MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 6, No. 3, Summer 2004, pp. 253–269
ISSN 1523-4614 | EISSN 1526-5498 | 04 | 0603 | 0253
DOI 10.1287/msom.1040.0045

Campaign Planning and Scheduling for Multiproduct Batch Operations with Applications to the Food-Processing Industry

Kumar Rajaram, Uday S. Karmarkar
Decisions, Operations and Technology Management, The Anderson School at UCLA, 110 Westwood Plaza, Los Angeles, California 90095-1481 {kumar.rajaram@anderson.ucla.edu, uday.karmarkar@anderson.ucla.edu}

We analyze planning and scheduling of multiproduct batch operations in the food-processing industry. Such operations are encountered in many applications including manufacturing of sorbitol, modified starches, and specialty sugars. Unlike discrete manufacturing, batch sizes in these operations cannot be set arbitrarily, but are often determined by equipment size. Multiple batches of the same product are often run sequentially in "campaigns" to minimize setup and quality costs.

We consider a multiproduct, single-stage, single-equipment batch-processing scheme and address the problem of determining the timing and duration of product campaigns to minimize average setup, quality, and inventory holding costs over a horizon. We formulate the deterministic, static version of this problem over an infinite horizon. We show that, in general, a feasible, finite, cyclic solution may not exist. We provide sufficient conditions for the existence of a finite cycle, use single-product problems to provide lower bounds on the costs for the multiproduct problem, and use them to test heuristics developed for this problem. Next, we modify this formulation to incorporate fixed cycles that may be necessary due to factors such as product obsolescence, perishability, or contracts with customers. We do this by allowing for disposal of excess stock so that finite cycles are always feasible, though they might not be optimal; we also develop bounds and heuristic solution procedures for this case. These methods are applied to data from a leading food-processing company. Our results suggest that our methods could potentially reduce total annual costs by about 7.7%, translating to an annual savings of around $7 million.

Key words: process industries; batch reactor scheduling; economic lot scheduling problem (ELSP); heuristics
History: Received: October 21, 2002; accepted: April 6, 2004. This paper was with the authors 7 months for 6 revisions.

1. Introduction

Batch operations are used to produce a wide variety of products in the food-processing industry, such as sorbitol, modified starches, and specialty sugars. In addition, batch operations are prevalent across a wide spectrum of process industries. Examples in these industries include autoclaves or stirred tank reactors in basic and specialty chemical processes, batch distillation in refining, kilns and furnaces in ceramics, hot metal production in steel-making, and lens grinding in optics manufacturing to name just a few.

Batch operations differ significantly from discrete manufacturing. The most significant difference from the planning and scheduling perspective is that batch sizes cannot be freely chosen with batch operations. Typically, the size of the equipment (reactor or processor) either will determine or significantly constrain batch size. Because reactors are standard pieces of equipment used for many different products, reactor sizes and optimal production quantities do not necessarily correspond for all of the products that can be produced on a given reactor. Consequently, several batches of a particular product may be run sequentially in a so-called "campaign" so as to produce an appropriate quantity, while avoiding part or all of the set-up costs incurred when switching between different products. In addition, with increasing campaign length, more batches can be pooled together, which reduces variance in conformance quality and associated rework costs. However, sufficient inventory of other products needs to be built up during this campaign to meet demand, which results in holding costs.

Finally, in certain applications, product disposal costs are incurred if production is far in excess of demand due to batch size and quality requirements.

This paper has been motivated by a study of the batch operations at a large food-processing company that manufactures wheat- and starch-based products such as glucose, sorbitol, modified starches, and gluten, with annual sales exceeding $2 billion. These products are used extensively as components in other food-processing industries (e.g., breweries, confectioneries, and bakeries), consumer product industries (e.g., cosmetics and toothpaste), as well as the paper, pharmaceutical, textile, and specialty chemical industries.

To produce these products, this company operates several industrial-scale batch processes. Because building these processes requires major capital investment, it is crucial that they constantly produce high volumes of output at the correct level of quality. To achieve this goal, these processes have high degrees of automation, are operated continuously, produce one product at a time, and are usually shutdown only a few times a year for scheduled maintenance. As product changeovers result in long downtimes and considerable setup costs, products are often produced in long campaigns and are inventoried. Reductions in downtimes and in costs associated with campaign switchovers, holding, and quality rework are critical at this company in particular and in the food-processing industry in general. This is because their products are typically commodities with market-defined prices, and profits can be increased only by reducing costs and by increasing output by minimizing downtimes.

The current procedures for campaign planning and scheduling at these processes rely on producing products in cycles. Here, the sequence of product campaigns in a cycle is chosen to minimize average campaign setup costs, whereas the campaign length of a product in a cycle is chosen so that the cycle production quantities are in the ratio of total aggregate product demand. While this procedure was easy to understand and implement, this company believed that downtimes and costs could be reduced further by using advanced analytical models.

The analytical models developed in this paper are based on the following two examples of batch operations used by this company. These examples were chosen because the products produced by these batch operations accounted for a significant portion of total profits at this company. The first example deals with the production of a variety of grades of sorbitol, which is used to produce several commodity-type products in the cosmetics and pharmaceutical industries. These grades are produced in a single-stage, single-equipment batch-processing scheme, where the stage and equipment correspond to a single batch reactor. The level of a key attribute distinguishes the grades of sorbitol and, typically, the duration of the reaction sets the level of the distinguishing attribute. There are significant setup costs and time associated with the start of a new campaign because the reactor has to be cleaned and, many times, reagents have to be changed, depending on the type of grade. However, there are no setup costs and time across batches in the same campaign. In addition, the probability of conformance of a product to meet the specified range of attribute level is increasing for the duration of the campaign. This is because more batches can be pooled together, thus reducing variability in this attribute and in associated rework costs. However, the inventory of the other grades must be sufficient to cover downstream demand, which is known and stable. The second example deals with manufacturing of modified starches. In this case, in addition to the characteristics of the previous example, the length of the production cycle is fixed by customer contracts, as these products have very specialized applications. We modify our model to consider this case.

The problem we study has similarities to the classic economic order quantity (EOQ) formulation. The multiproduct, static-deterministic EOQ problem with continuous batch sizes, also known as the economic lot scheduling problem (ELSP), has been studied extensively. Elmaghraby (1978) provides a comprehensive survey of the research on this problem. The ELSP has been shown to be NP-hard (Hsu 1983). Consequently, an effective method for computing the optimal solution to the general problem does not exist. The ELSP with discrete time periods, variable demand, and sequence-dependant setups is known as the "product-cycling" problem. Karmarkar and Schrage (1985) consider the product-cycling problem under deterministic variable demand, while Rajaram and Karmarkar (2002) consider the product cycling problem under yield uncertainty.

The ELSP with integer batch sizes addressing the scheduling and planning of batch operations has a considerably smaller amount of literature. Reklaitis (1992) provides a comprehensive review of the literature in scheduling and planning of batch operations in the chemical processing industry. More recent work in this area includes Castro et al. (2001), Lin et al. (2002), and Orcun et al. (2001). Generally, these works deal with developing detailed mathematical programming formulations of this problem in specific applications and simplifying these formulations so that the problem is amenable to solution by commercially available software. Further, when heuristics are used to solve this problem, lower bounds are seldom provided to test the quality of these heuristics. Other research on this problem includes Dessouky and Kijowski (1997) who model the multiproduct batch scheduling problem as a mixed-integer program over a finite horizon. They do not explicitly consider setup times and employ a uniform discretization of time in which each unit of time is one production shift and all products have a processing time of one production shift. While these assumptions are rather restrictive in our application to the food-processing industry, they seem realistic in the application to a pesticide plant considered in that paper and also facilitate a solution to the problem via a polynomial time algorithm. Orcun et al. (1999) develop a mixed-integer programming formulation for the planning and scheduling of batch process plants under uncertain operating conditions and test this model on data from a baker's yeast production plant. They solve this problem using commercially available software and illustrate the impact of uncertainty on production planning and scheduling.

The problem considered here differs from the problems addressed by these papers in several aspects. First, it considers campaign setup times and costs, quality costs, and demand feasibility constraints. These aspects are of practical importance in several applications in the food-processing industry. Second, it explicitly develops a quality model that calculates quality costs as a function of the batches produced in a campaign. Third, much of the literature to date addresses problems that are either of small size or of a structure simple enough to be solved using commercially available math programming software. In contrast, the problem we consider is more complex. In our computational experience, we have observed that powerful commercial software tools cannot generate feasible solutions to even small problems. Consequently, we develop bounds and heuristics, designed to solve large-sized practical problems to near optimality. Finally, to the best of our knowledge, this is the first known investigation of the applicability of these methods to the modified starch- and sorbitol-processing industries. We validate this model using data from a large food-processing company.

This paper is organized as follows. In the next section, we formulate the multiproduct, infinite-horizon problem, examine its feasibility, and construct a Lagrangian relaxation that decomposes this problem to multiple single-product EOQ-type problems with unit batches and quality costs. These EOQ-type problems are solved by rounding down and rounding up the continuous solution to the nearest batch multiple and taking the better answer. We also develop two heuristics that allow us to convert the solution of the relaxed problem into a feasible solution to the multiproduct infinite-horizon problem. In §3, we modify this formulation to incorporate fixed cycles by allowing for disposal of excess stock so that finite cycles are always feasible, though they might not be optimal. We also develop bounds and heuristics for this modified problem. We report computational results in §4. In §5, we describe the application of our methods to data from the food-processing industry. The computational study and the application show that these heuristics and bounds work well. In the concluding section, we summarize our work and suggest future research directions.

2. Model Formulation, Lower Bounds, and Heuristics

Consider a batch-manufacturing facility producing p products and let $i \in I = \{1, \ldots, p\}$ index the set of products. In this facility, we consider the problem of multiproduct, single-stage, single-equipment batch-process planning and scheduling where demand is known and constant. We consider the case in which the batch process has a fixed capacity that must be sequenced in an all-or-nothing manner and in which any batch consists of a single product. We assume that for a given product, this batch process has known and fixed set-up times, production times, and costs of set-up, production, and holding inventory. In this context, we formulate the static, deterministic, average cost, multiproduct campaign planning and scheduling problem for an infinite horizon. The aim is to determine the timing and the duration (in terms of the number of batches) of campaigns on a single batch processor to minimize average setup, holding, and quality costs. These assumptions are consistent with the sorbitol production example described in the introduction. In addition, the static setting is a reasonable model for planning in many food-processing industries where demand may be fairly level because downstream operations are typically run at stable utilization levels. To determine the timing and the duration of the production campaigns at this batch manufacturing facility, we define the following variables.

$T =$ cycle length;
$T_i =$ cycle length for product i;
$n_i =$ number of campaigns of product i that are conducted over T; and
$m_{ij} =$ number of batches of product i in campaign j, where $j \in J = \{1, \ldots, n_i\}$.

We are given:

$D_i =$ demand rate for product i (units/time);
$B =$ batch size of reactor (units);
$t_i =$ processing time per batch of product i (units of time);
$\tau_i =$ setup time per campaign of product i (units of time);
$S_i =$ fixed cost for setting up a campaign for product i ($);
$R_i =$ rework cost per campaign of product i ($); and
$h_i =$ holding cost per unit of product i ($/unit).

In this model, "quality" refers to conformance to product specifications. Quality is affected by the duration of the campaign or, in effect, by the number of batches produced in a campaign. The costs of quality are the expected costs of correcting nonconforming products by reworking. Let $P_{ij}(m_{ij})$ represent the probability of nonconformance of product i in campaign j as a function of m_{ij}. Then, the expected costs of quality of product i in campaign j can be approximated as $R_i P_{ij}(m_{ij})$. We derive the function $P_{ij}(m_{ij})$ under certain assumptions after describing the model.

Define vector $\vec{m}_i = (m_{i1}, m_{i2}, \ldots, m_{in_i})$ and $I_i(\vec{m}_i)$ a scalar function of vector $\vec{m}_i$ representing the cumulative inventory when product i is produced across n_i campaigns in cycle T. We estimate the value of function $I_i(\vec{m}_i)$ later in this section and the total holding costs for product i during cycle T is $h_i I_i(\vec{m}_i)$. In addition, let $\{X\}$ represent occupancy constraints that ensure that no two campaigns overlap on the single batch reactor. We do not explicitly state these constraints as they are relaxed in our solution procedures. Finally, we assume that there is enough capacity to satisfy total demand. The condition for which this assumption holds is established in Proposition 2.

The campaign planning and scheduling problem (CPSP) can be represented by the following nonlinear mixed-integer program.

(CPSP)

$$Z = \text{Min}\left(\frac{\sum_{i=1}^{p} S_i n_i + \sum_{i=1}^{p} h_i I_i(\vec{m}_i) + \sum_{i=1}^{p}\sum_{j=1}^{n_i} R_i P_{ij}(m_{ij})}{T}\right), \tag{1}$$

subject to

$$\sum_{j=1}^{n_i} m_{ij} B = T D_i \quad \forall i, \tag{2}$$

$$T \geq \sum_{i=1}^{p} \tau_i n_i + \sum_{i=1}^{p}\sum_{j=1}^{n_i} t_i m_{ij}, \tag{3}$$

$$\{X\}, \tag{4}$$

$$T \geq 0, \qquad m_{ij}, n_i \in N^+ \quad \forall i. \tag{5}$$

Objective function (1) consists of the average setup, holding, and quality costs over a horizon. Constraint (2) ensures that the total production quantity of any given product during the cycle is equal to its total demand during this cycle. This constraint is essential to prevent infinite accumulation or infinite shortage in the infinite-horizon average cost problem. Constraint (3) ensures that the total cycle is sufficiently long to contain the setup and production times for all of the product campaigns. Reactor occupancy constraints are represented by (4), while nonnegativity and integrality constraints are represented by (5).

To estimate $P_{ij}(m_{ij})$ in the CPSP, we assume that the customer-specified product quality tolerances are specified as a fraction ψ_i of attribute level a_i associated with product i. Thus, the upper specification limit (USL) is $a_i(1+\psi_i)$, and the lower specification limit (LSL) is $a_i(1-\psi_i)$. We assume that errors ε_i associated with a_i are due to the composition of the product and are normally distributed with mean 0 and standard deviation σ_i. Thus, $\tilde{a}_{ijk}$, the distribution of the attribute level of the kth batch of product i in campaign j is also normally distributed with mean a_i and standard deviation σ_i.

Typically, during a production campaign j, m_{ij} batches of product i are mixed together. In addition to sorbitol production, mixing is commonly used in manufacturing products such as beverages, glucose, and modified starches in the food-processing industry. The objective of mixing is to minimize the overall variation of a product attribute by compensating for errors in this attribute in a single batch by using other batches. When batches are mixed, the distribution of attribute level of product i in campaign j is given by $\tilde{a}_{ij} = \sum_{k=1}^{m_{ij}} \tilde{a}_{ijk}/m_{ij}$. Note that $\tilde{a}_{ij}$ is normally distributed with mean a_i and standard deviation $\sigma_i/\sqrt{m_{ij}}$. Let $\Phi(\;)$ denote the cumulative distribution function of the standard normal variate. Then, by definition, $P_{ij}(m_{ij}) = P[\tilde{a}_{ij} \geq \text{USL}] + P[\tilde{a}_{ij} \leq \text{LSL}] = P[\tilde{a}_{ij} \geq a_i(1+\psi_i)] + P[\tilde{a}_{ij} \leq a_i(1-\psi_i)]$. Thus,

$$P_{ij}(m_{ij}) = 1 - \Phi\left(\frac{a_i\psi_i\sqrt{m_{ij}}}{\sigma_i}\right) + \Phi\left(\frac{-a_i\psi_i\sqrt{m_{ij}}}{\sigma_i}\right) = 2\Phi\left(\frac{-a_i\psi_i\sqrt{m_{ij}}}{\sigma_i}\right). \tag{6}$$

Next, we establish some properties of the CPSP.

Proposition 1. *A finite feasible cycle need not exist for the CPSP.*

Proof. Define

$$r_{i_1,i_2} = \frac{\sum_{j=1}^{n_{i_2}} m_{i_2 j}}{\sum_{j=1}^{n_{i_1}} m_{i_1 j}}, \qquad i_1 \neq i_2,\ \forall i_1, i_2 \in I.$$

Because $m_{ij} \in N^+\ \forall i, j$, it also follows that $r_{i_1,i_2} \in Q^+$, where Q^+ is the set of positive rational numbers. Next, from (2) we would require that

$$r_{i_1,i_2} = \frac{\sum_{j=1}^{n_{i_2}} m_{i_2 j}}{\sum_{j=1}^{n_{i_1}} m_{i_1 j}} = \frac{D_{i_1}}{D_{i_2}}.$$

The result follows from choosing $D_{i_1} = \pi W$ and $D_{i_2} = W$ for $W > 0$, such that $D_{i_1}/D_{i_2} = r_{i_1,i_2} \notin Q^+$. □

The proof for Proposition 1 implies that a necessary condition for the existence of a feasible finite cycle is that $D_{i_1}/D_{i_2} \in Q^+$, $i_1 \neq i_2$, $\forall i_1, i_2 \in I$. This result suggests that a finite cyclic schedule is likely to exist in practice and provides a basis for searching for this cyclic schedule from a theoretical and practical point of view. This result also implies that for a feasible finite cycle to exist, we must be able to express $D_i = \alpha_i/(\text{ir})\beta_i$, $\alpha_i, \beta_i \in N^+\ \forall i$ and that the irrational portion ir in this term is common across all products. Let β represent the least common multiple of $(\beta_1, \beta_2, \ldots, \beta_p)$ and $K = (\text{ir})\beta$. Proposition 2 establishes the duration of a feasible finite cycle.

Proposition 2. *Any finite feasible cycle T is of duration θKB, where $\theta \in R^+$ satisfies*

$$\theta \geq \sum_i \frac{\tau_i n_i}{K(B - \sum_i D_i t_i)}.$$

Proof. Consider a feasible cycle T that satisfies (2) so that $\sum_j m_{ij} = TD_i/B$, $\forall i$. Because $D_i = \alpha_i/((ir)\beta_i)$, α_i, $\beta_i \in N^+\ \forall i$ and because we require that $\sum_j m_{ij} \in N^+$, this implies that $T = \theta KB$. For this value of T, note that $\sum_j m_{ij} = \theta KD_i, \forall i$. In addition, for this value of T to represent a feasible cycle, we would require that (3) is satisfied. The result follows by noting that when $T = \theta KB$ and $\sum_j m_{ij} = \theta KD_i$, $\forall i$, (3) is satisfied when $\theta \geq \sum_i \tau_i n_i/(K(B - \sum_i D_i t_i))$. □

Note from Proposition 2 that $B > \sum_i D_i t_i$ for the feasibility of any finite production cycle in the CPSP. This is similar to the feasibility condition for the ELSP as shown by Maxwell (1964) and Dobson (1987), in which the continuous production rate for product i is $r_i = B/t_i$ and the necessary and sufficient condition for feasibility is $1 > \sum_i D_i/r_i$. This also implies that the fraction of time available per cycle for setups is $(1 - \sum_i D_i/r_i)$ and the total time per cycle available for setup is $T(1 - \sum_i D_i/r_i)$. Because the total time required for setups in a cycle in the CPSP is $\sum_i \tau_i n_i$, for feasibility we would require that $T(1 - \sum_i D_i/r_i) \geq \sum_i \tau_i n_i$. By substituting $r_i = B/t_i$ and $T = \theta KB$ in this inequality, we get $\theta \geq \sum_i \tau_i n_i/(K(B - \sum_i D_i t_i))$, the inequality in the statement of Proposition 2.

The CPSP can be considered as an ELSP-type problem with fixed batch sizes and rework costs. Because it is well known that the ELSP is NP-hard (Hsu 1983), it is highly unlikely that we can solve large, real-sized instances of the CPSP to optimality. We confirm this in our numerical experiments and in the application. Consequently, we develop heuristics to address this problem. Next, we construct a decomposition of the CPSP, which is used to develop lower bounds on this problem. These lower bounds are then used to evaluate the performance of the heuristics in the computational study and in the application.

2.1. Problem Decomposition and Lower Bounds

To decompose the CPSP into tractable subproblems, we first rewrite constraint (3) as

$$1 \geq \frac{\sum_{i=1}^{p} \tau_i n_i + \sum_{i=1}^{p} \sum_{j=1}^{n_i} t_i m_{ij}}{T}.$$

We then relax constraint (3) by introducing Lagrange multiplier $\lambda \geq 0$ and drop constraint (4). Finally, we rewrite constraint (2) as $\sum_{j=1}^{n_i} m_{ij}B = T_i D_i \ \forall i$ and $T_i = T \forall i$. This leads to the following relaxed problem.

$$\text{Min} \sum_{i=1}^{p} Z_i^{\text{LB}}(n_i, m_{ij}, T_i),$$

subject to

$$\sum_{j=1}^{n_i} m_{ij} B = T_i D_i \quad \forall i, \tag{A}$$

$$T = T_i \quad \forall i, \tag{B}$$

(5),

where

$$Z_i^{\text{LB}}(n_i, m_{ij}, T_i) = \frac{(S_i + \lambda\tau_i)n_i + h_i I_i(\vec{m}_i) + \sum_{j=1}^{n_i}(R_i P_{ij}(m_{ij}) + \lambda t_i m_{ij})}{T_i}.$$

We next drop constraint (B). This leads to p single-product subproblems. Observe that constraint (A) can be written as $T_i = (\sum_{j=1}^{n_i} m_{ij}B/D_i) \ \forall i$. We substitute this value of T_i in the expression for $Z_i^{\text{LB}}(n_i, m_{ij}, T_i)$ to get the ith single-product subproblem as

$$Z_i^{\text{LB}}(\lambda) = \underset{n_i, m_{ij} \in N^+}{\text{Min}} \left\{ Z_i^{\text{LB}}(n_i, m_{ij}) = \left(\frac{1}{D_i} \sum_{j=1}^{n_i} m_{ij} B \right)^{-1} \cdot \left((S_i + \lambda\tau_i)n_i + h_i I_i(\vec{m}_i) + \sum_{j=1}^{n_i} (R_i P_{ij}(m_{ij}) + \lambda t_i m_{ij}) \right) \right\}, \tag{7}$$

so that $Z^{\text{LB}}(\lambda) = -\lambda + \sum_{i=1}^{p} Z_i^{\text{LB}}(\lambda)$ represents a lower bound on the optimal value of the CPSP. To find a tight lower bound, we solve the Lagrangian dual problem $Z^{\text{LB}} = \text{Max}_{\lambda \geq 0}\{Z^{\text{LB}}(\lambda)\}$. The following proposition simplifies the solution to these subproblems.

Proposition 3. *The optimal solution to the single-product, static, infinite-horizon batch processor subproblem* (7) *has a fixed cycle solution with* $n_i = 1$.

Proof. At any production occasion, the inventory level at the point of completion of the first batch of the campaign must be zero. If not, production could be delayed until this was the case, and costs would have been reduced. Thus, the completion of the first batch of a campaign is always a regeneration point. This implies that each campaign must be of the same length and that each cycle consists of one campaign. □

By Proposition 3, $n_i = 1 \ \forall i$, so that $\vec{m}_i = (m_{i1}) = m_i$ and the optimal policy is a fixed cycle with a single campaign of m_i^* batches. To calculate m_i^*, we need to first estimate $I_i(m_i)$, the average inventory level for product i during a fixed cycle with a single campaign. In this regard, it is useful to consider Figure 1, which represents the inventory level over time during the fixed cycle. The shaded area in this figure represents $I_i(m_i)$.

From Figure 1, we get

$$I_i(m_i) = \frac{t_i}{2} \sum_{k=1}^{m_i - 1} [2kB - (2k-1)t_i D_i] + \frac{[m_i B - (m_i - 1)t_i D_i]^2}{2D_i} = \frac{m_i B[m_i B - (m_i - 1)t_i D_i]}{2D_i}.$$

We use the above expression for $I_i(m_i)$, set $n_i = 1$ and $P_i(m_i)$ from (6) in (7) to get the $(\text{CPSP})_i$

$$Z_i^{\text{LB}}(\lambda) = \underset{m_i \in N^+}{\text{Min}} \left\{ Z_i^{\text{LB}}(m_i) = \frac{(S_i + \lambda\tau_i)D_i}{m_i B} + \frac{h_i}{2}(m_i B - (m_i - 1)t_i D_i) + \frac{2R_i D_i \Phi(-a_i \psi_i \sqrt{m_i}/\sigma_i)}{m_i B} \right\} + \frac{\lambda t_i D_i}{B}. \tag{8}$$

Note that the $(\text{CPSP})_i$ represents the relaxed problem for product i. The difference between this relaxed

Figure 1 Inventory Positions for ith Product in the Infinite Cycle Case

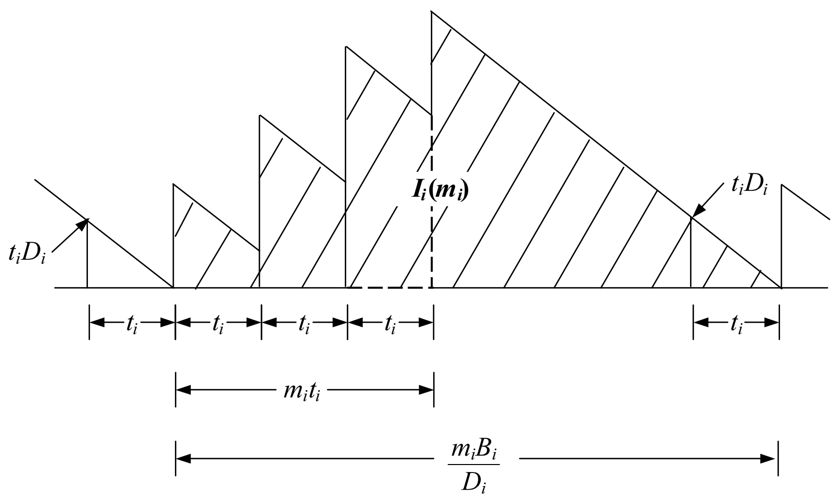

problem and an ELSP relaxation is (a) the different inventory calculation, (b) the quality cost, and (c) the restriction to integer multiples of B.

Next, consider the expression for $Z_i^{\text{LB}}(m_i)$ described in (8). The first term in this expression represents the setup costs per unit time, while the second term represents the holding cost per unit time. This is similar to the cost function of the classical EOQ formulation, modified to incorporate fixed batches (Johnson and Montgomery 1974). The third term represents the quality rework costs per unit time. Thus, the problem of determining $m_i^* = \arg\min_{m_i \in N^+}\{Z_i^{\text{LB}}(m_i)\}$ can be considered to be an EOQ-type problem extended to incorporate quality costs. Next, we develop some results, which are useful in computing m_i^*.

Proposition 4. *$Z_i^{\text{LB}}(m_i)$ is discrete convex in m_i.*

Proof. To show that $Z_i^{\text{LB}}(m_i)$ is discrete convex in m_i, it is sufficient to show that $Z_i^{\text{LB}}(m_i)$ is convex for continuous m_i. To see this result, note that

$$\Phi\left(\frac{-a_i\psi_i\sqrt{m_i}}{\sigma_i}\right) = \int_{-\infty}^{-a_i\psi_i\sqrt{m_i}/\sigma_i} e^{-Z^2/2}\,dz$$

and

$$\frac{d^2Z_i^{\text{LB}}(m_i)}{dm_i^2} = \frac{2(S_i+\lambda\tau_i)D_i}{Bm_i^3} + \frac{4R_iD_i}{Bm_i^3}\int_{-\infty}^{-a_i\psi_i\sqrt{m_i}/\sigma_i} e^{-Z^2/2}\,dz + \frac{10R_iD_ia_i\Psi_i}{4B\sigma_i m_i^{5/2}} e^{-(1/2)(a_i^2\Psi_i^2 m_i/\sigma_i^2)} + \frac{R_iD_ia_i^3\Psi_i^3}{2B\sigma_i^3 m_i^{3/2}} e^{-(1/2)(a_i^2\Psi_i^2 m_i/\sigma_i^2)} \geq 0. \quad \square$$

Let $\tilde{m}_i$ be the optimal continuous solution to the $(\text{CPSP})_i$. Because $Z_i^{\text{LB}}(m_i)$ is convex for continuous m_i, standard search techniques such as the golden section method (Luenberger 1984) are sufficient to compute $\tilde{m}_i$. Then from Proposition 4, it follows that the optimal number of batches in the fixed cycle $m_i^* = \arg\min(Z_i^{\text{LB}}(\lfloor\tilde{m}_i\rfloor), Z_i^{\text{LB}}(\lceil\tilde{m}_i\rceil)$ and $Z_i^{\text{LB}}(\lambda) = Z_i^{\text{LB}}(m_i^*) + \lambda t_iD_i/B$.

To find a tight lower bound, we next consider $Z^{\text{LB}}(\lambda) = -\lambda + \sum_{i=1}^p Z_i^{\text{LB}}(\lambda)$. Observe that in this expression for $Z^{\text{LB}}(\lambda)$, the first term is linear and decreasing in λ, while the second term is linear and increasing in λ. Therefore, we solve for $Z^{\text{LB}} = \text{Max}_{\lambda\geq 0}\{Z^{\text{LB}}(\lambda)\}$ by conducting a simple single-dimensional line search on λ. We use lower bound Z^{LB} to evaluate the quality of the heuristics developed to address this problem, which is described next.

2.2. Upper Bounds and Heuristic Solutions

In general, the solution provided by the lower bound in §2.1 might not be feasible to the CPSP due to the violation of constraints (3) and (4). Because the CPSP is related to the ELSP, it seems plausible that one could use the continuous solution to the ELSP (using one of the many well-known heuristics developed for this problem) and derive a feasible solution to the CPSP by rounding up or down this continuous solution to the nearest batch size multiple. However, rounding the continuous solution is seldom a good solution to the CPSP, because this could lead to an infeasible solution to the CPSP, as illustrated by the example in the appendix. In addition, ELSP heuristics do not consider quality rework costs, which are significant in the food-processing industry.

To ensure feasability and to incorporate the impact of quality rework costs, we develop the following two heuristics that are specially adapted to the needs of the CPSP.

Model-Based Heuristic. In the model-based heuristic, we use the optimal lower bound solution m_i^* to construct a production plan. This heuristic consists of two phases. In the first phase, we determine the total production quantity for each product during the entire cycle, and in the second phase we determine the sequence and production quantity for each campaign in the cycle. The details of each of these phases are as follows.

Phase 1. In this phase, we use m_i^* and compute $T_i^* = m_i^* B/D_i$, $\forall i$. Let $T^*_{(\text{l.c.m.})}$ represent the least common multiple of $(T_1^*, T_2^*, \ldots, T_p^*)$ and $n_i^{(h_1)} = T^*_{(\text{l.c.m.})}/T_i^*$. Let $\gamma^{(h_1)}$ represent the smallest positive integer such that $\gamma^{(h_1)} \geq \sum_i \tau_i n_i^{(h_1)}/(T^*_{(\text{l.c.m.})} - \sum_i t_i n_i^{(h_1)} m_i^*)$. Then, we set the total heuristic production quantity for the ith product during the cycle as $m_i^{(h_1)} = \gamma^{(h_1)} n_i^{(h_1)} m_i^*$, $\forall i$. Here, $\gamma^{(h_1)}$ can be regarded as a scaling factor that ensures the cycle time $T^*_{(\text{l.c.m.})}$ is sufficient to accommodate the campaign setup and production time across all products.

The procedure outlined above is similar to the basic period approach used to construct a solution procedure for the ELSP (Elmaghraby 1978). In this approach, one permits varying cycles as integer multiples of basic periods subject to a feasibility constraint. Here, T_i^* can be regarded as the basic period for product i, while $T^*_{(\text{l.c.m.})}$ can be considered as the varying cycle. Observe that $T^*_{(\text{l.c.m.})}$ is an integer multiple of the basic period and the feasibility of this varying cycle is imposed by $\gamma^{(h_1)} \geq \sum_i \tau_i n_i^{(h_1)}/(T^*_{(\text{l.c.m})} - \sum_i t_i n_i^{(h_1)} m_i^*)$.

Phase 2. To determine the production sequence of these products in a given cycle, we define $n_{\max}^{(h_1)} = \max_i(n_i^{(h_1)})$ and construct $n_{\max}^{(h_1)}$ serial sequence buckets. For product i, we choose $n_i^{(h_1)}$ instances to minimize the inventory-holding costs in a cycle by maximizing the time difference between any adjacent production instances across all production instances and producing $\gamma^{(h_1)} m_i^*$ batches in each production instance. To pick the production instances, define $g_i = \lfloor (n_{\max}^{(h_1)} - n_i^{(h_1)})/n_i^{(h_1)} \rfloor$ and $r_i = (n_{\max}^{(h_1)} - n_i^{(h_1)}) - g_i n_i^{(h_1)}$. We leave a gap of g_i buckets between each of the first $(n_i^{(h_1)} - r_i)$ production instances and a gap of $g_i + 1$ buckets between each of the remaining r_i production instances. This procedure is repeated for all the products and ensures that batches are spread out. Nearly identical procedures have been employed for the ELSP (Dobson 1987). Finally, to make sure that products within a bucket are sequenced in a well-defined and consistent manner, we sequence these products in increasing order of $m_i^* t_i$.

Feasibility Heuristic. In the feasibility heuristic, we use the feasibility conditions of the CPSP, summarized in Proposition 2, to construct a production plan. This heuristic also consists of two phases. The total production quantities for each product during the cycle are determined in the first phase, while the sequence and production quantity for each campaign in the cycle is determined in the second phase. These details of these phases are as follows.

Phase 1. In this phase, we consider (8), set λ to $\lambda^* = \arg\max_{\lambda \geq 0}\{Z^{\text{LB}}(\lambda)\}$ and use Proposition 2 to set $m_i = \theta K D_i$. This results in

$$Z_i = \min_{\theta \in R^+}\left\{ Z_i(\theta) = \frac{(S_i + \lambda^* \tau_i)}{\theta K B} + \frac{h_i}{2}(\theta K D_i B - (\theta K D_i - 1) t_i D_i) + \frac{2 R_i D_i \Phi\big(-a_i \psi_i \sqrt{\theta K D_i}/\sigma_i\big)}{\theta K D_i B} \right\}.$$

It is important to observe that $Z_i^{\text{LB}}(\theta)$ is convex in θ. This follows by applying the implicit function theorem and the proof of Proposition 4 and noting that

$$\frac{d^2 Z_i(\theta)}{d\theta^2} = \frac{d^2 Z_i^{\text{LB}}(m_i)}{dm_i^2} \times \frac{dm_i}{d\theta} + \frac{dZ_i^{\text{LB}}(m_i)}{dm_i} \times \frac{d^2 m_i}{d\theta^2} = \frac{d^2 Z_i^{\text{LB}}(m_i)}{dm_i^2} K D_i \geq 0.$$

Because the sum of convex functions is convex, we use this result and a golden section search to find θ^*, the value of θ that optimizes $Z(\theta) = \sum_i Z_i(\theta)$. Let $n_i^{(h_2)} = \lfloor \theta^* K D_i / m_i^* \rfloor + 1$ and let $\gamma^{(h_2)}$ be the smallest integer such that $\gamma^{(h_2)} \geq \sum_i \tau_i n_i^{(h_2)}/(\theta^* K (B - \sum_i D_i t_i))$. Then we set $m_i^{(h_2)} = \gamma^{(h_2)} n_i^{(h_2)}$, $\forall i$. Here again, $\gamma^{(h_2)}$ can be regarded as a scaling factor that ensures the cycle time $\theta^* KB$ is sufficient to accommodate the campaign setup and production time across all products.

Phase 2. Define $n_{\max}^{(h_2)} = \max_i(n_i^{(h_2)})$ and construct $n_{\max}^{(h_2)}$ serial sequence buckets. For product i, choose $n_i^{(h_2)}$ buckets such that the time difference between any adjacent production instances across all production instances is maximized. Production instances are picked using $n_{\max}^{(h_2)}$ and $n_i^{(h_2)}$ with the same procedure described in the second phase of the model-based heuristic. For all but the last production instance, produce $\gamma^{(h_2)} m_i^*$ batches of product i. In the last production instance, produce $m_i^{(h_2)} - (n_i^{(h_2)} - 1)\gamma^{(h_2)} m_i^*$ batches. This procedure is repeated for all products; within a bucket, products are sequenced in increasing order of $m_i^* t_i$.

3. Incorporation of Fixed Cycles

In certain applications, even if a finite feasible cycle exists, it may not be possible to implement the associated production plan. This occurs when the cycle has to be shortened to accommodate factors such as product obsolescence, perishability, or contracts with customers that are of shorter duration than the cycle. To determine a production plan that incorporates these aspects across all products, it is now necessary to consider a fixed cycle and to dispose of excess product, if needed, so that finite cycles are always feasible, though they might not be optimal. In practice, disposal may be necessary if production exceeds demand due to batch size and quality requirements. However, we have observed in practice that disposal usually takes place after the product is delivered to the customer to deal with potential contingencies during delivery and small increases in customer demand. Examples of fixed cycle production with disposal include the manufacturing of modified starches in the food-processing industry. These products typically have very specialized applications and thus the duration of the production cycles are often limited by customer contracts.

To model this case, we modify the CPSP by letting T_i represent the *fixed* and *given* planning horizon for the product. Let variable d_i denote the disposal quantity of product i and let C_i represent the disposal cost per unit of product i. Typically, C_i is comprised of additional handling and waste water treatment costs associated with product disposal. Here again, we assume that $B > \sum_i D_i t_i$ to guarantee feasibility of the finite cycle production schedule. The finite cycle campaign planning and scheduling problem (FCCPSP) is represented by the following nonlinear mixed-integer program

(FCCPSP)

$$\widetilde{Z} = \mathrm{Min} \sum_{i=1}^{p} \left(\frac{S_i n_i + h_i I_i(\vec{m}_i) + \sum_{j=1}^{n_i} \{R_i P_{ij}(m_{ij})\} + C_i d_i}{T_i} \right), \tag{9}$$

subject to (3), (4), and

$$\sum_{j=1}^{n_i} m_{ij} B = T_i D_i + d_i \quad \forall i, \tag{10}$$

$$d_i \geq 0, \qquad m_{ij}, n_i \in N^+ \quad \forall i. \tag{11}$$

Objective function (9) consists of the average setup, holding, quality, and disposal costs over a fixed horizon. Constraint (10) ensures that the production quantity of any given product during its production cycle is equal to its demand and disposal quantity during this cycle. This constraint is key in distinguishing the FCCPSP from the CPSP, as now we are always guaranteed a finite, feasible cycle because production in excess of demand can be disposed of at a fixed penalty. Nonnegativity and integrality conditions are represented by constraint (11).

3.1. Problem Decomposition and Lower Bounds

To decompose the FCCPSP into tractable subproblems, we relax constraint (3) by introducing Lagrange multiplier $\mu \geq 0$ and drop constraint (4). This leads to p single-product subproblems of the form

$$\widetilde{Z}_i^{\mathrm{LB}}(\mu) = \mathop{\mathrm{Min}}_{d_i, n_i, m_{ij}} \Bigg\{ \widetilde{Z}_i^{\mathrm{LB}}(n_i, m_{ij}, d_i) \frac{1}{T_i} = \Bigg((S_i + \mu \tau_i) n_i + h_i I_i(\vec{m}_i) + \sum_{j=1}^{n_i} \{R_i P_{ij}(m_{ij}) + \mu t_i m_{ij}\} + C_i d_i \Bigg),$$

subject to (10) and (11). (12)

Figure 2 Inventory Positions for *i*th Product in the Finite Cycle Case

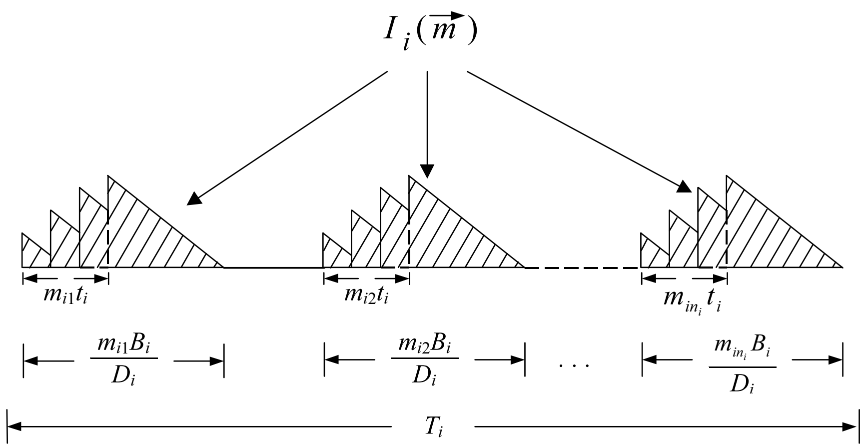

Note that $\widetilde{Z}^{\mathrm{LB}}(\mu)=-\mu+\sum_{i=1}^{p}\widetilde{Z}_i^{\mathrm{LB}}(\mu)$ represents a lower bound on the optimal value of the FCCPSP. To compute a tight lower bound, we solve $\widetilde{Z}^{\mathrm{LB}}=\mathrm{Max}_{\mu\geq 0}\{\widetilde{Z}^{\mathrm{LB}}(\mu)\}$. We next calculate $I_i(\vec{m}_i)$, the average inventory level over the cycle of the product. In this calculation, it is useful to consider Figure 2, which represents the inventory level over time when n_i campaigns of product i are conducted over a fixed cycle of length T_i. In this figure, shaded regions represent the average inventory level for product i during a particular campaign, while the sum of the shaded regions represents the average inventory for this product during the fixed cycle.

We use this figure to compute

$$I_i(\vec{m})=\sum_{j=1}^{n_i}\frac{m_{ij}B(m_{ij}B-(m_{ij}-1)t_iD_i)}{2D_i}. \tag{13}$$

We substitute (6) and (13) into (12) to get the relaxed problem for product i, which we call the $(\mathrm{FCCPSP})_i$

$$\begin{aligned}\widetilde{Z}_i^{\mathrm{LB}}(\mu)&\\=\operatorname*{Min}_{d_i,n_i,m_{ij}}\Bigg[\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)&=\frac{1}{T_i}\Bigg\{(S_i+\mu\tau_i)n_i\\&+\frac{h_i}{2D_i}\sum_{j=1}^{n_i}m_{ij}B(m_{ij}B-(m_{ij}-1)t_iD_i)\\&+\sum_{j=1}^{n_i}\left(2R_i\Phi\left(\frac{-a_i\psi_i\sqrt{m_{ij}}}{\sigma_i}\right)+\mu t_im_{ij}\right)+C_id_i\Bigg\}\Bigg],\end{aligned}$$

subject to (10) and (11). (14)

Proposition 5. *$\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)$ is jointly convex in m_{ij} and d_i.*

Proof. Observe that for continuous m_{ij}

$$\frac{\partial^2\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)}{\partial m_{ij}^2}=2B(B-t_iD_i)+\frac{R_i}{2}e^{-(1/2)(a_i^2\Psi_i^2m_{ij}/\sigma_i^2)}\cdot\left[\frac{a_i\Psi_i}{\sigma_im_{ij}^{3/2}}+\frac{a_i^3\Psi_i^3}{\sigma_i^3m_{ij}^{1/2}}\right]\geq 0,\quad\forall j,$$

$$\frac{\partial^2\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)}{\partial d_i^2}=0,\quad\forall j,$$

$$\frac{\partial^2\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)}{\partial d_i\partial m_{ij}}=\frac{\partial^2\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)}{\partial m_{ij}\partial d_i}=0\quad\forall j.$$

The above results also imply that the Hessian matrix corresponding to $\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)$ is positive semidefinite. This in turn implies that $\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)$ is jointly convex in m_{ij} and d_i. □

For a given n_i, let $\tilde{m}_{ij}$ and d_i^* optimize $\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)$. Proposition 5 ensures that a standard technique such as the conjugate gradient method (Luenberger 1984) is sufficient to compute $\tilde{m}_{ij}$ and d_i^*. Because $\widetilde{Z}_i^{\mathrm{LB}}(n_i,m_{ij},d_i)$ is convex in m_{ij}, it follows that $\widetilde{Z}_i^{\mathrm{LB}}(n_i, m_{ij}, d_i)$ is discrete convex in m_{ij}. Thus, the optimal number of batches of product i in campaign j is $m_{ij}^*=\mathrm{argmin}(\widetilde{Z}_i^{\mathrm{LB}}(n_i, \lfloor\tilde{m}_i\rfloor, d_i^*), \widetilde{Z}_i^{\mathrm{LB}}(n_i, \lceil\tilde{m}_i\rceil, d_i^*))$. This procedure is fast enough that we can find n_i^*, the optimal number of campaigns for product i by enumeration of all values $n_i=1,2\ldots$ subject to (10). Note that in this procedure the largest number of campaigns we need to consider is bounded by $\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij}^*, d_i^*)/S_i$ for any value of n_i for which we have evaluated $\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij}^*, d_i^*)$. This is because using a number of campaigns larger than this would incur setup costs greater than the total cost of a known solution. Finally, $\widetilde{Z}_i^{\text{LB}}(\mu) = \widetilde{Z}_i^{\text{LB}}(n_i^*, m_{ij}^*, d_i^*)$, $\widetilde{Z}^{\text{LB}}(\mu) = -\mu + \sum_i \widetilde{Z}_i^{\text{LB}}(\mu)$ and we solve $\widetilde{Z}^{\text{LB}} = \text{Max}_{\mu \geq 0}\{\widetilde{Z}^{\text{LB}}(\mu)\}$ by conducting a simple single-dimensional line search on μ.

Proposition 6. *Consider the* (*FCCPSP*)$_i$, *the relaxed problem for product* i *as described in Equation* (14). *The optimal solution to the* (*FCCPSP*)$_i$ *satisfies the following conditions.*

(a) *At the start of campaign for product* i, *the inventory level for this product is* $(\tau_i + t_i)D_i$ *so that inventory is zero when the first batch of this campaign is ready.*

(b) *Any sequence of campaigns in a cycle can be arbitrarily reordered without changing the solution.*

(c) $|m_{ij_1} - m_{ij_2}| < 2$, $\forall j_1, j_2 \in J$.

Proof. (a) If this is not true, then one can always postpone the start of the campaign until the point when inventory is $(\tau_i + t_i)D_i$ and reduce inventory-holding costs, contradicting the optimality of the solution.

(b) This result follows by noting that the value of $\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij})$ represented by (14) is invariant to the order of $(m_{i1}, m_{i2}, \ldots, m_{in_i})$.

(c) Let $m_{ij_1}^{(1)} = k$, $m_{ij_1}^{(2)} = k+1$, $m_{ij_2}^{(1)} = k+2$, $m_{ij_2}^{(2)} = k+1$, $f(m_{ij_1}^{(1)}, m_{ij_2}^{(1)}) = \widetilde{Z}_i^{\text{LB}}(n_i, m_{ij_1}^{(1)}, d_i) + \widetilde{Z}_i^{\text{LB}}(n_i, m_{ij_2}^{(1)}, d_i)$ and $f(m_{ij_1}^{(2)}, m_{ij_2}^{(2)}) = \widetilde{Z}_i^{\text{LB}}(n_i, m_{ij_1}^{(2)}, d_i) + \widetilde{Z}_i^{\text{LB}}(n_i, m_{ij_2}^{(2)}, d_i)$. We are required to show that $f(m_{ij_1}^{(2)}, m_{ij_2}^{(2)}) < f(m_{ij_1}^{(1)}, m_{ij_2}^{(1)})$ or $2\widetilde{Z}_i^{\text{LB}}(n_i, k+1, d_i) < \widetilde{Z}_i^{\text{LB}}(n_i, k, d_i) + \widetilde{Z}_i^{\text{LB}}(n_i, k+2, d_i)$. This follows from the discrete convexity of $\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij}, d_i)$ in m_{ij} established by Proposition 5 and the subsequent discussion. □

We next state a regularity condition that simplifies the computation of $\widetilde{Z}^{\text{LB}}$. This condition is based on examining the production of modified starches and other products such as beverages in the food-processing industry, where an additional batch of product is seldom produced purely to reduce rework costs of the entire campaign. This is because the disposal costs of this additional batch typically exceed the reduction in rework costs due to the mixing of an additional batch. In the context of the model, this observation can be represented by the following regularity condition,

$$C_i B > 2R_i\left(\Phi\left(\frac{-a_i\psi_i\sqrt{m_{ij}+1}}{\sigma_i}\right) - \Phi\left(\frac{-a_i\psi_i\sqrt{m_{ij}}}{\sigma_i}\right)\right), \quad \forall m_{ij} \in N^+. \tag{15}$$

Let $m_i = \sum_{j=1}^{n_i} m_{ij}$ and $m_i^* = \sum_{j=1}^{n_i} m_{ij}^*$. When (15) holds, note that $m_i^* = \lceil T_i D_i / B \rceil$. This is because if $m_i > m_i^*$, by (15), total costs can be reduced. Conversely, $m_i < m_i^*$ is not feasible because it would violate (10). We use this result along with Proposition 6 to calculate $\widetilde{Z}^{\text{LB}}$ by the following procedure. For a given $n_i = 1, 2, \ldots, m_i^*$, compute $b_i = \lfloor m_i^*/n_i \rfloor$ and $r_i = m_i^* - b_i n_i$. Set $m_{ij} = b_i$ for $j = 1$ to $(n_i - r_i)$, $m_{ij} = b_i + 1$ for $j = (n_i - r_i + 1)$ to n_i, $d_i = m_i^* B - T_i D_i$ and compute $\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij}, d_i)$. For $n_i = 1, 2, \ldots, m_i^*$, repeat this procedure and calculate $\widetilde{Z}_i^{\text{LB}}(\mu) = \min_{n_i}\{\widetilde{Z}_i^{\text{LB}}(n_i, m_{ij}, d_i)\}$. Finally, we solve for $\widetilde{Z}^{\text{LB}} = \text{Max}_{\mu \geq 0}\{-\mu + \sum_{i=1}^{p} \widetilde{Z}_i^{\text{LB}}(\mu)\}$ by conducting a line search on μ. We use lower bound $\widetilde{Z}^{\text{LB}}$ to evaluate the quality of the heuristics used to solve the FCCPSP.

3.2. Upper Bounds and Heuristic Solutions

Because the solution provided by the lower bound in §3.1 might not be feasible to the FCCPSP, we developed heuristics to generate feasible solutions to this problem. We first present an intuitive operational heuristic, which is followed by a more detailed backward assignment heuristic.

Operational Heuristic. In the operational heuristic, we use the process inventory constraints and the actual inventory positions of the products after producing a product batch to decide whether to continue with the campaign for the current product or switch over to a different product. To describe this procedure, we introduce the following notation.

$V_i^{(\max)}$ = maximum permissible inventory for the ith product (units stored).

$V_i^{(\min)}$ = minimum permissible inventory for the ith product (units stored).

$V_{it}^{(a)}$ = actual total inventory of the ith product at the beginning of time t (units stored).

Assume that a batch of the pth product has just been produced at time t. We continue with the campaign for product p if the following constraints are not violated

$$V_{pt}^{(a)} + B - t_p D_p \leq V_p^{(\max)} \tag{16}$$

$$V_{it}^{(a)} - t_p D_i - (\tau_i + t_i) D_i \geq V_i^{(\min)}, \quad \forall i \neq p. \tag{17}$$

Constraint (16) imposes the condition that if we produce another batch of product p, at the end of production, the actual inventory at the tank and the net inventory buildup due to this batch should always be lower than its maximum permissible inventory. Constraint (17) enforces the condition that the actual inventory and the inventory depletion for each product not in production should always be greater than its minimum permissible inventory. Here, inventory depletion for a given product is calculated during production of a batch of product p and during production of a batch of this product assuming that it is immediately produced after product p. In practice, the maximum permissible inventory level is set to the maximum storage allocated for a given product, while the minimum permissible inventory level is set to zero. Rajaram et al. (1999) develop a method to determine the optimum allocation of storage to products in a large-scale industrial process.

To implement this heuristic using real time data on actual inventory positions at time t, we check the feasibility of these m constraints. If none are violated, we continue to produce another batch of product p. When this batch is produced, we restart this procedure at time $t+t_p$. If constraint (16) is violated, we switch to the product for which constraint (17) is tightest. Otherwise, if constraint (17) is violated for one or more products, we switch to the product with the greatest associated violation. To initialize this heuristic, we start with a campaign of the product with the lowest initial inventory or highest holding cost.

It is important to recognize that in this heuristic we minimize the number of setups between campaigns and maximize the duration of a production campaign by initiating a switch only when the demand-dependant boundary conditions represented by these constraints are violated. A similar heuristic is developed and implemented for continuous flow refining processes in the food-processing industry (Rajaram and Karmarkar 2002).

Backward Assignment Heuristic. Assume that at time T demand for all products has been satisfied so that there are $\sum_{j=1}^{n_i} m_{ij} B$ units of finished goods inventory for product i. In the backward assignment heuristic, we assign this inventory back to the raw material stage of the process through the reactor so that we have zero batches in finished goods inventory for each product. This backward assignment is necessary to maintain demand feasibility. In this heuristic, we first assign the product that would lead to the lowest campaign setup and total inventory costs. When a batch of this product has been assigned, we continue with the assignment of this product if it still continues to result in the lowest costs; otherwise, we start the campaign of the product that leads to the lowest costs. This procedure is repeated until the finished goods inventory for all products is zero. We formalize this heuristic using the following steps.

Step 0. *Initialization.* At time T, one batch of product p is assigned back to the raw material stage if $p = \arg\min_{i=1 \text{to} m}\{\partial_i\}$, where: $\partial_i = (\sum_{r\neq i} h_r V_{rT}^{(a)} + K_i + h_i(V_{iT}^{(a)} - B))/(\tau_i + t_i)$. Set $T \to T - t_p$ and go to Step 1.

Step 1. We assign another batch of product p if $\{\sum_{r\neq i} h_r V_{rT}^{(a)} + h_p(V_{pT}^{(a)} - B)\}/t_p < \text{Min}_{i\neq p} \partial_i = (K_i + \sum_{r\neq i} h_r V_{rT}^{(a)} + h_i(V_{iT}^{(a)} - B))/(\tau_i + t_i)$. Set $T \to T - t_p$ and go to Step 2. Otherwise, we assign a batch of product q, where $q = \arg\min_{i\neq p}\{\partial_i\}$, set $T \to T - t_q$, and go to Step 2.

Step 2. If $V_{iT}^{(a)} = 0$, we remove product i from the set of products considered for assignment and go to Step 1. If $V_{iT}^{(a)} = 0$, $\forall i$, stop. Products are then produced in exactly the reverse order of the assignment, so that we start with the last batch of the assignment first and produce the first batch of the assignment last.

4. Computational Study

To test the performance of the heuristics and the lower bounds developed in §§2 and 3 of this paper, we used data from a large food-processing company. Here, data from the sorbitol-production process was used to test the CPSP, while data from the modified starch process was used for the FCCPSP. We first describe the results for the CPSP and then present the corresponding results for the FCCPSP.

The data set used to test the CPSP consisted of all input parameters required for this problem for eight products at a sorbitol-production process at this company. The data set included cost parameters such as the fixed cost for setting up a campaign, rework cost per campaign, and holding costs, and process parameters such as batch size, processing times per batch of product, and setup time per campaign of the product, and, finally, demand rate for each product. Data were provided over a three-year period. We were also provided with data on the attribute errors of each product during this period. We used this data to conclude that the distribution of attribute errors for each product was normally distributed with a large degree of confidence (χ^2 test holds at an $\alpha \leq 0.05$ level across all 8 products). We then computed σ_i, the standard deviation of the distribution of attribute error for product i, which was required in the calculation of the expected cost of rework associated with each product.

Recall that the objective function of the CPSP consists of the switching, holding, and rework costs across all campaigns. To analyze the relative proportions of these costs, we first calculate the ratios of the rework cost to the setup cost per campaign for each product. We then define $K_1 = (\sum_{i=1}^{8} R_i/S_i)/8$ as the average of the ratios across the 8 products. Next, we calculate the ratio of the holding costs per batch to the setup costs per campaign for each product and similarly define $K_2 = (\sum_{i=1}^{8} Bh_i/S_i)/8$ as the average of these ratios across these products. For the reference data set, $K_1 = 0.3$ and $K_2 = 0.1$.

We tested how sensitive our heuristics and the lower bound were to the scale of these cost parameters. To perform this analysis across all products, we scaled the rework costs by 1/3, 1/2, 2, and 3 (i.e., changing K_1 by these factors), scaled the holding cost by the same factors (i.e., changing K_2 by the same factors), and finally scaled the setup costs by these factors. Our scaling factors were chosen by roughly estimating such costs across a variety of food-processing industries that face the CPSP based on informal discussions with managers in these industries. This scaling, in effect, ensured that the cost proportions of these data were representative across the spectrum of food-processing industries. Note that our scaling procedure results in 64 (i.e., $4 \times 4 \times 4$) data sets generated from the reference problem.

We tried to solve the CPSP corresponding to these data sets using leading commercial software programs such as GAMS (Brooke et al. 1992) and CPLEX (1995) loaded on a Dell Optiplex PC. However, we aborted our runs after these programs ran for over four weeks on these subproblems and were not able to generate feasible solutions. This provides validation for developing and using the heuristics and bounds developed to address this problem. We then used the model-based and feasibility heuristic (described in §2.2) to solve the CPSP and to develop the production plan and its associated costs for these 64 data sets. We also computed a lower bound on the costs for each data set using the scheme developed in §2.1. All of these analyses were done using Matlab (MathWorks Inc. 1998) and a specialized C program. Each run was solved within a few minutes on a Dell desktop PC. We define the performance gap of the heuristic as the increase in the cost of a heuristic solution from the lower bound solution expressed as a percentage of the lower bound solution. Across the 64 data sets, the gap of the model-based heuristic ranged from 1% to 4%, averaging around 3%. The gap of the feasibility heuristic ranged from 2% to 5%, averaging around 3.8%.

We wanted to better understand the circumstances under which percentage gaps increase. This could provide us with insights to improve the heuristics and the lower bound. We observed from our analysis that these gaps are uniformly higher when setup costs are higher or when inventory holding and rework costs are lower than the reference case. Conversely, the gaps are significantly lower when setup costs are lower or when inventory holding and rework costs are higher than the reference case. It is important to note that these gaps were reduced in these heuristics largely because the lower bounds became tighter. For instance, the solution provided by the lower bounds increased by an average of 1.5% from the reference case, while the solution provided by the model-based heuristic decreased by an average of 0.7% and the feasibility heuristic decreased by an average of 0.5% from the reference case. This indicates that both of these heuristics are fairly stable for a wide variation in cost parameters, while there could still be potential to improve the lower bounds.

We repeated this analysis for the FCCPSP by using data from a modified starch-production process at this company, which produces over 300 products. Here, the corresponding values of the ratios K_1 and K_2 were 0.7 and 0.8, respectively. Using the same procedure described above, we scaled the holding, rework, and setup costs and formed 64 data sets. Here again, GAMS or CPLEX were not able to generate feasible solutions across any of the data sets. We then tested the operational and backward assignment heuristic (described in §3.2) to develop the production plan and to calculate the corresponding costs across the 64 data sets. We used the procedure developed in §3.1 to develop lower bounds on the solution to this problem and to evaluate the performance of these heuristics. The gap for the operational heuristic ranged from 2% to 5%, averaging around 3.3%. The gap for the backward assignment heuristic ranged from 3% to 6%, averaging around 4.5%. The changes in the gaps of these heuristics in response to changes in cost parameters were similar to those described for the heuristics for the CPSP. Again, we found that these heuristics were stable across a wide range of data, while there was potential to improve the lower bound.

5. Application

We have applied the methods in this paper to production data from the food-processing company discussed in the introduction. The assumptions of the CPSP and FCCPSP are consistent with the operating environment at the sorbitol- and modified starch-production processes, respectively. We tested the CPSP on data provided to us from the sorbitol-production process. The data available to us included three-year data on all parameters used in the CPSP from five processes located in five countries. Note that the data set from the process with the largest daily output was used as the reference in the computational study for the CPSP described in the previous section. We were also provided with data on the attribute errors of each product at each site during this period. We used this data to verify that the distribution of attribute errors for each product was normally distributed with a large degree of confidence (χ^2 test holds at $\alpha \leq 0.10$ level across all products). This data was then used to calculate the standard deviation of the distribution of attribute error for a given product, required in the calculation of the expected cost of rework associated with each product.

We solved the CPSP for the data from these five processes using the model-based and the feasibility heuristics. We used the procedure described in §2.1 to compute lower bounds. The lowest gaps were provided when we used the model-based heuristic. The results summarized in Table 1 indicate that solutions provided by this heuristic were within 4% of the lower bound.

We used the solution provided by the model-based heuristic at a given process and calculated the total campaign setup, holding, and rework costs that would have resulted had the production plan suggested by our method been implemented. We compared our costs with actual annual costs at these processes and found that our method would have reduced total costs by at least 7% at all of these sites. Had this approach been implemented, the total annual cost savings at all these processes would have been around \$3 million. Individual percentage and absolute cost savings for these processes are also summarized in Table 1.

To better understand the underlying reasons for the lowered costs resulting from the model-based heuristic, we compared the solution of the heuristic to the solution based on the existing operating procedure. This comparison found that our solution suggested fewer, but longer, campaigns for each product. While this increased the holding costs for each product, this increase was offset by a significant reduction in campaign setup costs. In addition, there was a notable decrease in expected quality rework costs. This is because, with increased campaign length, more batches could be mixed together, which, in turn, reduced the probability of nonconformance of a product produced in a given campaign.

We also tested the FCCPSP on data over three years from seven modified starch processes from seven different countries, where the data from the process with the largest output was used in the computational study for this problem. We used the operational

Table 1 Percentage Gaps of Heuristics from Tightest Lower Bound and Cost Savings when Applied to Sorbitol Data

Site	Percentage gap using model-based heuristic	Cost reduction (%)	Cost reduction (million \$)
1	3.5	8.0	0.5
2	4.0	7.0	0.4
3	2.5	9.0	0.7
4	3.3	8.0	0.6
5	2.0	9.0	0.8

Table 2 Percentage Gaps of Heuristics from Tightest Lower Bound and Cost Savings when Applied to Modified Starch Data

Site	Percentage gap using operational heuristic	Cost reduction (%)	Cost reduction (million $)
1	3.0	8.0	0.7
2	4.7	7.0	0.4
3	2.5	9.0	0.8
4	5.0	6.0	0.3
5	4.5	7.0	0.6
6	4.8	7.0	0.5
7	4.0	8.0	0.7

and backward assignment heuristics (§3.2) to develop the production plan at each of these processes and used the lower bound (§3.1) to calculate the gaps associated with the heuristic. The lowest gaps were provided by the operational heuristic. The results, summarized in Table 2, indicate that solutions provided by this heuristic were within 5% of the lower bound. We used the solution provided by the operational heuristic at a given process and calculated the total campaign setup, holding rework, and disposal costs that would have resulted had the production plan suggested by our method been implemented. We compared our costs with actual annual costs at these processes and found that our method would have reduced total costs by at least 6% at each site. Had this approach been implemented, the total annual cost savings at all these processes would have been around $4 million. Individual percentage and absolute cost savings for these processes are also summarized in Table 2. Here again, we observed that the solution provided by the operational heuristic had fewer, but longer, campaigns across all products than the solution based on the current operating procedure. Furthermore, the increase in holding costs due to the longer campaigns in the operational heuristic was offset by the reduction in campaign setup and expected quality rework costs to such an extent that it lowered total costs when compared to the current procedure.

The main insight that can be drawn from our analysis of the CPSP and the FCCPSP at this company is that it is more profitable to run fewer and longer product campaigns than current practice. This insight was important for this company because typically products were produced more frequently and in shorter campaigns than necessary to minimize finished goods inventories. The focus was on minimizing finished goods inventory because such inventory was more visible than aspects such as setups and rework to the management, who consequently felt that reducing this inventory could lead to the greatest cost reduction. The results of our analysis show that this may not always be true. In particular, it may be profitable to carry more inventory at higher holding costs, especially if these cost increases are much smaller than the savings from the resulting reduction of campaign setup and expected quality rework costs.

6. Summary and Future Research

In this paper, we consider the problem of planning scheduling multiproduct, single-stage, single-equipment batch operations in the food-processing industry. Such operations are encountered in many applications including manufacturing of sorbitol, modified starches, and specialty sugars. Unlike discrete manufacturing, batch sizes in these operations cannot be set arbitrarily, but are often determined by equipment size. Multiple batches of the same product are often run sequentially in campaigns to minimize setup and quality costs.

We considered the problem of determining the timing and duration of product campaigns to minimize average setup, quality, and inventory holding costs over a horizon. First, we formulated the deterministic static version of this problem over an infinite horizon and showed that, in general, a feasible finite cyclic solution might not exist. We then provided sufficient conditions for the existence of a finite cycle, used single-product problems to provide lower bounds on the costs for the multiproduct problem, and employed them to test heuristics developed for this problem. Next, we modified this formulation to incorporate fixed cycles, which might be necessary due to factors such as product obsolescence, perishability, or contracts with customers, by allowing for disposal of excess stock so that finite cycles are always feasible, though they might not be optimal. We also developed bounds and heuristic solution procedures for this case. These methods were applied to data from five sorbitol-production processes and seven modified starch-production processes at a leading food-processing company. Our results suggest that our methods could potentially reduce total annual costs by about 7.7%, translating to an annual savings of around \$7 million across all these processes.

Our future research will address some important extensions to this problem. One natural extension of this problem is to the multiproduct, multistage, multi-reactor case found in applications such as bakeries, confectioneries, and breweries. The increased complexity of this problem would undoubtedly require different heuristics and different procedures for determining lower bounds. Another extension could be to the problem of scheduling a multiproduct batch reactor with yield uncertainty and operator learning. This problem typically occurs across a variety of batch processes in the food-processing industry that use biochemical reagents. Determining the optimal campaign in this case is more complex, as this decision would depend on the realized yield from each batch. In addition, yields typically improve with increased duration of the campaign, due both to learning and to the nature of the processes. Finally, extensions to the batch reactor planning and scheduling problem arise when product demand is variable due to seasonal or broader economic factors. Depending on whether such demand variability is a known or unknown, the models and solution procedure that are required to address this case could be significantly different.

Acknowledgments

The authors would like to thank Professor Leroy B. Schwarz, a senior editor, and two referees for their constructive comments.

Appendix

Example

In this example, we illustrate how rounding the continuous solution to the ELSP might lead to an infeasible solution to the CPSP. We consider a three-product problem with the following parameters.

Product i	Setup cost (\$) S_i	Holding cost (\$/unit) h_i	Demand rate (units/month) D_i	Setup time per campaign (months) τ_i	Processing time per batch (months) t_i
1	100	1	2,000	0.005	0.0875
2	200	1.5	3,000	0.01	0.0583
3	300	2	4,000	0.015	0.0438

Furthermore let $B=700$ and for simplicity let $R_i=0$ $\forall i$. To determine a finite, feasible cycle of length T, we first determine the continuous solution to the ELSP using the approach of Hanssman (1962). In this approach, we impose the rule that each product is produced once in each cycle. This is equivalent to requiring that the number of runs per month be the same for all products, say N. Then, the total costs are $N\sum_{i=1}^{p} S_i + (1/2N)\sum_{i=1}^{p} h_i D_i(1-D_i/r_i)$, where $r_i = B/t_i$ is the production rate for the ith product. It is easy to show that the optimal number of runs minimizing the total costs is $\widetilde{N}^* = \sum_{i=1} h_i D_i(1-D_i/r_i)/2\sum_{i=1}^{p} S_i$. Because we require $\widetilde{N}^*$ to be an integer, we check the two bracketing integers and pick N^*, the one resulting in the lowest total cost. Then $T=1/N^*$ and the optimal lot size for item i is $Q_i^* = D_i/N^*$. Using these formulas in this example, we calculate that $N^*=3$, $T=1/3$, $Q_1^*=666.67$, $Q_2^*=1{,}000$, and $Q_3^*=1{,}333.33$. Note that $\sum_{i=1}^{3}(\tau_i + Q_i^*/r_i) = 0.28 < T$, $Q_1^* = TD_1$, $Q_2^* = TD_2$, and $Q_3^* = TD_3$, so that this represents a feasible solution to the ELSP.

To translate this solution to a feasible solution to the CPSP, we need to consider batch size $B=700$, set $n_i = 1$ $\forall i$, and round up Q_1^* to 700 so that $m_{11}=1$, round down Q_2^* to 700 so that $m_{21}=1$, and round up Q_3^* to 1,400 so that $m_{31}=2$. However, this would not represent a feasible solution to the CPSP because $m_{11}B \neq TD_1$, $m_{21}B \neq TD_2$, and $m_{31}B \neq TD_3$.

References

Brooke, A., D. Kendrick, A. Meeraus. 1992. *GAMS: A User's Guide*. The Scientific Press, San Francisco, CA.

Castro, P., A. Barbosa-Pova, H. Matos. 2001. An improved RTN continuous-time formulation for the short-term scheduling of multipurpose batch plants. *Indust. Engrg. Chemistry Res.* **40**(9) 2059–2068.

CPLEX Optimization Inc. 1995. *CPLEX Version 4.0 Manual*. Incline Village, NV.

Dessouky, M. M., B. A. Kijowski. 1997. Production scheduling of single-stage multi-product batch chemical processes with fixed batch sizes. *IIE Trans.* **29** 399–408.

Dobson, G. 1987. The economic lot-scheduling problem: Achieving feasibility using time-varying lot sizes. *Oper. Res.* **35**(5) 764–771.

Elmaghraby, S. 1978. The economic lot scheduling problem (ELSP): Review and extensions. *Management Sci.* **24**(6) 587–598.

Hanssmann, F. 1962. *Operations Research in Production and Inventory Control.* John Wiley and Sons, New York.

Hsu, W. 1983. On the general feasibility test of scheduling lot sizes for several products on one machine. *Management Sci.* **29**(1) 93–105.

Johnson, L. A., D. C. Montgomery. 1974. *Operations Research in Production Planning, Scheduling and Inventory Control*. John Wiley and Sons, New York.

Karmarkar, U. S., L. Schrage. 1985. The deterministic dynamic product cycling problem. *Oper. Res.* **33**(2) 326–345.

Lin, X., C. Floudas, S. Modi, N. Juhasz. 2002. Continuous-time optimization approach for medium-range production scheduling of a multiproduct batch plant. *Indust. Engrg. Chemistry Res.* **41**(16) 3884–3906.

Luenberger, D. G. 1984. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA.

MathWorks Inc. 1998. *The Mathlab Curriculum Series*. *Users Guide*. Prentice Hall, Englewood Cliffs, NJ.

Maxwell, W. L. 1964. The scheduling of economic lot sizes. *Naval Res. Logist. Quart.* **11** 89–124.

Orcun, S., A. Discioglu, O. Hortacsu. 1999. Optimal planning and scheduling of batch plants with operational uncertainties: An industrial application to baker's yeast production. *Comput. Chemical Engrg.* **23** S183–S186.

Orcun, S., A. Discioglu, O. Hortacsu. 2001. General continuous time models for production planning and scheduling of batch processing plants, mixed integer linear program formulations and computational issues. *Comput. Chemical Engrg.* **25** S371–S389.

Rajaram, K., U. S. Karmarkar. 2002. Product cycling with uncertain yields: Analysis and application to the process industry. *Oper. Res.* **50**(4) 680–691.

Rajaram, K., R. Jaikumar, F. Behlau, F. Van Esch, C. Heynen, R. Kaiser, A. Kuttner, I. Van De Wege. 1999. Robust process control at Cerestar's refineries. *Interfaces* **29**(1) 30–48.

Reklaitis, G. V. 1992. Overview of scheduling and planning of batch process operations. *Proc. NATO ASI Batch Process System Engrg.* Antalya, Turkey.