic* (DBH).

Our assumption that all the changeover times are equal may be realistic in some cases, but may not hold in other cases. Assuming different changeover times would augment the state space of the MDP and would make the need for a heuristic solution procedure even more pressing. Yet, we don't expect that it would change the structure of the optimal policy.

The remainder of this paper is organized as follows. In Section 2, we review the literature on the SELSP that is most closely related to our work. In Section 3, we present the MDP model of the original N-grade problem, and we outline the value iteration method used to solve it. The DBH procedure for solving problems with more than three grades is presented in Section 4. In Section 5, we present numerical results for problem examples with 2-5 grades. For the 2- and 3-

grade examples, we use the exact solution procedure to obtain insights into the structure of the optimal changeover policy. For the 4- and 5-grade examples, we compare the performance of the DBH solution procedure against that of the exact procedure. In Section 6, we perform a further numerical investigation. First, we compare the performance of the DBH method against the performance of three simpler parameterized heuristics. Then, we compare the performance of the DBH and the exact solution procedures for the case where the FG inventory buffer consists of a number of separate general-purpose silos capable of storing any grade as long as it is not mixed with any other grade. Finally, we draw our conclusions in Section 7.

2. Literature review

The SELSP has received considerable attention in the literature because of its theoretical and practical importance. A comprehensive review of related works can be found in Sox et al. (1999) and more recently Winands et al. (2011). From these reviews, it is apparent that there have been two approaches for tackling the SELSP. One approach is to develop a cyclic schedule, i.e., a fixed production sequence, usually using a deterministic approximation of the stochastic problem, and then develop a control rule for the stochastic problem to pursue that schedule. The attractiveness of this approach lies on its ability to provide a practical solution for problems with a large number of products, as it breaks up the difficult dynamic scheduling problem into two easier to manage sub-problems, namely, sequencing and lot sizing, which are solved sequentially. A drawback of this approach, however, is that it may not respond effectively to random changes in demand, as was mentioned earlier.

The other approach, which we follow in this paper, is to develop a dynamic scheduling rule that determines which product to produce based on the current state of the system. Such a rule may be a simple heuristic or may be derived from an optimal control analysis of the problem. The literature on this approach, particularly the track that adopts an optimal control perspective, is scarce, because of the immense difficulty of obtaining an analytical solution even for problems of small size, and the computational challenge of numerically solving problems of realistic size.

One of the first exploratory papers on the SELSP is that by Vergin and Lee (1978). They examine simple dynamic sequencing heuristics for the SELSP with changeover costs but no changeover times. Graves (1980) looks at the SELSP as a discrete-time stochastic control problem with dynamic sequencing. He first solves a one-product problem with inventorybackorder 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.

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-MDP, 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.

Bruin and van der Wal (2010b) propose a two-step approach for solving discrete-time, multi-item production systems with generally distributed processing times, Poisson demands, unit changeover times, inventory holding costs and lost sales. In the first step, they find a good fixed-cycle production schedule, using a methodology that they develop in Bruin and van der Wal (2010a). In the second step, they obtain a one-step improvement of this schedule by minimizing the future expected costs under the assumption that after this decision, the original cyclic policy will be resumed.

In a recent paper, Löhndorf and Minner (2012) model the SELSP with compound Poisson demand and deterministic processing and changeover times as an infinite-horizon average-cost semi-MDP and develop an approximate value iteration method for solving it. This method is based on approximating the differential cost function with a linear combination of piecewise-constant functions and using a simulation-based stochastic gradient algorithm to update the weights of the constant function segments. Based on a numerical study, they conclude that for problems of small size (up to three products) a particular version of the approximate value iteration method outperforms simpler fixed-cycle and base-stock policies that have been optimized with a global search, but for larger problems, a fixed-cycle policy with preemption is the most reliable choice and the base-stock policy is more often the better choice.

Löhndorf et al. (2012) extend the model of Löhndorf and Minner (2012) to include sequence-dependent changeover times and propose three cycling policies with base stock levels for solving the problem. Based on numerical experimentation, they conclude that a "balanced-cycle" policy, which balances between minimizing setup times and achieving a sequence with a regular production pattern, outperforms the simpler common-cycle and fixed-cycle policies.

Finally, Karmarkar and Yoo (1994) and Sox and Muckstadt (1997) 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.

There has also been a somewhat parallel stream of works on the dynamic scheduling of failure-prone flexible manufacturing systems that are based on a flow control approach. In much of that literature, it is assumed that the production capacity changes randomly due to machine failures and repairs, while the demand rate remains constant. When the manufacturing system is not perfectly flexible but requires setups, several authors have looked at setup scheduling policies that use "corridors" in the product surplus/backlog space to determine the timing of the setup changeovers in order to guide the trajectory in the desired direction (e.g., Sharifnia et al., 1991, Liberopoulos and Caramanis, 1997, Elhafsi and Bai, 1997).

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 as such is more closely related to Graves (1980), Qiu and Loulou (1995), and Löhndorf and Minner (2012). It is also closely related to Sharifnia et al. (1991), Liberopoulos and Caramanis (1997), and Elhafsi and Bai (1997), as we use a qualitatively similar approach and obtain a similar setup changeover policy. Our work differs from previous works in that it considers a new variant of the SELSP, where the FG inventory buffer has finite storage capacity, the facility produces continuously at a constant rate, and the only allowable changeovers are from one grade to the next lower or higher grade. The latter feature renders problems with a large number of grades amenable to heuristic solution procedures that are based on decomposing the original problem into several smaller (i.e., with fewer grades) sub-problems which are computationally easier to solve and then constructing the final solution for the original problem by combining parts of the sub-problem solutions. We develop such a procedure in Section 4. Finally, we should point out that the feature of sequence restricted changeovers that we consider should not be confused with the completely different feature of sequence dependent changeover times that has received wider attention in the literature (e.g., see Karalli and Flowers, 2006 and Löhndorf et al., 2012).

3. 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 $n + 1$, $n = 1, \ldots, N - 1$. The changeover time is one period. In each time period, the production facility produces P units of the grade that it is set up for at the beginning of the period. The quantity produced is stored in a common FG inventory buffer which has 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, where D_n, $n = 1, \ldots, N$, is the demand for grade n. The demands D_n are discrete random variables with known stationary joint probability distribution. For each grade n, the part of the demand that cannot be satisfied from FG inventory, if any, is lost, incurring a lost-sales cost of \$ CL_n per unit of unsatisfied demand. For the purposes of short- to medium-term scheduling that we consider in this paper, we assume that P is fixed and equal 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 (called the "setup" state) and x_n, $n = 1, \ldots, N$, 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, $n = 1, \ldots, N$, such that

$$0 \leq \sum_{n=1}^{N} x_n \leq X \tag{1}$$

For each setup state s, the number of inventory level states $(x_1, \ldots, x_N)$ is equal to the number of ways X indistinguishable balls can be distributed to $N + 1$ labeled urns, where the $(N + 1)$st urn is for the unused or empty inventory storage space. This number is equal to Comb$(X + N, N)$; hence the size of the state space is $N(X+N)!/(X!N!)$.

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 $U(1) = \{1, 2\}$, $U(s) = \{s - 1, s,$

$s + 1\}$, $s = 2, \ldots, N - 1$, and $U(N) = \{N - 1, N\}$. If the decision is to initiate a changeover, then the new setup of the facility, i.e., after the changeover is completed, will be in effect at the beginning of the next period, since the changeover time is one period. A decision to initiate a changeover at the beginning of a period incurs a changeover cost of \$ 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

$$s' = u$$

$$x_n' = (x_n + p(\mathbf{y})\cdot I_{n=s} - D_n)^+,\ n = 1, \ldots, N$$

where $p(\mathbf{y})$ is the amount of material added to the FG buffer after the facility produces P units, minus any spillage, and before the demand is realized, and is given by

$$p(\mathbf{y}) \equiv \min\left(P, X - \sum_{n=1}^{N} x_n\right) \tag{2}$$

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) expected average cost per period. To find such a policy, we need to solve Bellman's dynamic programming equation, which for our problem can be written as

$$J + V(\mathbf{y}) = \min_{u\in U(s)} T_u(V(\mathbf{y})) \tag{3}$$

where J is the optimal (minimum) expected average cost per period, $V(\mathbf{y})$ is the optimal differential cost starting from state $\mathbf{y}$, and $T_u(\cdot)$ is a mapping defined as $T_u(V(\mathbf{y})) \equiv E_D\{g(\mathbf{y},u,\mathbf{D}) + V(\mathbf{y}')\}$. The minimizer of the Bellman equation determines the optimal policy when the system is in state $\mathbf{y}$, denoted by $\mu^*(\mathbf{y})$.

To solve Bellman's equation, we use the method of successive approximations of the optimal differential cost functions, which is well-known as the value iteration method. We denote by $V_k(\mathbf{y})$ the value of the optimal 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

$$V_{k+1}(\mathbf{y}) = T(V_k(\mathbf{y})) - T(V_k(\hat{\mathbf{y}})) \tag{4}$$

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. This state must be recurrent to guarantee that the value iteration method converges to an optimal policy. Note that in each iteration the optimal differential cost of the special state is reset to zero. Assuming that the iteration scheme converges to some values $V(\mathbf{y})$, then from recursion (4), 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 (3) reveals that $J = T(V(\hat{\mathbf{y}}))$.

To implement the value iteration method, at each iteration $k = 1, 2, \ldots$, 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.

4. DBH solution procedure

Although the exact method presented in the preceding section can in principle determine the optimal policy for any number of grades, it becomes computationally very demanding for more than three grades. In this section, we propose a heuristic procedure – the DBH – that decomposes any N-grade problem, $N > 3$, into several 3-grade sub-problems and then combines the sub-problem solutions (determined by the exact method) to construct a final policy for the original problem.

The DBH procedure that we propose works as follows. Let S denote the original N-grade problem. For each grade n, $n = 2, \ldots, N-1$, 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 $1, \ldots, n-1$, and the high grade is the composite of all grades that are higher than n, i.e., grades $n+1, \ldots, N$; 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 $s_n \in \{n-1, n, n+1\}$ and w_n and z_n are the "aggregate" inventory levels of the low and high composite grades, respectively, which represent in some aggregate way the total of their individual components through some function h, i.e., $w_n = h(x_1, \ldots, x_{n-1})$ and $z_n = h(x_{n+1}, \ldots, x_N)$. It should be understood that when s_n is equal to $n-1$ or $n+1$ in sub-problem S_n, the facility is set up to produce one of the grades that comprise the low grade or the high grade, respectively. 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 presented in the previous section to obtain the optimal policy of each sub-problem S_n, denoted by $\mu_n^*(\mathbf{y}_n)$. The DBH then constructs the changeover policy for the original N-grade problem, denoted by $\mu^h(\mathbf{y})$, by combining parts of the optimal policies of the sub-problems, as follows:

$$\mu^h(1, x_1, \ldots, x_N) = \mu_2^*(1, w_2, x_2, z_2)$$

$$\mu^h(n, x_1, \ldots, x_N) = \mu_n^*(n, w_n, x_n, z_n), \quad n = 2, \ldots, N-1$$

$$\mu^h(N, x_1, \ldots x_N) = \mu_{N-1}^*(N, w_{N-1}, x_{N-1}, z_{N-1})$$

Next, we discuss how to determine an appropriate form for function h.

First, note that in S_2, $w_2 = h(x_1)$, i.e., w_2 is the aggregate inventory level of a single grade, namely grade 1; therefore, it is reasonable to simply set $h(x_1) \equiv x_1$ so that $w_2 = h(x_1) = x_1$. Similarly, in S_{N-1}, we set $h(x_N) \equiv x_N$, so that $z_{N-1} = h(x_N) = x_N$. Let us next focus on w_n, $n > 2$, as z_n is obtained in a symmetric way.

An obvious choice for the aggregate inventory level of the composite of grades 1, …, $n-1$ in sub-problem S_n is to set it equal to the sum of the inventory levels of the individual grades, i.e., set $w_n \equiv x_1 + \ldots + x_{n-1}$. This is a reasonable choice, especially with respect to estimating potential spill-over costs; however, it fails to detect the situation where the sum $x_1 + \ldots + x_{n-1}$ is high, implying that the composite grade has a low risk of stocking out, yet one (or more) of its individual components, x_1, …, x_{n-1}, is (are) low, implying that the corresponding individual grade(s) has(ve) a high risk of stocking out, which may lead to significant lost-sales costs. We refer to this situation as the "imbalance problem," because one or more of the individual inventory levels are much lower than the average.

To illustrate the imbalance problem, suppose that in an N-grade problem, where $N > 4$, the facility is currently set up to produce grade 4 and that the inventory levels of grades 1-4 are $x_1 = 15$, $x_2 = 15$, $x_3 = 0$, $x_4 = 6$. 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 from solving S_4 might indicate that it is optimal for the facility not to change over to the low composite grade, because there is plenty of it (30 units) in storage compared to the inventory level of the middle grade 4, which is much lower (6 units). What the DBH fails to see here is that although w_4 is relatively high, its individual components are quite imbalanced – in fact, one of them, namely x_3, is zero. In this case, unless the facility changes over to grade 3, a heavy lost-sales cost is likely to be incurred in the current and in the following period.

In case of an imbalance among the individual inventory levels $x_1, \ldots, x_{n-1}$, we need to find an alternative definition for the aggregate inventory level, w_n, which reflects this imbalance. To this end, let $ILS_n(x_1, \ldots, x_{n-1})$ denote the total expected lost sales over the individual grades $1, \ldots, n-1$, as a function of $x_1, \ldots, x_{n-1}$, i.e.,

$$ILS_n(x_1, \ldots, x_{n-1}) \equiv \sum_{i=1}^{n-1} E\left[\left(D_i - x_i\right)^+\right]$$

Also, let $CLS_n(w_n)$ denote the expected lost sales of the composite of grades $1, \ldots, n-1$, as a function of the aggregate inventory level, w_n, i.e.,

$$CLS_n(w_n) \equiv E\left[\left(\left(\sum_{i=1}^{n-1} D_i\right) - w_n\right)^+\right]$$

The expected lost sales is a measure of the effect of stock-outs. A desirable property of w_n is that it should yield the same measure of the effect of stock-outs as that yielded by the individual grades, i.e., w_n should satisfy

$$CLS_n(w_n) = ILS_n(x_1, \ldots, x_{n-1}) \tag{5}$$

If $ILS_n(x_1, \ldots, x_{n-1}) = 0$, then there is little risk of a grade stocking out; hence, there is no imbalance problem. In this case, setting $w_n = x_1 + \ldots + x_{n-1}$ would make $CLS_n(w_n) = ILS_n(x_1, \ldots, x_{n-1}) = 0$. If $ILS_n(x_1, \ldots, x_{n-1}) > 0$, on the other hand, then there is a more significant risk that one or more grades may stock out; hence, there is an potential imbalance problem. In this case, setting $w_n = x_1 + \ldots + x_{n-1}$ would make $CLS_n(w_n) \leq ILS_n(x_1, \ldots, x_{n-1})$; hence, w_n would underestimate the effect of a possible stock-out, due to the imbalance. To remedy this, we should choose a smaller value of w_n that satisfies (5). Let $v_n(x_1, \ldots, x_{n-1})$, denote the value of w_n that satisfies (5), i.e.,

$$v_n(x_1, \ldots, x_{n-1}) \equiv CLS_n^{-1}\left(ILS_n(x_1, \ldots, x_{n-1})\right)$$

It can be easily shown that $v_n(x_1, \ldots, x_{n-1}) \leq x_1 + \ldots + x_{n-1}$. To compute $v_n(x_1, \ldots, x_{n-1})$, we need the probability distribution of the demand of the composite grade, which we can derive by convolving the probability distributions of the demands of the individual grades. We propose a faster alternative that is based on approximating $ILS_n(x_1, \ldots, x_{n-1})$ and $CLS_n(w_n)$ by the following expressions, respectively:

$$ILS_n(x_1,\ldots,x_{n-1}) \approx \sum_{i=1}^{n-1}\left(E[D_i]-x_i\right)^+ \tag{6}$$

$$CLS_n(w_n) \approx \left(\left(\sum_{i=1}^{n-1} E[D_i]\right)-w_n\right)^+$$

If $ILS_n(x_1, \ldots, x_{n-1}) > 0$, the above approximations allow us to approximate v_n by the following expression:

$$v_n(x_1,\ldots,x_{n-1}) \approx \sum_{i=1}^{n-1} E[D_i]-ILS_n(x_1,\ldots,x_{n-1}) = \sum_{i=1}^{n-1}\min\left(x_i, E[D_i]\right) \tag{7}$$

Although $v_n(x_1, \ldots, x_{n-1})$ reflects the imbalance between the individual inventory levels $x_1, \ldots, x_{n-1}$, it may be significantly smaller than the sum $x_1 + \ldots + x_{n-1}$, which is the natural candidate for the aggregate inventory level of the composite grade. With this in mind, we propose to set w_n equal to a linear combination of $x_1 + \ldots + x_{n-1}$ and $v_n(x_1, \ldots, x_{n-1})$ (rounded to the nearest integer), if $ILS_n > 0$; otherwise, set it equal to $x_1 + \ldots + x_{n-1}$, namely,

$$w_n = h(x_1,\ldots,x_{n-1}) = \begin{cases} x_1+\ldots+x_{n-1}, & \text{if } ILS_n(x_1,\ldots,x_{n-1})=0 \\ \alpha v_n(x_1,\ldots,x_{n-1})+(1-\alpha)(x_1+\ldots+x_{n-1}), & \text{if } ILS_n(x_1,\ldots,x_{n-1})>0 \end{cases} \tag{8}$$

where α is a coefficient, such that $0 \le \alpha \le 1$, and $ILS_n(x_1, \ldots, x_{n-1})$ and $v_n(x_1, \ldots, x_{n-1})$ are approximated by (6) and (7), respectively. Note that if $\alpha = 0$, the rule prescribed by (8) becomes $w_n = x_1 + \ldots + x_{n-1}$, irrespectively of whether there is an imbalance problem or not, whereas if $\alpha = 1$, w_n is set equal to $x_1 + \ldots + x_{n-1}$, if $ILS_n(x_1, \ldots, x_{n-1}) = 0$, and equal to $v_n(x_1, \ldots, x_{n-1})$ if $ILS_n(x_1, \ldots, x_{n-1}) > 0$. It is reasonable to expect that the smaller the imbalance problem, the smaller the optimal value of α, and the better the performance of the DBH. In Section 5.3, we investigate the performance of the DBH as a function of coefficient α.

The DBH policy that we described above, as any feedback policy, satisfies an expression similar to Bellman's equation (3), without the minimization, i.e., it satisfies

$$J^h + V^h(\mathbf{y}) = T_{u^h}(V^h(\mathbf{y})) \tag{9}$$

where J^h is the expected average cost per period and $V^h(\mathbf{y})$ is the differential cost starting from state $\mathbf{y}$, when the DBH policy $u^h = \mu^h(\mathbf{y})$ is used. Note that J^h, $V^h(\mathbf{y})$, and $\mu^h(\mathbf{y})$ also depend on α, but we omitted this dependence here for notational simplicity.

One way to evaluate the DBH policy is to use the method of successive approximations of the functions $V^h(\mathbf{y})$ (value iteration). More specifically, if we denote by $V_k^h(\mathbf{y})$ the values of the differential cost function at the k^{th} iteration, then the values at the $(k + 1)^{\text{th}}$ iteration are obtained from the previous iteration by a recursion similar to (4), without the minimization, i.e.,

$$V_{k+1}^h(\mathbf{y}) = T_{u^h}(V_k^h(\mathbf{y})) - T_{u^h}(V_k^h(\hat{\mathbf{y}})) \tag{10}$$

Note that as in (4), at each step of iteration (10), the differential cost of the special state $\hat{\mathbf{y}}$ is reset to zero. Assuming that the iteration scheme converges to some value $V^h(\mathbf{y})$, for each state $\mathbf{y}$, then the expected average cost per period of the DBH policy is given by $J^h = T_{u^h}(V^h(\hat{\mathbf{y}}))$.

An alternative way to evaluate the DBH policy is to use simulation. Our numerical experience for 4-grade and 5-grade problems showed that simulation is faster than the method of value iteration by as much as 100 times.

5. Numerical results

In this section, we present numerical results for problem examples with 2-5 grades. First, we solve some indicative 2-grade and 3-grade examples using the exact solution procedure. For these examples, we briefly discuss the optimal changeover policy and performance. More cases and discussion of the results for the 2-grade and 3-grade examples can be found in Hatzikonstantinou (2009). Then, we solve 4-grade and 5-grade examples using both the exact and the DBH solution procedures. We discuss the performance and computational efficiency of the DBH procedure, and we explore how they are affected by the distribution of the relative market size of the different grades and the size of weight α in expression (8). In all the examples, we set $\hat{\mathbf{y}}$ to be the state where $s = 1$ and $x_n = 0$, $n = 1, \ldots, N$. For practical purposes, this state will be recurrent under mild assumptions on the distribution of the demand and the relative lost-sales cost of grade 1. For example, if $\Pr(D_n > P) > 0$, $n = 1, \ldots, N$, state $x_n = 0$, $n = 1, \ldots, N$ will be accessible from any other state. In addition, if the lost-sales cost of grade 1 is comparable to that of the other grades, the set-up state $s = 1$ should be visited infinitely often under the optimal policy.

5.1 2-grade example

First, we consider a 2-grade example ($N = 2$), where $P = 5$, and the demand distribution for the two grades is given in Table 1; the last column shows the coefficient of variation (C.V.) which is defined as the standard deviation over the mean.

Table 1: Probability distribution of demand, $\Pr(D_n=i)$, for the 2-grade example

	i								
n	0	1	2	3	4	5	6	$E[D_n]$	C.V.$[D_n]$
1	0.1	0.15	0.15	0.2	0.15	0.15	0.1	3	0.6055
2	0.15	0.15	0.4	0.15	0.15	0	0	2	0.6124

We run the value iteration procedure outlined in Section 3 for six combinations of storage capacity, X, and cost rate parameters, CC, CS, CL_1 and CL_2, where we assumed that $CL_1 = CL_2 = CL$. The results are shown in Table 2, spread in three rows for each case. The first row shows the number of iterations until convergence, denoted by k_c, for convergence tolerance criterion $\varepsilon = 0.001$, the total CPU time in hours on an Intel Pentium PC at 2.99 GHz with 1 GB RAM, and the resulting optimal expected average cost per period, J. The second row shows the per period expected average number of changeovers, units spilled over, and lost sales for each grade, denoted by $E[C]$, $E[S]$, $E[L_1]$ and $E[L_2]$, respectively. These quantities are related to J by the expression $J = CC{\cdot}E[C] + CS{\cdot}E[S] + CL{\cdot}E[L_1] + CL{\cdot}E[L_2]$. The third row shows the elements of the inventory level vector that minimizes $V(s, x_1, x_2)$, denoted by $(x_1^*(s), x_2^*(s))$, for $s = 1, 2$.

Table 2: Results for the 2-grade example

				$X = 40$				$X = 60$				$X = 80$			
Case	CC	CS	CL	k_c	CPU	J	-	k_c	CPU	J	-	k_c	CPU	J	-
				$E[C]$	$E[S]$	$E[L_1]$	$E[L_2]$	$E[C]$	$E[S]$	$E[L_1]$	$E[L_2]$	$E[C]$	$E[S]$	$E[L_1]$	$E[L_2]$
				$x_1^*(1)$	$x_2^*(1)$	$x_1^*(2)$	$x_2^*(2)$	$x_1^*(1)$	$x_2^*(1)$	$x_1^*(2)$	$x_2^*(2)$	$x_1^*(1)$	$x_2^*(1)$	$x_1^*(2)$	$x_2^*(2)$
1	1	5	5	253	0.4045	0.9804	-	642	2.2917	0.6168	-	1210	7.5323	0.4494	-
				0.1991	0.0777	0.0412	0.0374	0.1285	0.0486	0.0248	0.0243	0.0943	0.0353	0.0180	0.0177
				1	21	22	0	2	32	34	0	2	44	46	0
2	2	5	5	240	0.3808	1.1616	-	599	2.0880	0.7327	-	1125	7.2427	0.5343	-
				0.1662	0.0825	0.0419	0.0415	0.1067	0.0517	0.0252	0.0270	0.0781	0.0377	0.0174	0.0206
				1	22	23	0	1	35	36	0	1	48	49	0

From the results, it can be seen that the values of k_c and CPU range from 240 iterations in 0.3808 hours (~23 min), in case 2 ($X = 40$), to 1210 iterations in 7.5323 hours, in case 1 ($X = 80$).

As expected, k_c and CPU are increasing in X, whereas $E[C]$, $E[S]$ $E[L_1]$, $E[L_2]$ and J are decreasing. Also, J is increasing in the cost rate parameters.

Figure 1 shows the optimal changeover policy as a function of inventories x_1 and x_2, for cases 1 and 2 of Table 2, for X = 40, and is representative of all other cases. The optimal policy partitions the inventory space in several regions, where each region is characterized by a different changeover action, as described in Table 3. If the inventory level vector is in region *c*, the facility changes over from one grade to the other in each period ("chattering"). If the inventory level vector is in region *b*, the facility keeps producing the grade it is set up for. As a result, the inventory level vector moves on a trajectory, which is more or less parallel to the outer face of the triangular state space, until it crosses one of the borders of region *b*, entering region *a* or *d*. At this point, the facility changes over to the other grade, and the inventory level vector reverses its direction, entering region *b* again, heading for the other border. Note that region *b* is wider in case 2, where the changeover cost is greater, indicating that in case 2, the facility produces longer runs (campaigns) of each grade with less frequent changeovers. In fact, the widening up of region *b* in case 2 is so big that it makes region *c* disappear.

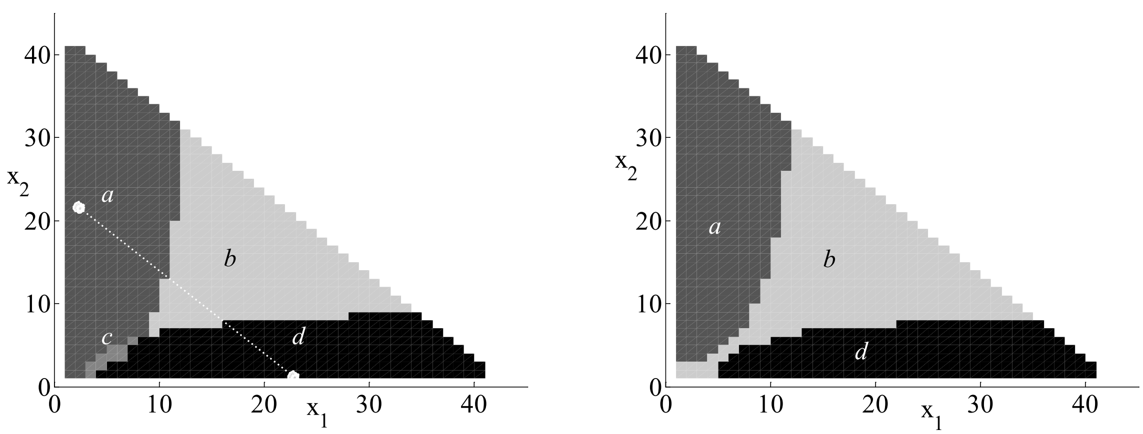

Figure 1: Optimal changeover policy for cases 1 (left) and 2 (right) of Table 2, for X = 40

The borders of region *b* towards its wider end tend to align themselves to the orthogonal lines $x_1 = c_1$ and $x_2 = c_2$, respectively, where c_1 and c_2 are some constants. This means that when the facility is set up for, say, grade 2, it will change over to grade 1, if x_1 drops below c_1, irrespectively of the value of x_2, as long as x_2 is high.

Table 3: Optimal policy $\mu^*(s, (x_1, x_2) \in R)$ for the 2-grade example

	s		
R	1	2	Description
a	1	1	Changeover to grade 1
b	1	2	Do not changeover
c	2	1	Changeover to the other grade
d	2	2	Changeover to grade 2

In all cases, the minimizer of $V(s, x_1, x_2)$, $(x_1^*(s), x_2^*(s))$, for any given setup state s is such that $x_n^*(s) \approx 0$, for $n = s$, and $0 << x_n^*(s) << X$, for $n \neq s$. For example, for the case where $X = 40$, $x_1^*(1) = 1$ and $x_2^*(1) = 21$. The vector (1, 21) can be thought of as the ideal value of the inventory level vector $(x_1(1), x_2(1))$, meaning that if the facility is set up to produce grade 1 ($s = 1$), the best place to be in terms of inventory is to have only one unit of grade 1 and 21 units of grade 2. This will allow a long production campaign of grade 1 with a low risk of stocking out of grade 2 or spilling over the produced grade 1, because of the lack of storage space. The line connecting the inventory level vectors that minimize the differential cost functions of the two setup states (i.e., the line connecting inventory points (1, 21) and (22, 0)) is shown as a dotted line in Figure 1 (left) for case 1. Under the optimal changeover policy, the inventory level vector moves on average back and forth along the segment of that dotted line that falls in region b, as the facility changes over from one grade to the other, whenever the inventory level vector enters region a or d.

The optimal inventory level of the grade that the facility is not set up for acts as a “safety stock”. Although not shown here for space considerations, it is increasing in CL to better hedge against stock-out occurrences, and decreasing in CS to better hedge against spill-over occurrences. From Table 2, it can be seen that this safety stock is insensitive to CC. CC primarily affects the width of region b and therefore the “cycle stock” and changeover frequency.

5.2 3-grade example

Next, we consider a 3-grade ($N = 3$) example that originated from the real dynamic scheduling application of a continuous-flow PET resin processing plant that produces three grades, presented in Liberopoulos et al. (2009). The production rate and the FG storage capacity are $P = 6$ and $X = 115$. The distribution of the discretized demand for the three grades is given in Table 4. The cost rate parameters that we used are $CC = 1$, $CS = CL_1 = CL_2 = 2$, to reflect the fact that the plant manager wishes to avoid frequent changeovers, but is even more wary about material spill-over and lost sales.

Table 4: Probability distribution of demand, $\Pr(D_n=i)$, for the 3-grade example

	i												
n	0	1	2	3	4	5	6	7	8	9	10	$E[D_n]$	C.V.$[D_n]$
1	0.1676	0.1429	0.3214	0.1538	0.1016	0.0604	0.0247	0.0110	0.0137	0.0027	0.0000	2.3159	0.7625
2	0.5000	0.1648	0.1071	0.0824	0.0604	0.0302	0.0220	0.0137	0.0027	0.0110	0.0055	1.4231	1.4254
3	0.1519	0.2652	0.2956	0.0718	0.0663	0.0525	0.0442	0.0138	0.0276	0.0028	0.0083	2.2901	0.9067

We solved the problem optimally using the value iteration method outlined in Section 3. The method converged after 533 iterations that took 269 hours on an Intel Pentium PC at 2.99 GHz with 1 GB RAM, for convergence tolerance criterion $\varepsilon = 0.01$. The resulting optimal expected average cost per period, J, is 0.4522. As in the 2-grade example, the optimal changeover policy partitions the inventory space in several regions, each characterized by a different optimal changeover action. Figure 2 shows the optimal policy as a function of inventory levels x_1 and x_3, for different values of x_2. The optimal changeover action in each region is given in Table 5.

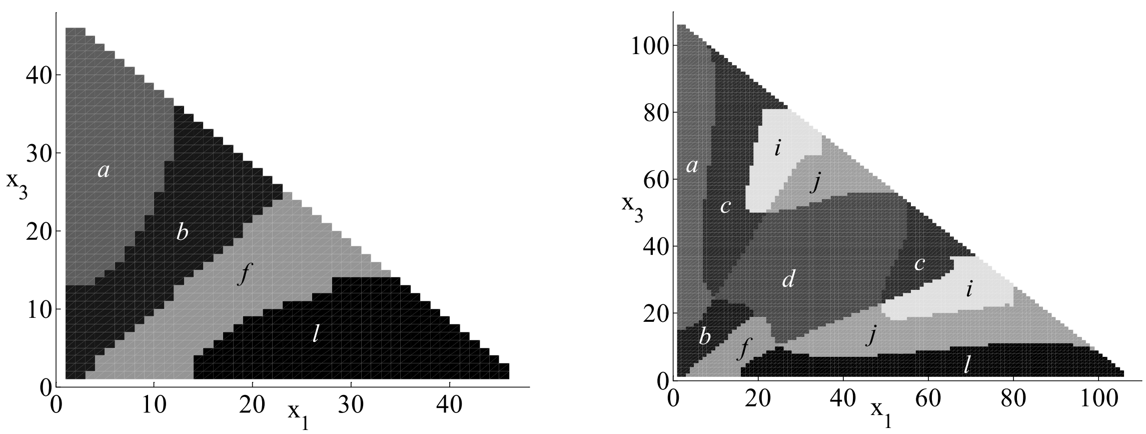

Figure 2: Optimal changeover policy for $x_2 = 70$ (left) and $x_2 = 10$ (right), for the 3-grade example

Figure 2 (left), shows the optimal changeover policy for x_2 (= 70) >> $x_1 + x_3$. It can be seen that if $(x_1, x_3) \in a$, in which case $x_2 >> x_3 >> x_1$, then the production facility must change over to the next lower grade so that it is eventually set up for grade 1, because x_1 is significantly lower than x_2 and x_3. If $(x_1, x_3) \in b$, in which case $x_2 >> x_3 > x_1$, then the facility must change over to grade 1, if it is set up for grade 2. If it is set up for grade 3, however, it need not change over to grade 2 (to eventually change over to grade 1), because x_1 is not that much lower than x_3 to justify the cost of such a changeover. The optimal changeover policies in regions l and f are symmetric to those in regions a and b, respectively, with the roles of x_1 and x_3 being reversed.

Table 5: Optimal policy $\mu^*(s, R)$ for the 3-grade example

R	s = 1	2	3	Description
a	1	1	2	Changeover to the next lower grade
b	1	1	3	If set up for grade 2, changeover to grade 1
c	1	2	2	If set up for grade 3, changeover to grade 2
d	1	2	3	Do not changeover
e	1	3	2	If set up for grade 2 or 3, changeover to grade 3 and 2, respectively
f	1	3	3	If set up for grade 2, changeover to grade 3
g	2	1	2	If set up for grade 1 or 3, changeover to grade 2; if set up for grade 2, changeover to grade 1
h	2	1	3	If set up for grade 1 or 2, changeover to grade 2 and 1, respectively
i	2	2	2	Changeover to grade 2
j	2	2	3	If set up for grade 1, changeover to grade 2
k	2	3	2	If set up for grade 1 or 3, changeover to grade 2; if set up for grade 2, changeover to grade 3
l	2	3	3	Changeover to the next higher grade

Figure 2 (right), shows the optimal changeover policy for x_2 (= 10) << $x_1 + x_3$. It can be seen that, in addition to regions a, b, f and l, seven new regions enter the picture, and the overall optimal policy looks rather complicated. Moreover, the same changeover action may apply to more than one regions in different parts of the state space. The explanation of the optimal policy in different regions can be found in Hatzikonstantinou (2009).

Table 6 shows the elements of the inventory level vector that minimizes the differential cost function $V(s, x_1, x_2, x_3)$, for each set up state s. As in the 2-grade example, the minimizing inventory level of the grade being produced is close to zero, whereas the minimizing inventory level of the grades not being produced are positive and quite big. In fact, the further away (in terms of number of changeovers) a grade is from the grade produced, the higher its minimizing inventory level. For example, when the production is setup to produce grade 1 (s = 1), $x_1^*(1) = 2$, which is close to zero, $x_2^*(1) = 22$, which is significantly higher than zero, and $x_3^*(1) = 47$, which is even higher.

Table 6: Minimizing inventory level, $x_n^*(s)$, for the 3-grade example

	n		
s	1	2	3
1	2	22	47
2	22	3	56
3	47	31	3

5.3 4-grade and 5-grade examples

Finally, we consider 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, $j = A, B, C, D$, whose distributions are given in Table 7.

Table 7: Probability distribution of demand, $\Pr(D_j = i)$, for the 4-grade and 5-grade examples

	i					
j	0	1	2	3	$E[D_j]$	C.V.$[D_j]$
A	0.4	0.5	0.05	0.05	0.75	1.0220
B	0.25	0.5	0.25	0	1	0.7071
C	0.25	0.25	0.5	0	1.25	0.6633
D	0.05	0.2	0.45	0.3	2	0.4183

For each example (N = 4 and N = 5), we consider four different cases. In each case, the set of the probability distributions of the demands of the different grades is the same and such that the total expected demand is equal to the production rate. The difference between the cases is in the order in which these distributions appear in the chain of allowable changeovers. For instance, in all cases of the 4-grade example, we assume that the demands of two of the grades are identically distributed to random variable D_B, which has an expected value of 1, and the demands of the other two grades are identically distributed to random variable D_D, which has an expected value of 2. In other words, two grades have low demand and two grades have high demand. In case 1, the grades with the low demand are the end grades, 1 and 4, whereas the grades with the high demand are the middle grades, 2 and 3. To indicate this order we use the notation "*B-D-D-B*". In case 2, the order is *D-D-B-B*, which means that grades 1 and 2 have high demand and grades 3 and 4 have low demand, and so on. Hence, each case represents a different way that total expected demand is distributed among the individual grades.

First, we solved each case optimally using the value iteration procedure described in Section 3, for convergence tolerance criterion ε = 0.001. Then, we solved each case using the DBH procedure described in Section 4. In the implementation of the DBH procedure we used expressions (6) and (7) to approximate $ILS_n(x_1, \ldots, x_{n-1})$ and $v_n(x_1, \ldots, x_{n-1})$, respectively, and expression (8) to estimate the aggregate inventory levels of the composite grades, w_n, for 11 values of α ranging from 0 to 1 with a step size of 0.1. In all cases, we assumed that $CC = CS = CL_n = 1$, $n = 1, \ldots, 5$, and $P = 6$.

The results for the 4-grade example, for $X = 30$, are shown in Table 8. The CPU times reported are in hours on an Intel Core i7 PC at 2.67 GHz with 3 GB RAM. For the DBH, we show the total CPU time in hours that it took to solve the $(N - 2)$ 3-grade sub-problems and generate the DBH policy – which is what counts here – but not the time it took to evaluate the DBH policy. As was mentioned at the end of Section 4, the time it takes to evaluate the DBH policy using the value iteration method is significant, whereas the alternative of using discrete-time system simulation is much faster. In all cases of the 4-grade problem, we used the value iteration method. The optimal value of α, among the 11 values examined, is denoted by α^*, and the corresponding expected average cost per period is denoted by $J^h(\alpha^*)$. The last column of Table 8 shows the percent cost increase between the DBH and the optimal policy.

Table 8: Performance of the exact and DBH procedures for the 4-grade example for $X = 30$

	Demand	Exact			DBH			% cost
Case	pattern	k_c	CPU	J	CPU	α^*	$J^h(\alpha^*)$	increase
1	*B-D-D-B*	55	8.57	1.0034	0.0461	0.7	1.2442	24.00
2	*D-D-B-B*	156	24.29	1.0927	0.0574	0.5	1.2253	12.13
3	*D-B-D-B*	187	29.12	1.1835	0.0748	0.1	1.3207	11.59
4	*D-B-B-D*	110	17.13	1.2881	0.1040	0.1	1.3139	1.96

In case 1, the grades with the highest expected demands are in the middle of the chain of allowable changeovers, whereas in case 4, they are at the two ends of the chain. Hence, in case 1 the dispersion of the total expected demand among the individual grades is relatively low, because most of the time the production facility will be changing over between the highly demanded grades, 2 and 3, which are adjacent. In case 4, on the other hand, the dispersion of the total expected demand is relatively high, because most of the time the production facility will be changing over between the highly demanded grades, 1 and 4, which are spaced 3 grades apart. Cases 2 and 3 are intermediate cases.

From the results, it can be seen that as we move from case 1 to case 4, i.e., as the dispersion of the total demand among the individual grades increases, the expected average cost per period increases, because the number of changeovers needed to effectively meet the demands for all the grades increases. To see this, note that in case 1, every time the facility must change over between the highly demanded grades, 2 and 3, one changeover is needed, namely, 2→3. In case 4, however, when the facility must change between the highly demanded grades, 1 and 4, three costly but inevitable changeovers are needed, namely 1→2, 2→3, and 3→4. During the latter two changeovers, the lowly demanded grades, 2 and 3, are each produced for one period. These inevitable single-period production runs prevent the inventory levels of grades 2 and 3 from dropping too much on average and causing a significant imbalance among the inventory levels of all the grades. This suggests that the bigger the dispersion of the total demand among the individual grades, the smaller the imbalance problem. Moreover, as was mentioned in Section 4, the smaller the imbalance problem, the smaller the optimal value of α, and the better the performance of the DBH. This explains why, as we move from case 1 to case 4, α^* and the percent cost increase between the DBH policy and the optimal policy decrease. Actually, in all cases, except case 1, $J^h(\alpha)$ is relatively insensitive to parameter α, as can be seen from Figure 3.

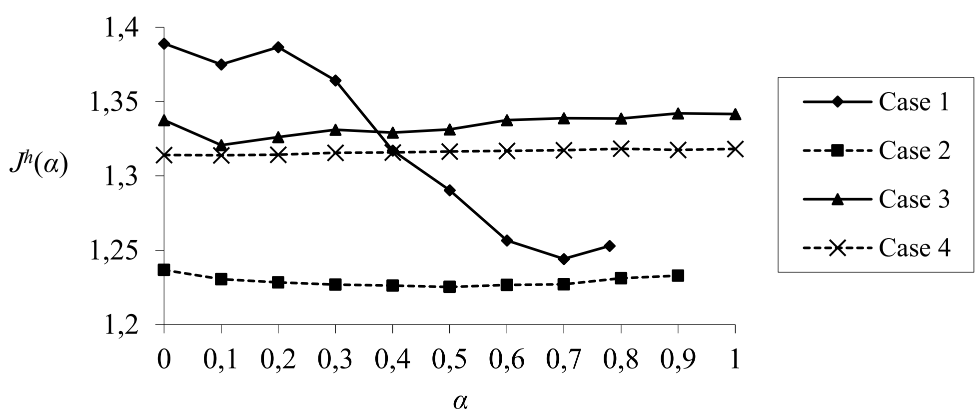

Figure 3: Expected average cost per period of the DBH, $J^h(\alpha)$, vs. α, for the 4-grade example

Finally, the cost increase when using the DBH instead of the exact method is 12.43% on average and ranges between 2.00% for case 4 and 24% for case 1. Note, however, that the DBH method is between 160 and 420 times faster than the exact method.

The results for the 5-grade example, for $X = 20$, are shown in Table 9. In all cases, we used discrete-time system simulation to evaluate the DBH policy. To obtain each estimate $J^h(\alpha)$ and its 95% confidence interval, denoted by "c.i.", we run 60 simulations, each with a time horizon of 100,000 time units. In all cases, the width of the resulting c.i. was approximately 1‰ of the estimated $J^h(\alpha)$ value.

Table 9: Performance of the exact and DBH procedures for the 5-grade example for $X = 20$

Case	Demand pattern	Exact			DBH			% cost increase
		k_c	CPU	J	CPU	α^*	$J^h(\alpha^*)$ (95% c.i.)	
1	*A-C-D-C-A*	35	38.05	2.6520	0.0242	0.1	3.0355 ± 0.0016	14.46
2	*D-C-C-A-A*	71	78.70	3.0016	0.0488	0.1	3.4512 ± 0.0015	14.98
3	*D-C-A-A-C*	129	141.21	3.4916	0.0670	0	3.8759 ± 0.0020	11.00
4	*D-A-C-A-C*	129	140.39	3.6572	0.0670	0	3.9348 ± 0.0020	7.59

The results shown in Table 9 are qualitatively similar to those obtained for the 4-grade example. Namely, the smaller the imbalance problem, the smaller the optimal value of α, and the better the performance of the DBH. In all cases of this example, α^* is quite small and $J^h(\alpha)$ is slightly increasing and relatively insensitive to parameter α, at least for values of α smaller than 0.5, as can be seen from Figure 4.

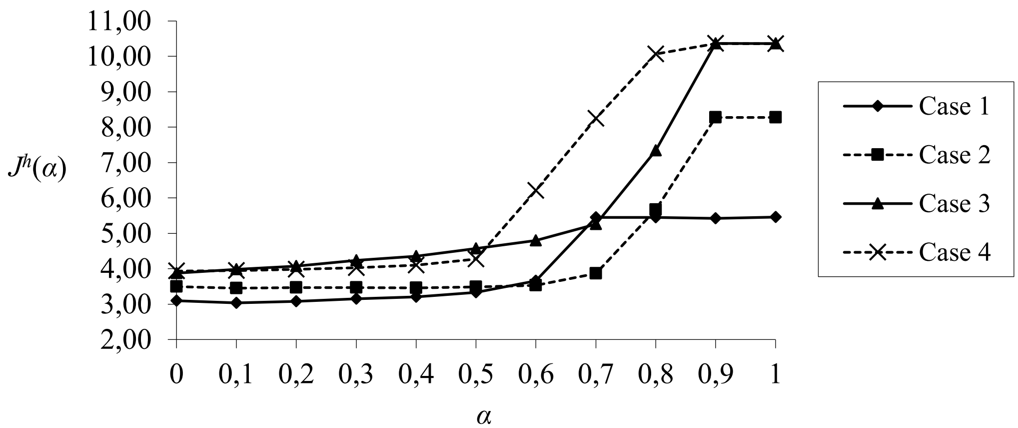

Figure 4: Expected average cost per period of the DBH, $J^h(\alpha)$, vs. α, for the 5-grade example

The cost increase when using the DBH instead of the exact method ranges between 7.59% for case 4 and 14.98% for case 2 and is 12.01% on average, which is practically the same as the average cost difference in the 4-grade example. The DBH method, however, is between 600 and 1700 times faster than the exact method, which is quite significant.

6. Further numerical investigation

In this section, we further investigate the performance of the DBH. First, we compare it against the performances of three simpler parameterized heuristics, in order to better access its value. Then, we compare the performances of the DBH and the exact solution procedures for the case where the FG inventory buffer is not common but consists of a number of separate general-purpose silos capable of storing any grade as long as it is not mixed with any other grade.

6.1 Comparison of the DBH with three simple parameterized heuristics

In Section 5.3, we compared the performance of the DBH solution procedure against the performance of the exact solution procedure, for 4-grade and 5-grade examples. For problems with more than 5 grades, although the DBH solution can be readily obtained, it is practically impossible to compare its performance against that of the exact solution, because it is impossible to even start the value iteration method to find the exact solution, as the state space grows dramatically with the number of grades and there is simply not enough computer memory to store it. For example, for a problem with N = 6 and X = 20, the state space contains 6(20 + 6)!/(20!6!) = 1,381,380 points.

To the best of our knowledge, in the relatively scarce literature on dynamic scheduling approaches to the SELSP, no paper except for Löhndorf and Minner (2012) compares the performance of heuristic solutions against the overall optimal dynamic global lot-sizing optimal policy, precisely because the curse of dimensionality makes it practically impossible to solve for the optimal policy. Löhndorf and Minner (2012) compare the performance of the policies that they develop against the optimal policy for a tractable case with three products. All the other related papers that we are aware of evaluate and compare heuristic lot-sizing policies only.

One such example is Gascon et al. (1994) that compares six different heuristics for the classical SELSP. The classical SELSP differs from our problem in that changeovers are allowed from any item to any other item and not only to neighboring items, production is halted during changeovers, there is a FG inventory holding cost instead of a hard FG inventory storage capacity limit, it is possible to halt production (e.g., if the inventory levels are high), and the production rate is larger than the total expected demand rate. The heuristics that Gascon et al. compare are based on the notion of following economic rotation cycles but with the possibility of temporary departing from these cycles whenever potential stockout situations occur. The results show that for the stationary demand case (they also look at non-stationary demand), when the utilization of the facility is high, the heuristic proposed by Vergin and Lee (1978) yields the smallest average changeover cost at the expense of a relatively high average inventory holding cost.

In the Vergin and Lee heuristic, production changes over to the product with the fewest days of stock on hand or most days of backorders, if that item has fewer than L days of stock on hand, where L is a policy parameter. Otherwise, if the item currently produced does not exceed its maximum absolute inventory level (a policy parameter) or its maximum relative inventory level (a computed parameter), then the production of that item continues in the next period. If none of these three conditions hold, then the production facility is idled in the next period.

In what follows, we propose three simple, practical, parameterized heuristics that are inspired by the Vergin and Lee heuristic. We then evaluate the performances of these heuristics for the 4-grade and 5-grade examples that were presented in Section 5.3, and we compare these performances against those of the DBH and the exact solution procedures. All three heuristics are global in that the changeover policy that they prescribe depends on the entire state of the system, $\mathbf{y} = (s, x_1, \ldots, x_N)$; in other words, the setup state in the next period, s', is a function of the system state $\mathbf{y}$ in the current period. As there is no inventory carrying cost and no possibility to change the production rate, in all the heuristics, there are no conditions for idling the facility, and the changeover decision is based only on a condition that detects potential stockout situations.

Minimum individual coverage heuristic (MICH)

Step 1: Given $\mathbf{y} = (s, x_1, \ldots, x_N)$, compute $n^* = \arg\min_n \{x_n / E[D_n]\}$

Step 2: If $x_{n^*} / E[D_{n^*}] \le L$ then $s' = s + \operatorname{sign}(n^* - s)$

The MICH first finds the individual grade with the smallest expected number of periods that can be covered from on hand stock, n^*. If the expected number of periods that can be covered from on hand stock for that grade is less than or equal to a threshold value, L, then the facility changes over to the next neighboring grade in the direction of n^*, aiming to eventually reach n^*. Of course, if $n^* = s$, then there is no changeover. L is a policy parameter which acts as a safety time (analogous to a safety stock).

Because of the changeover-to-neighbor-only restriction, a seemingly more suitable variant of the MICH would be to replace the actual expected number of periods that can be covered from on hand stock for grade n, $x_n/\mathrm{E}[D_n]$, with the "effective" expected number of periods that can be covered from on hand stock, where the latter is defined as the actual expected number of periods that can be covered by stock on hand minus the number of periods needed to start producing grade n, $|n - s|$. In other words, the variant replaces $x_n/\mathrm{E}[D_n]$ with $x_n/\mathrm{E}[D_n] - |n - s|$ in the rule. We tried this variant, but it performed very poorly, as it produced a changeover policy where production was limited to the middle grades only. The reason is that the variant seems to embark to changeover to the most extreme grade, due to the $|n - s|$ term, but never reaches that grade, because the extreme grade changes during the course.

The MICH is one of the simplest heuristics that one may think of. A potential problem with it is that it considers the individual grades in isolation without taking into account which side of the current grade (low or high) they are in. Thus, it may happen that the individual grade with the smallest expected number of periods that can be covered from on hand stock is on the low side of the current grade, but there are many other grades on the high side with slightly higher but still quite small expected number of periods that can be covered from on hand stock. In such a case, it might be better to changeover to the next high grade than to the next low grade. The next heuristic is designed to address this issue.

Minimum average aggregate coverage heuristic (MAACH)

Step 1: Given $\mathbf{y} = (s, x_1, \ldots, x_N)$, compute $m_l = \frac{1}{s-1}\sum_{n=1}^{s-1}\frac{x_n}{E[D_n]}$, $m_s = \frac{x_s}{E[D_s]}$, and

$$m_h = \frac{1}{N-s}\sum_{n=s+1}^{N}\frac{x_n}{E[D_n]}$$

Step 2: Compute $m_{\min} = \min\{m_l, m_s, m_h\}$

Step 3: If $m_{\min} \le L$ then $s' = s + \mathrm{sign}(m_l - m_h)\cdot\mathrm{sign}(m_s - m_{\min})$

The MAACH first finds the average expected number of periods that can be covered by the aggregate on-hand inventory of the low, current, and high grades, where the low and high grades are defined as in the DBH. Then it computes the minimum of the three averages. Note that if the current grade is an extreme grade (1 or N), then either the low or high grade term does not exist, so it is omitted from the minimization. If the minimum of the averages is less than or equal to a threshold value, L, then the facility changes over in the direction of the grade that this minimum corresponds to (low, current, or high). Of course, if this grade is s, then there is no changeover.

The MAACH takes into account the fact that the candidate for changeover grades are on one or the other side (low or high) of the current grade; however, it may suffer from the "imbalance problem" that we described in Section 4, because it may happen that average expected number of periods that can be covered by the aggregate on-hand inventory of a composite grade (e.g., the low grade) is relatively high, signaling that there is no need to changeover towards that composite, but the expected number of periods that can be covered by one of the components of the composite grade is small. The next heuristic is designed to address this issue.

Maximum average aggregate shortfall heuristic (MAASH)

Step 1: Given $\mathbf{y} = (s, x_1, \ldots, x_N)$, compute $m_l = \frac{1}{s-1}\sum_{n=1}^{s-1}\left(L - \frac{x_n}{E[D_n]}\right)^+$, $m_s = \left(L - \frac{x_s}{E[D_s]}\right)^+$,

and $m_h = \frac{1}{N-s}\sum_{n=s+1}^{N}\left(L - \frac{x_n}{E[D_n]}\right)^+$

Step 2: Compute $m_{\max} = \max\left\{m_l, m_s, m_h\right\}$

Step 3: $s' = s + \mathrm{sign}(m_h - m_l)\cdot\mathrm{sign}(m_{\max} - m_s)$

The MAASH first finds the average expected shortfall of the number of periods that can be covered by the aggregate on-hand inventory of the low, current, and high grades, from a threshold value, L. Then it computes the maximum of the three averages. Note that if the current grade is an extreme grade (1 or N), then either the low or high grade term does not exist, so it is omitted from the maximization. Also note that if a term is zero, then it should also be omitted from the maximization. The facility changes over in the direction of the grade that this maximum corresponds to (low, current, or high). Of course, if this grade is s or if all terms are omitted from the maximization, then there is no changeover. Note that when L = 1, the expressions for the average expected shortfalls of the low and high grades, m_l and m_h, resemble expression (6) used in the DBH, except that (6) refers to the total expected lost sales expressed in product units, whereas the expressions for m_l and m_h, refer to the average expected lost sales (when L = 1) expressed in periods.

All the heuristics presented above are practical and easy-to-use. The changeover decision that they prescribe can be computed on the spot for any given $\mathbf{y}$, and does not have to be precalculated as is the case with the DBH. On the downside, all three heuristics require the sizing of the threshold parameter L.

Tables 10 and 11 show the expected average cost of the exact procedure, the DBH, and the three parameterized heuristics, for the 4-grade and 5-grade examples presented in Section 5.3. The optimal value of L, denoted L^*, is shown in parentheses, for the three parameterized heuristics. L^* was found by successively trying different values of L in increments of one, starting from $L = 0$, until an increase in the expected average cost was observed. The percent cost increase of the heuristics with respect to the exact procedure is shown in square brackets. The results for the three heuristics were obtained by simulation, whereas those for the exact method and the DBH were taken from Tables 8 and 9, respectively.

Table 10: Expected average cost of the exact procedure, the DBH, and the three parameterized heuristics for the 4-grade example for $X = 30$

Case	Demand pattern	Exact	DBH [% cost increase]	MICH (L^*) [% cost increase]	MAACH (L^*) [% cost increase]	MAASH (L^*) [% cost increase]
1	*B-D-D-B*	1.0034	1.2442 [24.00]	1.1652 (4) [16.13]	1.3420 (6) [33.75]	1.1988 (4) [19.47]
2	*D-D-B-B*	1.0927	1.2253 [12.13]	1.8309 (3) [67.56]	1.7019 (6) [55.75]	1.5580 (5) [42.58]
3	*D-B-D-B*	1.1835	1.3207 [11.59]	4.5802 (1) [287.00]	1.6064 (5) [35.73]	1.5861 (6) [34.02]
4	*D-B-B-D*	1.2881	1.3139 [2.00]	4.3619 (1) [238.63]	1.8826 (4) [46.15]	1.8801 (10) [45.96]

Table 11: Expected average cost of the exact procedure, the DBH, and the three parameterized heuristics for the 5-grade example for $X = 20$

Case	Demand pattern	Exact	DBH [% cost increase]	MICH (L^*) [% cost increase]	MAACH (L^*) [% cost increase]	MAASH (L^*) [% cost increase]
1	*A-C-D-C-A*	2.6520	3.0355 [14.46]	3.7225 (4) [40.37]	4.0426 (5) [52.44]	3.8389 (1) [44.75]
2	*D-C-C-A-A*	3.0016	3.4512 [14.98]	4.2896 (3) [42.91]	4.3290 (5) [44.22]	4.1798 (2) [39.25]
3	*D-C-A-A-C*	3.4916	3.8759 [11.01]	5.5017 (2) [57.57]	4.4721 (4) [28.08]	4.4376 (6) [27.09]
4	*D-A-C-A-C*	3.6572	3.9348 [7.590]	6.7949 (3) [85.80]	4.3932 (5) [20.12]	4.3770 (7) [19.68]

From the results, it can be seen that the DBH significantly outperforms the three heuristics in all cases except case 1 of the 4-grade example, where the DBH has the poorest performance and the MICH and MAASH incur slightly smaller costs. More specifically, for the 4-grade example, the average cost increase (with respect to the optimal solution) for the MICH, MAACH, and MAASH is 152.33%, 42.85%, and 35.51%, respectively, whereas the average cost increase for the DBH is 12,42%. For the 5-grade example, the average cost increase for the three heuristics is 56.66%, 36.22%, and 32.70%, respectively, whereas the average cost increase for the DBH is

12,01%. This comparison shows that the complexity of the changeover decision problem, which is due to a large extent to the changeover-to-neighbor-only restriction and the inability to change the production rate, cannot be effectively tackled with simple rules, even if these rules are global and use an optimized threshold parameter to hedge against future uncertainty. The DBH seems to respond to this complexity much more adequately. Among the three heuristics, the MAASH performs slightly better than the MAACH, which itself performs much better than the MICH, except for case 1, where the MICH outperforms all other heuristics.

6.2 Investigation of the case of separate general-purpose FG silos

Thus far, we have assumed that the FG inventory buffer is common and therefore each grade can be stored in any discrete quantity as long as the total inventory over all grades does not exceed the common buffer capacity X. In this section, we assume that grades are stored separately in M general-purpose, equal-sized FG silos, i.e., storage spaces that can store any grade as long as it is not mixed with another grade. The total storage capacity is still X, which means that each silo has a capacity of X/M. Figure 5 illustrates the inventory state space for a 2-grade system for different values of M. As can be seen, the number of silos affects the outer facet of the inventory storage state space. Clearly, the larger M, the more flexible the usage of the total storage capacity. The right-most drawing in Figure 5 shows the inventory state space for the limiting case where $M = X$, which represents the case of the single common FG inventory buffer that we have considered thus far in the previous sections.

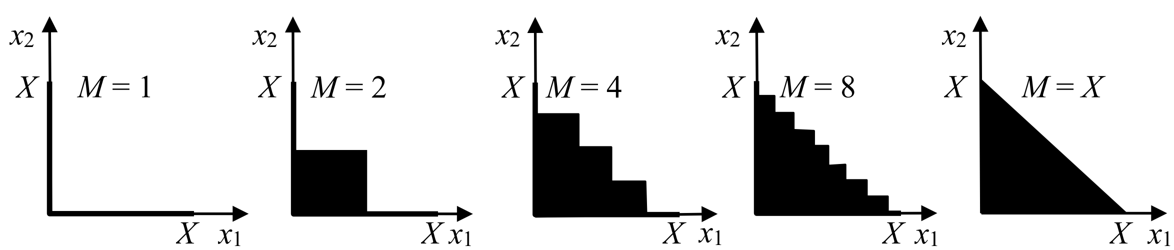

Figure 5: Inventory state space for a 2-grade system for different values of M

When $M < X$, the definitions of the inventory state space and available storage capacity must be modified. More specifically, the set of allowable inventory levels is determined by all integers x_n, $n = 1, \ldots, N$, such that

$$0 \leq \sum_n \left\lceil \frac{x_n}{X/M} \right\rceil \leq M \tag{11}$$

where $\lceil x \rceil$ symbolizes the ceiling of x, i.e., the smallest integer which is not smaller than x. Note that in the limiting case where $M = X$, expression (11) becomes identical to (1).

Moreover, when the state of the system at the beginning of a period is $\mathbf{y}$, where $\mathbf{y} \equiv (s, x_1, \ldots, x_N)$, the amount of material that is added to storage after the facility produces P units, minus any spillage, and before the demand is realized is given by

$$p(\mathbf{y}) \equiv \min\left(P, X - x_s - \left(\sum\nolimits_{n \neq s} \left\lceil \frac{x_n}{X/M} \right\rceil \right) X/M \right) \tag{12}$$

Again note that in the limiting case where $M = X$, expression (12) becomes identical to (2).

The rest of the expressions that we developed for the case where $M = X$ remain the same when $M < X$. It should be pointed out, however, that an additional layer of approximation is introduced in the DBH procedure that we developed in Section 4, regarding the definitions of the aggregate inventory state space and available storage capacity of the composite grades. This additional approximation stems from the fact that in each 3-grade sub-problem S_n, each composite grade (low and high) is treated as a single grade and can therefore be stored in a single silo, whereas in the original N-grade problem S, it consists of many individual grades which have to be stored in separate silos.

To explore the impact of M on the performance of the optimal policy and on the relative performance of the DBH policy, we reran the 4-grade and 5-grade examples that were presented in Section 5.3 for different values of M. The results are shown in Tables 12 and 13. The first line in each case shows the results for $M = X$ and is taken from Tables 8 and 9, respectively, for the two examples.

From the results, it can be seen that in both examples, as M decreases, the expected average cost per period (both for the optimal and the DBH policy) increases significantly (70-110% for the 4-grade example and 30-40% for the 5-grade example). This is expected, because as was mentioned earlier, the larger M, the more flexible the usage of the total storage capacity. It should be noted, however, that the case where the number of silos is approximately equal to the number of grades (M = 5) is rather extreme for practical purposes. At the same time, as M decreases, the CPU to generate the optimal policy decreases, mainly because the state space decreases in M.

Table 12: Performance of the exact and DBH procedures for the 4-grade example for $X = 30$ and different values of M

Case	Demand pattern	M	Exact k_c	Exact CPU	Exact J	Heuristic CPU	Heuristic α^*	Heuristic $J^h(\alpha^*)$	% cost difference
1	B,D,D,B	30	55	8.57	1.0034	0.0461	0.7	1.2442	24.00
		15	49	6.37	1.1018	0.0461	0.7	1.3439	21.97
		10	44	4.71	1.2280	0.0372	0.7	1.4913	21.44
		5	44	2.54	1.7191	0.0214	0.5	2.1092	22.69
2	D,D,B,B	30	156	24.29	1.0927	0.0574	0.5	1.2253	12.13
		15	147	19.04	1.2194	0.0597	0.1	1.3342	9.41
		10	125	13.87	1.3876	0.0537	0.1	1.4976	7.93
		5	68	3.92	2.0260	0.0308	0	2.4463	20.75
3	D,B,D,B	30	187	29.12	1.1835	0.0748	0.1	1.3207	11.59
		15	158	20.42	1.3433	0.0761	0.1	1.4755	9.84
		10	154	16.52	1.5472	0.0664	0.1	1.7005	9.91
		5	56	3.22	2.1962	0.0403	0	2.8479	29.67
4	D,B,B,D	30	110	17.13	1.2881	0.1040	0.1	1.3139	2.00
		15	103	13.25	1.4664	0.1068	0	1.4881	1.48
		10	92	9.86	1.6902	0.0964	0	1.7162	1.54
		5	59	3.39	2.5215	0.0598	0	2.7484	9.00

Table 13: Performance of the exact and DBH procedures for the 5-grade example for $X = 20$ and different values of M

Case	Demand pattern	M	Exact k_c	Exact CPU	Exact J	Heuristic CPU	Heuristic α^*	Heuristic $J^h(\alpha^*)$	% cost increase
1	A,C,D,C,A	20	35	38.05	2.6520	0.0242	0.1	3.0355 ± 0.0016	14.46
		10	38	26.36	2.9816	0.0187	0.1	3.3421 ± 0.0016	11.21
		5	39	11.58	3.7019	0.0110	0.1	4.0832 ± 0.0016	11.03
2	D,C,C,A,A	20	71	78.70	3.0016	0.0488	0.1	3.4512 ± 0.0015	14.98
		10	65	45.30	3.2787	0.0383	0.1	3.6564 ± 0.0019	11.52
		5	50	14.80	3.9435	0.0217	0.1	4.4732 ± 0.0021	11.34
3	D,C,A,A,C	20	129	141.21	3.4916	0.0670	0	3.8759 ± 0.0020	11.01
		10	102	71.39	3.8277	0.0532	0	4.3633 ± 0.0019	13.99
		5	44	13.06	4.5761	0.0351	0	5.3384 ± 0.0020	16.66
4	D,A,C,A,C	20	129	140.39	3.6572	0.0670	0	3.9348 ± 0.0020	7.59
		10	101	71.12	4.0060	0.0547	0	4.3579 ± 0.0018	8.78
		5	54	16.58	4.7827	0.0342	0	5.2690 ± 0.0025	10.17

Finally, in the 4-grade example, the cost increase, when using the DBH instead of the exact method, slightly drops as M decreases, except for the extreme sub-case where M is very small ($M = 5$). In the 5-grade example, the cost increase drops as M decreases in the cases where it is relatively large (cases 1 and 2), whereas it increases in the cases where it is relatively small (cases 3 and 4), but the differences are not that significant.

7. Conclusions

We studied a new version of the SELSP, for which we developed an MDP model. For problems with 2 and 3 grades, we numerically solved the MDP problem and obtained useful insight into the structure of the optimal changeover policy.

For problems with N grades, $N > 3$, we developed a heuristic solution procedure, called DBH, which is based on decomposing the original multi-grade problem into $(N-2)$ 3-grade sub-problems, numerically solving each sub-problem, and constructing the final policy for the original problem by combining parts of the sub-problem optimal solutions. We tested the DBH for problems with 4 and 5 grades. For the 4-grade examples, the DBH procedure was 160-420 times faster than the numerical procedure for solving the exact problem and the DBH solution performed on average 12.42% worse than the exact solution. For the 5-grade examples, the DBH procedure was 600-1700 times faster than the numerical procedure for solving the exact problem and the DBH solution performed on average 12.75% worse than the exact solution. The fact that the performance of the DBH solution is more or less the same for the 4-grade and 5-grade problems is an encouraging sign for problems with more than 5 grades. The numerical results showed that the DBH is best fit when the dispersion of the total expected demand among the individual grades is high. The DBH performed significantly better than three simpler dynamic heuristic rules that we tested. Moreover, its performance did not seem to change significantly when we assumed that the grades are stored in several separate, general-purpose, equal-sized FG silos, as opposed to a common FG buffer.

The DBH has a lot of room for improvement in terms of its computational performance, if one enhances the performance of its main building block, which is the value iteration algorithm for solving 3-grade problems. In a recent work, Kalantzis (2012) compares the computational performance of the standard value iteration algorithm against a modified algorithm, called minimum difference criterion value iteration (Herzberg and Yechiali, 1994), which is based on applying relaxation at the value functions in each iteration to accelerate the standard algorithm. For the 2-grade SELSP, he also develops an action elimination procedure which he uses within the modified algorithm. He reports significant computational savings for the 2-grade examples, and he finds that the modified value iteration compared to the standard value iteration can save on average 35% and 18% of CPU time for 3-grade and 4-grade examples, respectively.

A possible direction for future research would be to try to develop a better heuristic that somehow uses the optimal differential cost functions of the sub-problems, although we should point out that our initial experimentation with this possibility has not been encouraging. An exact approach, with a less fine discretized buffer space might also be worth looking into, although we should caution that such an approach may result in distorting significantly key problem parameters, such as the production and demand rates, and therefore end up solving a problem which is different than the real problem.

We believe that one of the contributions of this manuscript is that it introduces a new variant of the SELSP encountered in process industries, namely that of a non-stop multi-grade production facility with sequence-restricted setup changeovers. Our hope is that it may generate some interest in the research community which may lead to the development of new and possibly better heuristic solutions. As can be seen from Figure 2 (right), the optimal policy even for a 3-grade problem can be quite complicated, so we expect that developing an efficient heuristic that performs much better than the one presented in this manuscript will be challenging.

Acknowledgements

The work in this paper was supported by grants MIS 03ED913 “Optimization of production planning and grade distribution of a PET resin chemical plant” and MIS 379526 “ODYSSEUS: A holistic approach for managing variability in contemporary global supply chain networks,” which were co-financed by the European Union (European Social Fund) and Greek national funds through the Operational Programs “Competitiveness” of the 3rd Support Framework of the European Community (Research Funding Program: PENED – Reinforcement program of human research manpower) and “Education and Lifelong Learning” of the National Strategic Reference Framework (Research Funding Program: THALES – Reinforcement of the interdisciplinary and/or inter-institutional research and innovation), respectively. The first grant was also co-financed by private funds from Artenius Hellas S.A. PET Industry.

References

Bruin, J., van der Wal, J. 2010a. A cyclic production scheme for muti-item production systems with backlog; Part 1. Report 2010-018, EURANDOM, Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands.

Bruin, J., van der Wal, J. 2010b. A dynamic control strategy for multi-item production systems with backlog; Part 2. Report 2010-018, EURANDOM, Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands.

Cooke, D.L., T.R. Rohleder. 2006. Inventory evaluation and product slate management in large-scale continuous process industries. Journal of Operations Management 24, 235-249.

Elhafsi, M., Bai, S.X. 1997. Optimal and near-optimal control of a two-part-type stochastic manufacturing system with dynamic setups. Production and Operations Management 6 (4) 419-438.

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

Gascon A., R.C. Leachman, P. Lefrancois. 1994. Multi-item, single-machine scheduling problem with stochastic demands: a comparison of heuristics. International Journal of Production Research 32 (3) 583-596.

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

Hatzikonstantinou, O. 2009. *Production Scheduling Optimization in a PET Resin Chemical Industry*. PhD Dissertation, Department of Mechanical Engineering, University of Thessaly, Greece.

Herzberg, M., Yechiali U. 1994. Accelerating procedures of the value iteration algorithm for discounted Markov decision processes, based on a one-step look-ahead analysis. Operations Research Letters. 42 (5) 940-946.

Karalli, S.M., Flowers, A.D. 2006. The multiple-family ELSP with safety stocks. Operations Research 54 (3) 523-531.

Kalantzis, G., 2012. Accelerating the Value Iteration Algorithm on the Stochastic Economic Lot Scheduling Problem for Continuous Multi-Grade Production. M.Sc. Thesis in Econometrics and Management Science with Specialization in Operational Research and Quantitative Logistics, Erasmus University, The Netherlands.

Karmarkar, U.S., Yoo, J. 1994. The stochastic dynamic product cycling problem. European Journal of Operational Research 73 (2), 360-373.

Liberopoulos, G., Kozanidis, G. Hatzikonstantinou, O. 2009. Production scheduling of a multi-grade PET resin plant. Computers and Chemical Engineering 34 (3), 387-400.

Liberopoulos, G., Caramanis, M. 1997. Numerical investigation of optimal policies for production flow control and set-up scheduling: lessons from two-part-type failure prone FMSs. International Journal of Production Research 35 (8), 2109-2133.

Löhndorf, N., Minner, S. 2012. Simulation optimization for the stochastic economic lot scheduling problem. IIE Transactions doi: 10.1080/0740817X.2012.662310.

Löhndorf, N., Riel, M., Minner, S. 2012. Simulation optimization for the stochastic economic lot scheduling problem with sequence-dependent setup times. Optimization Online http://www.optimization-online.org/DB_HTML/2012/10/3661.html.

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

Sharifnia, A., Caramanis, M., Gershwin, S.B. 1991. Dynamic setup scheduling and flow control in manufacturing systems. Discrete Event Dynamic Systems 1 (2) 149-175.

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

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

Tousain, R.L., O.H. Bosgra. 2006. Market-oriented scheduling and economic optimization of continuous multi-grade chemical processes. Journal of Process Control 16 (3) 291-301.

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

Vergin, R.C., Lee, T.N. 1978. Scheduling rules for the multiple product single-machine system with stochastic demand. INFOR 16 (1), 64-73.

Winands, E.M.M., Adan, I.J.B.F., van Houtum, G.J. 2011. The stochastic economic lot scheduling problem: a survey. European Journal of Operational Research 210 (1) 1-9.