PRODUCT CYCLING WITH UNCERTAIN YIELDS: ANALYSIS AND APPLICATION TO THE PROCESS INDUSTRY

KUMAR RAJARAM and UDAY S. KARMARKAR

Decision, Operations, and Technology Management,
The Anderson School, University of California at Los Angeles,
110 Westwood Plaza, Los Angeles, California 90095-1481,
kumar.rajaram@anderson.ucla.edu

(Received February 2000; revision received November 2000; accepted May 2001)

We formulate the dynamic product-cycling problem with yield uncertainty and buffer limits to determine how much product to produce at what time to minimize total expected switching, production, inventory storage, and backorder costs. A "restricted" Lagrangian technique is used to develop a lower bound and a model-based Lagrangian heuristic. We also develop an operational heuristic and a greedy heuristic. The operational heuristic has been implemented at seven refineries at Cerestar, Europe's leading manufacturer of wheat-and corn-based starch products in the food-processing industry. This has already reduced total costs by around 5 percent or $3 million annually at these sites. Tests of the Lagrangian heuristic on data from these refineries during this period have shown the potential to further reduce total costs by at least 2 percent or about $1 million. In addition, the Lagrangian heuristic has provided an objective basis to evaluate the economic impact of several strategic decisions involving issues such as buffer expansion, variability reduction, and product selection.

1. INTRODUCTION

In many process industries, a single, capital intensive, highly specialized process such as a refinery manufactures only a few products. Typically, such processes are dedicated to a single product at a time, and incur significant downtimes and costs when products are switched and production yields are uncertain. Examples of such manufacturing processes can be found across many process industry sectors including food processing, glass and paper manufacturing, chip fabrication, metal fabrication, plastic extrusion and moulding, and semicontinuous chemical processes found in the petrochemical and pharmaceutical industry. The key production-planning problem in these processes is to determine which products to produce at which time to meet demand in the most cost-effective manner. This type of problem is commonly called a "product-cycling" problem because the process switches between products, although products are not usually produced in the same repetitive cycle.

In these processes, while one product is being produced, inventory for all other products is being decreased by demand. Since these products are usually commodities and are intermediary in nature, short-term demand is known, but could be variable during the planning horizon. On the other hand, production yield is uncertain due to the characteristics of the chemical processes in many of these operations. However, using historical production data, we can derive the distribution of yield realizations for each product during a production period. There are variable costs incurred for holding inventory between periods, for backorders and for production. There are limits on the total amount of inventory that can be accumulated; beyond these limits, significant handling costs are incurred. At each product changeover, fixed-switching costs are incurred and production time is lost. In addition, the switching costs and times may be sequence dependent. In this paper, we formulate and solve the resulting dynamic product-cycling problem with yield uncertainty and buffer limits to determine which product to produce at which times to minimize total expected switching, production, inventory storage (including holding and handling), and backorder costs.

The static, constant demand, deterministic yield, continuous time version of the product-cycling problem is well known (Elmaghraby 1978, Dobson 1987). Bourland et al. (1997) consider this problem when the yield rate is deterministic but increasing with the duration of the production run. The discrete period, dynamic version has received less attention. This problem with deterministic yield and demand has been considered by Karmarkar and Scharge (1985), Salomon et al. (1989), Fleischmann (1990), Salomon (1990), and Pochet and Wolsey (1991). The static discrete-period problem with deterministic yield and stochastic demand has been studied by Graves (1980), who develops heuristic solution techniques. The dynamic version of this problem has been studied by Karmarkar and Yoo (1994), who present decomposition techniques and report computational studies on small-scale problems.

The problem considered here differs from the problems addressed by these papers in several ways. First, while there has been a considerable body of work dealing with lot sizing under yield uncertainty (Yano and Lee 1995), the authors are not aware of any work that formulates and solves the dynamic product-cycling problem with uncertain yield, encountered in many process industry settings.

Subject classifications: Production planning: product cycling with yield uncertainty. Production application: process industry, food-processing sector. Production heuristics: Lagrangian heuristic, operational heuristic.
Area of review: Manufacturing, Service, and Supply-Chain Operations.

Second, much of the literature to date addresses problems that are both small in size and 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 observe that powerful commercial software tools cannot generate feasible solutions to even small problem instances. 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 application of these methods to the food-processing industry. We consider the dynamic product-cycling problem with uncertain yields and buffer limits faced at seven refineries of Cerestar, Europe's leading manufacturer of starch and wheat-based products in the food-processing industry. The methodologies developed to address this problem are validated and implemented at these seven sites as part of a broader multiyear research project described in Rajaram et al. (1999). This work has had a major economic, strategic, and organizational impact at this company.

This paper is organized as follows. In the next section, we formulate the dynamic product-cycling problem with yield uncertainty and buffer limits as a stochastic dynamic program with mixed continuous and (0-1) integer variables. Except in special cases, the state space of this problem is too large to allow the use of conventional dynamic programming techniques. Consequently in §3, we employ a "restricted" Lagrangian decomposition approach to decouple the problem into more manageable pieces and provide a lower bound on expected costs. In §4, we use the solution of the decomposition to develop a model-based Lagrangian heuristic and also describe an operational heuristic and a greedy heuristic. The solutions provided by these heuristics provide an upper bound on optimal expected costs. In §5, we perform a computational study to test the performance of these heuristics. In §6, we describe the application of our method at seven of Cerestar's refineries. In the concluding section, we summarize our work and discuss its economic, strategic, and organizational impact at Cerestar.

2. MODEL FORMULATION

Consider a manufacturing process producing m products and let $i \in I = (1, \dots m)$ index the set of products. These products are produced during n time periods indexed by $t \in T = (1, \dots n)$ to meet a prespecified level of demand during each period. However, output during each period varies due to yield uncertainty. Product changeovers take place during the beginning of the period and require a fixed downtime. Only one type of product can be produced during each period. To decide what product, when, and how much to produce, define the variables:

$$Y_{it} = \begin{cases} 1 & \text{if product } i \text{ is produced in period } t, \\ 0 & \text{otherwise,} \end{cases}$$

$$Z_{jit} = \begin{cases} 1 & \text{if we switch from product } j \text{ to product } i \\ & \quad \text{at the begining of period } t, \\ 0 & \text{otherwise.} \end{cases}$$

We are given:

D_{it} = demand for product i during period t (units demanded/unit-time).
B_i = maximum storage level for product i (units stored).
S_{jit} = downtime incurred when we switch from product j to product i at the beginning of period t (fraction of length of period t). We assume that $S_{iit} = 0$.
K_{jit} = cost for switching over from product j to product i at beginning of period t. We assume that $K_{iit} = 0$.
C_{it} = cost of producing product i in period t (\$).
h_i = holding cost per unit of product i (\$/unit-time).
e_i = handling cost per unit of product i in excess of available buffer capacity (\$/unit-time).
p_i = backorder cost per unit of product i (\$/unit-time).
$\tilde{a}_i$ = random variable representing the yield distribution of product i during a production period (units-produced /unit-time), with $\bar{a}_i$ representing the expected value of this distribution.

We define $X_{it} = Y_{it} - \sum_j S_{jit} Z_{jit}$ as the fraction of time period t available for production if we switch over to product i at the beginning of period t. Let $I_{i(t-1)}$ units of inventory be available at the end of period $t-1$. During period t, $\tilde{a}_i X_{it}$ units of the ith product are produced. Thus, $\widetilde{I}_{it} = \tilde{a}_i X_{it} + I_{i(t-1)} - D_{it}$ units of inventory are available at the end of period t. For product i in period t, we represent the expected holding costs of producing more than demanded, handling costs when we exceed the allocated buffer size and backorder cost for producing less than demanded by $L_{it}(X_{it},\ I_{i(t-1)}) = E_{\tilde{a}_i}\{h_i(\widetilde{I}_{it})^+ + e_i(\widetilde{I}_{it} - B_i)^+ + p_i(-\widetilde{I}_{it})^+\}$. The dynamic product-switching problem with yield uncertainty is formulated as a stochastic dynamic program with mixed continuous and (0-1) integer variables, in which during the last period of the n period problem, we solve problem (P^n), where:

$$\xi_n(Y_{n-1}, I_{n-1}) = \operatorname*{Min}_{Y_n, Z_n} \sum_i \Big\{ C_{in} Y_{in} + \sum_j K_{jin} Z_{jin} + L_{in}(X_{in}, I_{i(n-1)}) \Big\} \qquad (P^n)$$

subject to

$$(1n) \quad X_{in} = Y_{in} - \sum_j S_{jin} Z_{jin} \quad \forall\, i,$$

$$(2n) \quad Z_{jin} \geqslant Y_{in} + Y_{j(n-1)} - 1 \quad \forall\, i,$$

$$(3n) \quad \sum_i Y_{in} = 1,$$

$$(4n) \quad Y_{in}, Z_{ijn} \in \{0, 1\} \quad \forall\, i,$$

$$(5n) \quad X_{in} \geqslant \quad \forall\, i.$$

In this problem, we minimize the total expected production, switching, inventory storage (including holding and handling) and backorder costs during the nth period. Constraint $(1n)$ calculates the time per period available for production. Constraint $(2n)$ ensures that $Z_{jin} = 1$ if we switch from product j to product i at the beginning of period n. Constraint $(3n)$ enforces that only one product is produced per period, while integrality and nonzero conditions are enforced by Constraints $(4n)$ and $(5n)$. For $t = 1, 2, \ldots n-1$, we solve problem P^t where:

$$\xi_t(Y_{t-1}, \widetilde{I}_{t-1}) = \min_{Y_t, Z_t} \sum_i \Big\{ C_{it} Y_{it} + \sum_j K_{jit} Z_{ijt} + L_{it}(X_{it}) + E_{\tilde{a}_i}[\xi_{t+1}(Y_t, \tilde{I}_t)] \Big\} \qquad (P^t)$$

subject to

$(1t)$ $\quad X_{it} = Y_{it} - \sum_j S_{jit} Z_{jit} \quad \forall i,$

$(2t)$ $\quad Z_{jit} \geqslant Y_{it} + Y_{j(t-1)} - 1 \quad \forall i,$

$(3t)$ $\quad \sum_i Y_{it} = 1,$

$(4t)$ $\quad Y_{it}, Z_{ijt} \in \{0, 1\} \quad \forall i,$

$(5t)$ $\quad X_{it} \geqslant 0 \quad \forall i.$

The optimal solution for the entire time horizon is obtained by solving problem (P^1) and has the value $\xi_1(Y_0, I_0)$. This stochastic-dynamic problem is difficult to solve due to the large state space, as the inventory and product state variables can take on a large set of values. It can be shown that even if we restrict I_n to a finite set of values by assuming a discrete distribution of yield, this set can take $(I_n)^n m^3 n$ values. Consequently, we elected to develop a Lagrangian relaxation of the problem to provide lower bounds and use this solution to construct and evaluate heuristics to address this problem.

In general, the choice of the specific multiplier values used in the relaxation could be deferred until the starting state in the given period is known. However, the multipliers would then effectively be functions of random variables representing the yield in the previous period. To make the computations tractable, we treat the multipliers as though they are fixed at the start of the problem. In effect, the choice of the multipliers is being restricted to the class of constant functions. This "restricted Lagrangian method" led to promising results when applied to stochastic multiperiod, multilocation inventory-allocation problems by Karmarkar (1987). However, the same technique applied to a combinatoric-stochastic problem by Karmarkar and Yoo (1994) was much less successful. An essentially similar decomposition technique in a continuous time, convex control setting was proposed by Kleindorfer (1973) and Kleindorfer and Glover (1973), who did not examine the computational viability of their method. In the next section, we apply this technique to decompose the problem and generate lower bounds.

3. PROBLEM DECOMPOSITION AND LOWER BOUNDS

To develop an appropriate decomposition, we first consider problem (P^n). This problem consists of a *process-switching problem* in which we decide what product to produce in each period based on minimizing switchover and variable production costs per period and a *production-inventory problem* in which we make a production decision based on minimizing expected inventory storage and backorder costs. These problems are linked together by Constraint $(1n)$. We relax this constraint by introducing multipliers μ_{in}, separating the process-switching problem from the production-inventory problem, which results in the following subproblems in period n:

$$W_n(Y_{n-1}, \mu_{in}) = \min_{Y_n, Z_n} \sum_i \Big\{ (C_{in} - \mu_{in}) Y_{in} + \sum_i \sum_j (K_{jin} + \mu_{in} S_{jin}) Z_{jin} \Big\} \qquad (PS^n)$$

subject to $(2n), (3n), (4n)$,

$$U_{in}(I_{i(n-1)}, \mu_{in}) = \min_{X_{in}} \mu_{in} X_{in} + L_{in}(X_{in}, I_{i(n-1)}) \qquad (PI^{in})$$

subject to $(5n)$,

$$\Pi_n(Y_{n-1}, I_{i(n-1)}, \mu_{in}) = W_n(Y_{n-1}, \mu_{in}) + \sum_i U_{in}(I_{i(n-1)}, \mu_{in}). \qquad (PR^t)$$

The relaxed problem (PR^n) consists of a process-switching problem (PS^n) and m production inventory problems (PI^{in}), $i = 1$ to m, each corresponding to an individual product. In period t, we solve:

$$W_t(Y_{t-1}, \mu_t) = \min_{Y_n, Z_n} \sum_i \Big\{ (C_{it} - \mu_{it}) Y_{it} + \sum_i \sum_j (K_{jit} + \mu_{it} S_{jit}) Z_{jit} + W_{t+1}(Y_t, \mu_{t+1}) \Big\} \qquad (PS^t)$$

subject to $(2t), (3t), (4t)$,

$$U_{it}(I_{i(t-1)}, \mu_{it}) = \min_{X_{in}} \mu_{it} X_{it} + L_{it}(X_{it}, I_{i(t-1)}) + E_{\tilde{a}_i}(U_{i(t+1)}(\widetilde{I}_{it}, \mu_{i(t+1)})) \qquad (PI^{it})$$

subject to $(5t)$,

$$\Pi_t(Y_{t-1}, I_{i(t-1)}, \mu_{it}) = W_t(Y_{t-1}, \mu_{it}) + \sum_i U_{it}(I_{i(t-1)}, \mu_{it}). \qquad (PR^t)$$

We solve the process-switching problem (PS^t) at $t = 1$ by a straightforward dynamic programming algorithm of complexity $O(m^2 n)$. This algorithm is described in detail in the Appendix.

The production-inventory problem for the ith product at the nth period is a problem similar to the classic newsboy problem. In this problem, demand for the product is known in the short term, but production quantity is unknown due to yield uncertainty. This type of problem is known to occur in several process industries including petroleum refining, petrochemical, glass, pharmaceutical, and food-processing industries where demand is known in the short term, but production varies due to uncertainty in the yield of the chemical reactions used in these processes. Bollapragada and Morton (1999) study the stationary version of this problem with demand uncertainty across an infinite horizon. We next develop the solution to (PI^{in}).

PROPOSITION 3.1. *Let X^*_{in} be the optimal solution to problem (PI^{in}). Then, if random variable $\tilde{a}_i$ has distribution function $F_i(a_i)$ and $\Lambda_i(k)=\int_0^k \tilde{a}_i \partial F_i(a_i)$ for any real k, X^*_{in} satisfies*:

$$\Lambda_i\left(\frac{D_{in}-I_{i(n-1)}}{X^*_{in}}\right)+\frac{e_i}{h_i+p_i}\Lambda_i\left(\frac{D_{in}+B_i-I_{i(n-1)}}{X^*_{in}}\right) = \frac{\mu_{in}+(h_i+e_i)\bar{a}_i}{h_i+p_i}.$$

PROOF. Let $\Psi_i(X_{in}, I_{i(n-1)}) = \mu_{in}X_{in}+L_{in}(X_{in}, I_{i(n-1)}) = \mu_{in}X_{in}+E_{\tilde{a}_i}\{h_i(\widetilde{I}_{in})^+ + e_i(\widetilde{I}_{in}-B_i)^+ + p_i(-\widetilde{I}_{in})^+\}$, where $\widetilde{I}_{in}=\tilde{a}_iX_{in}+I_{i(n-1)}-D_{in}$. Then,

$$\begin{aligned}\frac{\partial\Psi_i(X_{in}, I_{i(n-1)})}{\partial X_{in}} &= \mu_{in}+h_i\int_{\frac{D_{in}-I_{i(n-1)}}{X_{in}}}^{\infty}\tilde{a}_i\partial F(a_i)+e_i\int_{\frac{D_{in}+B_i-I_{i(n-1)}}{X_{in}}}^{\infty}\tilde{a}_i\partial F(a_i)\\ &\quad -P_i\Lambda_i\left(\frac{D_{in}-I_{i(n-1)}}{X_{in}}\right)\end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2\Psi_i(X_{in}, I_{i(n-1)})}{\partial X_{in}^2} &= (h_i+p_i)\frac{(D_{in}-I_{i(n-1)})^2}{X_{in}^3}\partial F\left(\frac{D_{in}-I_{i(n-1)}}{X_{in}}\right)\\ &\quad +e_i\frac{(D_{in}+B_i-I_{i(n-1)})^2}{X_{in}^3}\partial F\left(\frac{D_{in}+B_i-I_{i(n-1)}}{X_{in}}\right)\geqslant 0.\end{aligned}$$

Thus, the condition $\partial\Psi_i(X^*_{in}, I_{i(n-1)})/\partial X^*_{in}=0$ is both necessary and sufficient to establish the optimality of this problem at X^*_{in}. We use this condition along with the fact that $\bar{a}_i=\int_0^\infty\tilde{a}_i\partial F_i(a_i)=\Lambda_i(k)+\int_k^\infty\tilde{a}_i\partial F_i(a_i)$ and rearrange and group terms to derive the required result.

PROPOSITION 3.2. *$U_{in}(I_{i(n-1)}, \mu_{in})$ is a convex function of $I_{i(n-1)}$.*

PROOF.

$$\begin{aligned}U_{in}(I_{i(n-1)}, \mu_{in}) &= \mu_{in}X^*_{in}+L(X^*_{in}, I_{i(n-1)})\\ &= \mu_{in}X^*_{in}+E_{\tilde{a}_i}\{h_i(\widetilde{I}_{in})^+ + e_i(\widetilde{I}_{in}-B_i)^+\\ &\quad + p_i(-\widetilde{I}_{in})^+\},\end{aligned}$$

where $\widetilde{I}_{in}=\tilde{a}_iX_{in}+I_{i(n-1)}-D_{in}$.

$$\begin{aligned}\frac{\partial U_{in}(I_{i(n-1)}, \mu_{in})}{\partial I_{i(n-1)}} &= h_i\left(1-F\left(\frac{D_{in}-I_{i(n-1)}}{X_{in}}\right)\right)\\ &\quad +e_i\left(1-F\left(\frac{D_{in}+B_i-I_{i(n-1)}}{X_{in}}\right)\right)\\ &\quad -p_iF\left(\frac{D_{in}-I_{i(n-1)}}{X_{in}}\right),\\ \frac{\partial U^2_{in}(I_{i(n-1)}, \mu_{in})}{\partial I^2_{i(n-1)}} &= (h_i+p_i)\partial F\left(\frac{D_{in}-I_{i(n-1)}}{X_{in}}\right)\\ &\quad +e_i\partial F\left(\frac{D_{in}+B_i-I_{i(n-1)}}{X_{in}}\right)\geqslant 0.\end{aligned}$$

Thus, $U_{in}(I_{i(n-1)}, \mu_{in})$ is convex in $I_{i(n-1)}$.

Using these propositions and the fact that the expectation operator preserves convexity, we solve m multiperiod production-inventory problems (PI^{it}) at each period t using backward recursion for a particular choice of multipliers. The computational procedure for solving this problem is included in the Appendix. If switchover times are sequence independent (i.e., $S_{jit}=S_{it}$, for all i, j), this computation is greatly simplified since during any period, $X_{it}\in\{0, 1-S_{it}, 1\}$. To get a tight lower bound, we solve the Lagrangian dual problem (PD^t) defined as:

$$\bar{\xi}_t(Y_{t-1}, I_{i(t-1)}) = \operatorname*{Max}_{\mu_{it}}\Pi_t(Y_{t-1}, I_{i(t-1)}, \mu_{it}). \qquad (PD^t)$$

To solve the Lagrangian-dual problem (PD^t), in principle, we could use a subgradient search method over the multipliers to tighten the bound as much as possible. This procedure (e.g., Bertsekas 1995) has been successfully applied to address a variety of large-scale deterministic-combinatorial problems. In this problem, this subgradient depends on the *expected* production schedule for product i during the remaining $(n-t)$ periods. Thus, this procedure would be computationally prohibitive for a practical-sized problem.

To develop an alternate procedure, we note that multiplier μ_{ij} can be interpreted as a credit for switching to product i in period j. From any given heuristic solution in which we know what product we produce in a given time period, we estimate a trial value for μ_{ij} as:

$$\mu_{ij}=\begin{cases} p_i & \text{if } \sum_{t=1}^{j-1}\bar{a}_iY_{it}<\sum_{t=1}^{n}D_{it},\\ -nh_i & \text{if } \sum_{t=1}^{n}D_{it}\leqslant\sum_{t=1}^{j-1}\bar{a}_iY_{it}\leqslant\sum_{t=1}^{n}D_{it}+B_i,\\ -nh_i-e_i & \text{if } \sum_{t=1}^{j-1}\bar{a}_iY_{it}>\sum_{t=1}^{n}D_{it}+B_i.\end{cases}$$

Here, we set the multiplier (the credit gained from switching to product i) to the backorder cost if the expected quantity produced until period $j-1$ is smaller than the total demand required over the planning horizon. If we switch to this product when the quantity produced until period $j-1$ exceeds the total demand over the planning horizon, but within the buffer limits, this credit becomes a holding penalty for the excess inventory across the entire planning horizon. Finally, if the resulting inventory exceeds buffer limits, we would incur handling costs in addition to holding costs.

We use the lower bound generated by this procedure to evaluate the quality of the heuristics developed to address this problem. If this lower bound is used to generate a model-based Lagrangian heuristic, we perform this procedure in an iterative manner until the gap between the heuristic and lower bound either reaches an acceptable value or does not show any further reduction.

4. HEURISTICS AND UPPER BOUNDS

In general, the lower bound solution to problem (P^t) may not be feasible since Constraint ($1t$) is violated. Consequently, we developed heuristics to solve this problem. We first present an intuitive operational heuristic, which was initially developed for the application to enable planners to understand and develop confidence in our approach, improve upon current practice, and better appreciate the gains of implementing more detailed heuristics. Next, we describe a greedy heuristic that is more involved than the operational heuristic, but does not use the decomposition of problem. Finally, we develop a model-based Lagrangian heuristic that uses the solution of the decomposition of problem (P^t) described in §3 to develop a feasible solution.

Operational Heuristic

In the operational heuristic, we use the process-inventory constraints and the actual inventory positions of the products at the beginning of the period to decide whether to continue with the campaign for the current product or switch over to a different product during that period. 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 period t (units stored).

$\mathrm{T}_{i,sh}$ = time required to shut down the ith product under production. This is independent of the next product to be produced (fraction of length of period t).

Assume that the ith product is being produced in period $(t-1)$. At the beginning of period t, we define the following constraints:

$$V_{it}^a + (\bar{a}_i - D_{it})T_{i,sh} \leqslant V_i^{(\max)}, \tag{1}$$

$$V_{jt}^a - D_{jt}S_{ijt} \geqslant V_j^{(\min)} \quad \forall j \neq i. \tag{2}$$

Constraint (1) imposes the condition that while the ith product is being produced, actual inventory at the tank and the maximum expected inventory build-up during shutdown should always be lower than its maximum permissible inventory. The remaining constraints for each product not under production each enforces the condition that the actual inventory and the maximum expected inventory depletion during its switchover should always be greater than its minimum permissible inventory.

To implement this heuristic, using real time data on actual inventory positions at the beginning of period t, we check the feasibility of these m constraints. If none of them are violated, we continue to produce product i until the beginning of period $t+1$ and restart this procedure. If Constraint (1) is violated, we switch to that product for which Constraint (2) is tightest. Otherwise, if Constraint (2) is violated for one or more products, we switch to the product with the greatest associated violation.

It is important to recognize that in this heuristic we minimize the number of basic grade switches 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.

Greedy Heuristic

In the greedy heuristic, we use the expected cost at a given period represented by the production, switching, and expected inventory holding, backorder, and handling costs to decide which product to produce at that particular time period. To formalize the greedy heuristic, assume without loss of generality that product f was produced in period $(t-1)$ and product i is being considered for production in period t. Then, let $\overline{V}_{it}$ represent the expected inventory level of product i at the end of period t, where $\overline{V}_{it} = \bar{a}_i X_{it} + V_{it}^a - D_{it}$. The expected inventory backorder or holding and handling costs associated with $\overline{V}_{it}$ are represented by $\ddot{e}_{it} = h_i(\overline{V}_{it})^+ + e_i(\overline{V}_{it} - B_i)^+ + p_i(-\overline{V}_{it})^+$. Similarly, we let $\overline{V}_{jt} = V_{jt}^a - D_{jt}$ be the inventory level for product $j \neq i$ and let $\ddot{e}_{jt} = h_j(\overline{V}_{jt})^+ + e_j(\overline{V}_{jt} - B_i)^+ + p_j(-\overline{V}_{jt})^+$ represent the associated expected inventory cost. In the greedy heuristic, we produce product g at the beginning of period t, where

$$g = \underset{i=1\text{tom}}{\arg\min} \left\{ C_{it} + K_{fit} + \ddot{e}_{it} + \sum_{j\neq i} \ddot{e}_{jt} \right\}.$$

It is important to recognize that in this heuristic, we produce the product with results in the smallest expected costs during a given period.

Lagrangian Heuristic

In the process-switching problem, decisions on which product to produce and when to produce are made by minimizing production and switching costs. In the Lagrangian heuristic, we adjust the solution of the process-switching problem to ensure that each product realizes at least R_i runs, where

$$R_i = \frac{\sum_{t=1}^{n} D_{in}}{\bar{a}_i} + \left(n - \sum_{i=1}^{m} \frac{(\sum_{t=1}^{n} D_{in})}{\bar{a}_i}\right)\left(\frac{\overline{S}_i}{\sum_{i=1}^{m} \overline{S}_i}\right).$$

Here, $\bar{a}_i$ is the expected yield per period for product i, while $\overline{S}_i$ represents the average downtime incurred when we switch over to product i. The first term of this equation represents the minimum expected number of runs required to meet total demand during the planning horizon. The second term represents the residual time allocated to this product to compensate for switching downtimes. This allocation is achieved by partitioning the total remaining time across the horizon in proportion to the magnitude of average switching downtimes.

To adjust the solution of the process-switching problem so that each product realizes R_i runs, we define the following:

k = iteration index.
R_i^k = number of runs required at iteration k.
t^k = start time of iteration k.
I^k = index set of products at iteration k.

The planning horizon is adjusted to accommodate R_i runs of the ith product using the following procedure.

Step 0: Initialization: $k = 0$, $t^k = 0$ and $I^k = I$.

Step 1: Elimination:
1.1. Solve the process-switching problem starting at period t^k.
1.2. Let product p be the first product to reach its limit R_p in time period t.
1.3. Update $I^{k+1} = I^k - \{p\}$.
1.4. If $I^{k+1} = \{\ \}$, Stop. Else go to Step 2.

Step 2: Recalculation:
2.1. Update $R_i^{k+1} = R_i - R_i^k \quad \forall i å I^{k+1}$.
2.2. Set $k = k + 1$ and $t^k = t$ and go to Step 1.

Note that for a given product, by aiming to meet total demand during the planning horizon, this heuristic indirectly incorporates inventory storage and backorder costs, and thus embeds the production inventory aspect in the process-switching problem. It can be shown that this heuristic is of complexity $O(m^2n^2)$.

5. COMPUTATIONAL STUDY

To test the performance of these heuristics and of the lower bound we used a data set from Cerestar's largest glucose refinery located at Sas van Gent in the Netherlands. This data set consisted of all input parameters required for this problem for seven products (also known as basic grades) produced in the refining process. These included cost parameters such as production, switching, inventory storage (holding and handling), and backorder costs; process parameters such as yield distributions, switching downtimes, and buffer limits, and finally, demand for each basic grade. Data was provided for a three-month or 92-day duration. This process operates continuously on a three-shift basis and each period was set to one shift of eight hours. Thus, we had data across 276 continuous shifts or periods. The planning horizon was typically one week or 21 periods.

Recall that the objective function of the product-switching problem at any period consisted of production, switching, inventory backorder, and holding and handling costs. To analyze the relative proportions of these costs, define:

$$R_1 = \frac{\sum_{i=1}^{7} \frac{\overline{K}_i}{\overline{C}_i}}{7} \quad \text{and} \quad R_2 = \frac{\sum_{i=1}^{7} \frac{(p_i+h_i+e_i)}{(\overline{C}_i/\overline{D}_i)}}{7},$$

where

$$\overline{K}_i = \frac{\sum_{t=1}^{276}\left(\frac{\sum_{j\neq i} K_{jit}}{6}\right)}{276}, \quad \overline{C}_i = \frac{\sum_{t=1}^{276} C_{it}}{276}, \quad \text{and}$$

$$\overline{D}_i = \frac{\sum_{t=1}^{276} D_{it}}{276}.$$

Here, we first calculate the ratios of the average switching cost to the average production cost per period for each product across the entire data set. We then define R_1 as the average of the ratios across the seven products. Next, we calculate the ratio of the inventory costs consisting of unit backorder, and holding and handling costs to the average unit production costs and similarly define R_2 as the average of these ratios across these products. For the reference data set, $R_1 = 3.8$ and $R_2 = 2.1$.

We tested how sensitive our heuristics and bounding techniques were to the scale of these cost parameters. To perform this analysis, across all the products, we scaled the switching costs by 1/3, 1/2, 2, and 3 (i.e., changing R_1 by these factors), scaled the inventory cost by the same factors (i.e., changing R_2 by the same factors), and finally scaled the production costs by these factors. Our scaling factors were chosen by roughly estimating such costs across industries like petrochemicals, food processing, and pharmaceuticals based upon informal discussions with managers in these industries. This scaling in effect ensured that the cost proportions of this data were representative across this spectrum of industries. Note that our scaling procedure results in 64 (i.e., $4 \times 4 \times 4$) data sets generated from the reference problem.

For each of the 64 data sets, we used simulation to determine the production yield realization in each period. We next used the operational, greedy, and Lagrangian heuristic to develop the production plan and its associated expected cost corresponding to these yield realizations. We also computed a lower bound on expected costs for each data set using the scheme developed in §3. All of these analyses were done using Matlab (Mathworks Inc. 1998) and a specialized C program. Each run was solved within few minutes on a Dell desktop PC. We define the performance gap of the heuristic as the increase in the expected 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 operational heuristic ranged from 3.5 to 10 percent averaging around 7 percent, the gap for the greedy heuristic ranged from 6 to 20 percent averaging around 12 percent, while the gap of the Lagrangian heuristic ranged from 2.5 to 9 percent averaging around 4.5 percent. Due to the inherent randomness in yield, it was not possible to find an optimal solution to even small-sized test problems. Thus, it was difficult to assess the portions of the gap attributable to the heuristic and to the lower bound respectively.

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 production and switching costs are higher or while inventory storage and backorder costs are lower than the reference case. Conversely, the gaps are significantly lower when production and switching costs are lower, or while inventory storage and backorder costs are higher than the reference case. It is important to note that these gaps were reduced in the operational and Lagrangian heuristic largely because the lower bounds became tighter. For instance, the average solution provided by the lower bounds increased by an average of 2.8 percent, while the average solution provided by the operational heuristic decreased by 0.8 percent, and the Lagrangian heuristic decreased by 0.4 percent. This indicates that both of these heuristics are fairly stable across a wide range of variation in cost parameters, while there could still be potential to improve the lower bounds. However, we found that the average solution by the greedy heuristic decreased by 4 percent, suggesting that the performance of this heuristic was very sensitive to the relative size of the cost parameters.

In general, the performance of a heuristic to solve this problem is determined by how well it can incorporate the planning horizon, costs, and process constraints such as inventory-storage limits and yield variability. We found that the expected costs corresponding to the Lagrangian heuristic was on an average 6 percent lower than the greedy heuristic. The greedy heuristic considers costs directly and the process constraints indirectly through these costs and actual inventory levels. In addition to incorporating these aspects, the Lagrangian heuristic considers the entire planning horizon and anticipates any changes in the known demand patterns during this period. We also found that the average expected cost of the Lagrangian heuristic was around 3 percent lower than the operational heuristic. We believe this is because the Lagrangian heuristic, unlike the operational heuristic, explicitly considers costs and the entire planning horizon. Finally, we found that the operational heuristic outperformed the greedy heuristic on average by 2.5 percent. While both these heuristics are myopic, we found that since the operational heuristic is not sensitive to the scale of the cost parameters and is always within process constraints, it is better able to cope with varying demand patterns within the planning horizon. In light of these results and the fact that the Lagrangian heuristic captures all of the aspects of the greedy heuristic in addition to considering the entire planning horizon, we restrict our attention to the operational and Lagrangian heuristic in future analyses.

Other factors that could potentially affect the performance of the heuristics and the lower bound include the degree of variation of demand during the planning horizon and the level of yield variability. To analyze the effect of these parameters, we calculate the average coefficient of variation of the weekly demand for each product in the reference data and define θ_D as the average coefficient of variation across all the products. Similarly, we calculate the coefficient of variation of the yield distribution for each product and let θ_Y represent the average across all products. For the reference data set, we found that $\theta_D = 0.3$, while $\theta_Y = 0.4$. We scaled $\{\theta_D, \theta_Y\}$ by factors 1/3, 1/2, 2, and 3 generating 16 additional data sets from the reference problem.

Across these data sets, the performance gap of the operational heuristic averaged around 10 percent ranging from 6 to 15 percent. We observed that as θ_D increased the gaps increased, as the myopic nature of the heuristic was unable to capture the known, but variable, demand patterns. However, the lower bound did not change appreciably in these cases. The performance gap of the Lagrangian heuristic ranged from 3 to 10 percent averaging around 5 percent. Here, we observed that as the level of yield variability increased, both the costs of the lower bound and the heuristic increased. Finally, we found that across these data sets, the expected costs of the operational heuristic were around 4 percent higher than that of the Lagrangian heuristic. As expected, these costs were much higher than the average for the larger values of θ_D. On the other hand, when θ_D was scaled by 1/3 to reach its smallest value and θ_Y was scaled by 3 to achieve its largest value, we found that the expected cost of the operational heuristic was 1 percent lower than the Lagrangian heuristic. This shows that for processes with very small levels of demand variation within the planning horizon, but with an extremely high degree of yield variability, an operational heuristic using actual volumes and process constraints may be more effective than the Lagrangian heuristic that uses a distribution of yield realizations. However, for most real applications, we can expect the Lagrangian heuristic to outperform the operational heuristic as yield variability is seldom allowed to drift to such levels due to its negative impact on process productivity and product quality.

Generally in a stochastic-decision problem, it is not valid to judge the quality of a decision based on an outcome, as due to randomness a good outcome does not necessarily imply a good decision. However, in this study, since the evaluation of the heuristics were extensive, based on 64 sets of data over a three-month period, we are confident that they would perform well in a real application. In the final analysis, the real measure of performance of the heuristics is the quality of the decisions based on its solution, a question we consider in the Application (§6).

6. APPLICATION

We have implemented the operational heuristic as part of a broader research project described in Rajaram et al. (1999) to Cerestar, Europe's leading manufacturer of wheat- and starch-based products such as glucose, sorbitol, dextrose, and gluten with annual sales exceeding $2 billion. These products are used extensively as components in the food-processing industries (e.g., breweries, confectioneries, and bakeries), consumer-product industries (e.g., cosmetics and toothpaste), and other industries such as paper, pharmaceuticals, textiles, and specialty chemicals.

To produce these products, Cerestar operates over 20 different types of industrial-scale processes in 16 plants located in nine countries. These can be broadly classified into physical processes such as refining, separation, grinding, and extracting, and chemical processes, such as hydrogenating and modifying starch products. Since building these processes requires major capital investments, it is crucial that they constantly produce high volumes of output at the correct quality level. To achieve this goal, these processes are characterized by high degrees of automation, are operated continuously producing one product at a time, and are usually shut down only a few times a year for scheduled maintenance. As product changeovers result in long downtimes and considerable setup costs, products are produced in long campaigns and inventoried subject to buffer limits.

Since the mid-eighties, Cerestar has been aggressively improving its process-control and product-planning systems under an initiative called Social Technical Systems (STS). STS has been widely recognized as a benchmark for process improvement and has been adapted for similar processes within and across industries. While STS greatly improved the performance of existing planning systems at its processes, Cerestar believed that downtimes and costs due to switchovers, holding, and lost sales could further be reduced using advanced analytical models, which is what led to the involvement of the first author. Such reductions were critical in this industry, as products are commodities with market-defined prices and profits can be increased only by reducing costs and increasing outputs by minimizing downtimes.

We focused on the refining processes at Cerestar because refined products accounted for a large part of total profits. The production planner responsible for a particular refinery decides which product to produce during which shift based on demand requirements and inventory positions and provides this input to operators at the beginning of a shift. Product changeovers, if any, are conducted only at the beginning of an eight-hour shift. This policy reduces the overall level of control and coordination complexity during the shift and offers the potential for the operator to better understand and learn which parameter settings work well for a given product across the entire shift.

We were provided with all of the data required to validate our approach during a three-month period from seven of Cerestar's largest refineries located in five countries. Note that we used the largest data set (based on a refinery at Sas van Gent producing seven basic grade products) as the reference in the computational analysis described in the previous section. We used this data across a rolling planning horizon of one week to a generate lower bound on expected costs using the approach outlined in §3 and also generated an upper bound on expected costs using the Lagrangian and operational heuristics. The percentage gaps for the seven refineries are summarized in Table 1. These results suggest that both heuristics perform very favorably; the gaps associated with the model-based Lagrangian heuristic are on average around 1.5 percent smaller. We also computed the expected costs using the existing production-planning systems. Table 2 summarizes the percentage reduction in expected costs had these heuristics been implemented in this period for each refinery. The total cost savings across all refineries using the operational heuristic would have been around $0.55 million, while the corresponding figure using the Lagrangian heuristic would have been about $0.75 million.

At this time, the management at Cerestar has implemented the operational heuristic, since the planners understood and developed confidence in this method. The glucose refinery at Sas van Gent was chosen as the first test site because it was the flagship refinery of the company,

Table 1. Performance gaps of the heuristic on sample data over a three-month period from seven refineries.

Refinery Site	Performance Gap of the Operational Heuristic (Percent)	Performance Gap of the Lagrangian Heuristic (Percent)
Sas van Gent, The Netherlands	6	4
Manchester, United Kingdom	6	5
Martorell, Spain	7	6
Castelmassa, Italy	5	3
Haubordin, France	7	5
Krefeld, Germany	5	4
Barbie, Germany	4	3

Table 2. Potential reduction in expected costs using the heuristic solution on sample data over a three-month period from seven refineries.

Refinery Site	Potential Reduction in Expected Costs from the Operational Heuristic (Percent)	Potential Reduction in Expected Costs from the Lagrangian Heuristic (Percent)
Sas van Gent, The Netherlands	6	9
Manchester, United Kingdom	7	8
Martorell, Spain	8	9
Castelmassa, Italy	7	9
Haubordin, France	6	8
Krefeld, Germany	5	7
Barbie, Germany	4	7

producing the most profitable products and perceived to have the best planning- and process-control systems. If improvements could be realized at this refinery, these heuristics could clearly result in improvements at all other sites. This heuristic was implemented for a six-month period and led to a realized cost saving of about 3 percent or $0.4 million more then a comparable six-month period when the previous system was used. This successful implementation motivated Cerestar to implement this heuristic at six other refineries. Since October 1997, production planning at all these refineries is conducted using the operational heuristic. The cost saving after implementation at each site is summarized in Table 3. The total annual cost savings at these refineries is estimated to be about 5 percent or $3 million.

We wanted to calculate the additional benefit that would have accrued had the solution to the Lagrangian heuristic associated with the smallest gap from the lower bound been implemented in this period. To estimate this value, we assumed that the actual yield realized during this period is unchanged even when the actual production times of the product may have changed under the Lagrangian heuristic. We found that the Lagrangian heuristic offered the potential to further reduce total costs by at least 2 percent or about $1 million. Cerestar is currently evaluating the feasibility of implementing the Lagrangian heuristic.

We can use the Lagrangian heuristic to evaluate the cost implications of several strategic operational decisions at each refinery. For example, we can evaluate the benefits of expanding the end-process buffers at a given refinery by increasing the total buffer level at a given refinery and evaluating the changes in expected costs corresponding to the best (i.e., lowest gap) solution to the Lagrangian heuristic. Figure 1 represents the percentage reduction of expected costs (including the fixed costs for buffer expansion) versus the percentage increase in buffer size from the base case at the Sas van Gent refinery. Increasing the buffer size in effect increases parameter B_i, and reduces the expected handling costs in the loss function. Predictably, as observed from this figure, after a certain level of expansion, the fixed costs of expansion exceed the reduction in expected costs due to handling. These results serve as a guideline to choose the level of buffer expansion.

Table 3. Realized annual cost reduction after implementation of the operational heuristic.

Refinery Site	Cost Saving (Percent)	Cost Saving (Million $)
Sas van Gent, The Netherlands	7	0.8
Manchester, United Kingdom	5	0.3
Martorell, Spain	6	0.6
Castelmassa, Italy	8	0.3
Haubordin, France	7	0.5
Krefeld, Germany	6	0.3
Barbie, Germany	5	0.2

Figure 1. Reduction in total expected costs with increase in buffer volumes at the Sas van Gent refinery.

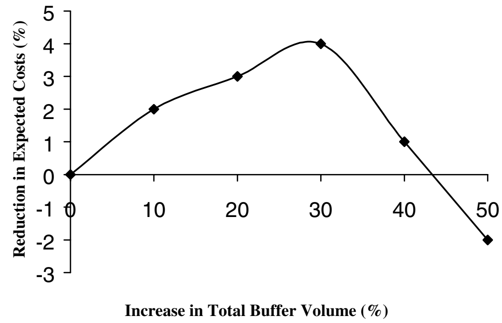

Second, we are able to quantify the benefit of reducing yield variability at a given refinery. This analysis is useful in justifying initiatives used to achieve variability reduction, including process-technology choice, better process-control technology, and better operational procedures for control as those described in Rajaram et al. (1999). To perform this analysis, we used the solution from the Lagrangian heuristic associated with the lowest gap to calculate expected costs for different levels of variability for a_i, the yield distribution for the ith product. Figure 2 represents the percentage reduction in expected costs when the coefficient of variation of the yield distribution is systematically reduced by the specified percentage for all products at the Sas van Gent refinery. This figure shows that even a small reduction in variability leads to a significant reduction in expected costs and larger reductions of yield variability often lead to only marginal levels of cost reductions. This suggests that rather than investing in capital-intensive process technology to achieve radical reduction in yield variability, the major cost benefits can be derived by focusing on operational procedures to achieve incremental reductions in yield variability. We also performed this analysis in all the other refineries to assess the benefit of achieving variability reduction. Here again, the major cost benefits from variability reduction can be derived from using operational procedures rather than investments in process technology to achieve this goal.

Figure 2. Reduction in expected costs with reduction in variability at the Sas van Gent refinery.

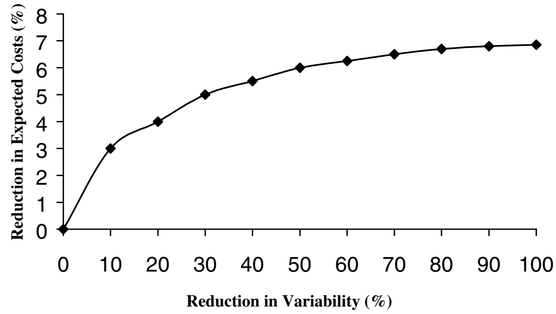

Finally, we can use the Lagrangian heuristic to develop a framework to rationalize the choice of products that are produced in a given refinery. Typically, these products are referred to as basic grades and are directly produced at the refinery. Basic grades typically vary by a particular attribute known as the Dextrose Equivalence (DE) level, and customer requirements for a particular DE level are met by blending basic grades. In principle, we could just produce two basic grades representing the extremes in DE level and produce the entire range of final customer products. Such a mechanism would reduce switchover costs, but could increase the costs of blending, which include costs of storage, operation, and rework at the blenders. We used the solution of the Lagrangian heuristic corresponding to the lowest gap to calculate the expected costs for all different product configurations for the Sas van Gent refinery. Here we first choose the products representing the extreme points in the range of the DE attribute and calculate the expected costs using this heuristic and the additional blending costs that are incurred due to this configuration. We use extensive enumeration to find the optimal basic grade configuration when we produce more than the two extreme basic grades. This analysis at the Sas van Gent refinery suggests that the lowest expected cost configuration occurs when we utilize only five out of the seven basic grades currently employed at this process. Similar analysis can be performed at the other refineries to rationalize and if necessary reconfigure the choice of basic grade portfolios.

7. DISCUSSION

In this paper, we consider the product-cycling problem with uncertain yields commonly found in many process-industry applications. The key production-planning problem in these applications is to determine which product to produce in a given process in order to minimize total production, switching, inventory storage, handling, and backorder costs. We formulate this problem as a stochastic-dynamic program with mixed continuous and (0-1) integer variables. We develop lower bounds by using a "restricted Lagrangian technique" that decomposes this problem into a process-switching problem and an individual production-inventory problem for each product across the planning horizon. This lower bound solution is used to develop a model-based Lagrangian heuristic. The gap in expected costs between the lower bound and the heuristic solution is improved by using the heuristic to set Lagrangian multipliers and then using the improved lower bound to generate the heuristic in an iterative manner. We also develop a simple and easy to implement operational heuristic and a greedy heuristic. Computational results show that the Lagrangian heuristic outperforms the greedy heuristic, since it includes the features of the greedy heuristic and considers the entire planning horizon. In addition, the Lagrangian heuristic outperforms the operational heuristic. We believe that this is due to the fact that the Lagrangian heuristic, unlike the operational heuristic, explicitly considers costs and the entire duration of the planning horizon in determining the production plan.

We validate these heuristics using data from seven Cerestar refineries and found that, compared to the existing practice, these heuristics offered the potential to significantly reduce costs at these sites. The operational heuristic has been implemented at all of these sites since October 1997. The total annual realized cost savings have been estimated at 5 percent or $3 million. Tests of the Lagrangian heuristic on data from these refineries during this period have shown the potential to further reduce total costs by at least 2 percent or about $1 million. Cerestar is currently evaluating the feasibility of incorporating this heuristic. In addition, the Lagrangian heuristic has been used to evaluate the cost implications of several strategic-operational decisions at these refineries, including choosing the best level of buffer expansion, quantifying the benefits of yield variability reductions, and rationalizing the selection of basic grades produced at these refineries. This analysis has significantly influenced several strategic-operational decisions at these sites.

The organizational impact of this work has been significant. Prior to our work at these refineries, production-planning decisions were made in a more reactive and ad hoc manner, while strategic-operational decisions were based more upon subjective experience, seniority, and anecdotal evidence. In contrast, our methods have provided a systematic, rigorous, and consistent framework to approach production-planning and strategic-operational decisions. Decisions are now made in a proactive and scientific manner based upon data and analysis. This in turn has led to greater transparency and diffusion of ideas across this multinational organization and has helped in benchmarking and standardization of best practices across these sites. More details on these aspects can be found in Rajaram et al. (1999). Cerestar is currently transferring these ideas in the production planning of newly acquired refineries in North America. These methods are also being used to evaluate buffer expansion and product-line selection decisions and variability reduction initiatives at these sites and in the design and installation of new processes.

In summary, the methodology described in this paper has had a major economic, strategic, and organizational impact at this company. Cerestar expects to maintain the gains we described and to increase them continuously several years into the future.

APPENDIX

Dynamic Programming Algorithm for the Process-Switching Problem

Step 1. Set $t = n$.

For $i = 1$ to m, compute

$$W_{it}(Y_{i(t-1)}, \mu_{it}) = \mathrm{Min}_{j=1\ to\ m}\{(C_{it} - \mu_{it}) + K_{jit} + \mu_{it} S_{jit}\}.$$

Go to Step 2.

Step 2. Set $t = n - 1$.
For $i = 1$ to m, compute
$W_{it}(Y_{i(t-1)}, \mu_{it})$
$= \mathrm{Min}_{j=1 \text{ to } m}\{(C_{it} - \mu_{it}) + K_{jit} + \mu_{it} S_{jit} + W_{j(t+1)}(Y_{jt}, \mu_{j(t+1)})\}$.
Go to Step 3

Step 3. Repeat Step 2 until $t = 1$.
At $t = 1$, the optimal switching cost is $W_1(Y_0, \mu_0) = \mathrm{Min}_{j=1 \text{ to } m}\{W_{i1}(Y_{i0}, \mu_{i1})\}$.

Dynamic Programming Algorithm for the Production-Inventory Problem

Additional Notation

α_{ik} is the yield per period for product i ($k = 1, \ldots, K$ possible values).
p_{ik} is the probability of yield α_{ik}.

Step 0. Initialization.
Define the range of $I_{i(t-1)}$ from $-\sum_{t=1}^{n} D_{it}$ to $\sum_{t=1}^{n} D_{it}$.

Step 1. Solve the inventory problem at period $t = n$ (last time step).
Set $I_{i(t-1)}$ to $-\sum_{t=1}^{n} D_{it}$.
Set $X_{it} = 0$.

Step 1.1.
$\widetilde{I}_{it}^{k} = a'_{ik} X_{it} + I_{i(t-1)} - D_{it}$.
$L_{it}^{k}(X_{it}, I_{i(t-1)})$
$= h_i(\widetilde{I}_{it}^{k})^{+} + e_i(\widetilde{I}_{it}^{k} - B_i)^{+} + p_i(-\widetilde{I}_{it}^{k})^{+}$.
$L_{it}(X_{it}, I_{i(t-1)}) = \sum_{k=1}^{K} p_{it} L_{it}^{k}(X_{it}, I_{i(t-1)})$.
$\Psi_{it}(X_{it}, I_{i(t-1)}) = \mu_{it} X_{it} + L_{it}(X_{it}, I_{i(t-1)})$.
If $X_{it} < 1$, increase X_{it} by 0.01 and go to Step 1.1, else go to Step 1.2.

Step 1.2.
Set $U_{it}(I_{i(t-1)}, \mu_{it})$
$= \mathrm{Min}_{X_{it}}\{\Psi_{it}(X_{it}, I_{i(t-1)})\}$.
If $I_{i(t-1)} < \sum_{t=1}^{n} D_{it}$ increase $I_{i(t-1)}$ by 1 and go to Step 1.1, else set $t = n - 1$ and go to Step 2.

Step 2. Solve the inventory problem at period t.
Set $I_{i(t-1)}$ to $-\sum_{t=1}^{n} D_{it}$.
Set $X_{it} = 0$.

Step 2.1.
$\widetilde{I}_{it}^{k} = a'_{ik} X_{it} + I_{i(t-1)} - D_{it}$.
$L_{it}^{k}(X_{it}, I_{i(t-1)})$
$= h_i(\widetilde{I}_{it}^{k})^{+} + e_i(\widetilde{I}_{it}^{k} - B_i)^{+} + p_i(-\widetilde{I}_{it}^{k})^{+}$.
$L_{it}(X_{it}, I_{i(t-1)}) = \sum_{k=1}^{K} p_{it} L_{it}^{k}(X_{it}, I_{i(t-1)})$.
$\Psi_{it}(X_{it}, I_{i(t-1)}) = \mu_{it} X_{it} + L_{it}(X_{it}, I_{i(t-1)})$.
If $X_{it} < 1$, increase X_{it} by 0.01 and go to Step 2.1, else go to Step 2.2.

Step 2.2.
Set $U_{it}(I_{i(t-1)}, \mu_{it})$
$= \mathrm{Min}_{X_{it}}\{\Psi_{it}(X_{it}, I_{i(t-1)})\}$.
$X_{it}^{*} = \arg\min\{U_{it}(I_{i(t-1)}, \mu_{it})\}$.
$\widetilde{I}_{it}^{k} = \tilde{\alpha}_{ik} X_{it}^{*} + I_{i(t-1)} - D_{it}$.
$E_{\tilde{a}_i}\{U_{i(t+1)}(\widetilde{I}_{it}, \mu_{i(t+1)})\}$
$= \sum_{k=1}^{K} p_{ik} U_{i(t+1)}(\widetilde{I}_{it}^{k}, \mu_{i(t+1)})$.
$U_{it}(I_{i(t-1)}, \mu_{it}) = \mu_{it} X_{it}^{*} + L_{it}(X_{it}^{*}, I_{i(t-1)}) + E_{\tilde{a}_i}\{U_{i(t+1)}(\widetilde{I}_{it}, \mu_{i(t+1)})\}$.
If $I_{i(t-1)} < \sum_{t=1}^{n} D_{it}$ increase $I_{i(t-1)}$ by 1 and go to Step 2.1. Else set $t = t - 1$ and go to Step 3.

Step 3. Repeat Step 2 until $t = 1$ for each product. Then go to Step 4.

Step 4. At $t = 1$, the optimal expected production-inventory cost is $\sum_{i=1}^{m} U_{i1}(I_{i0}, \mu_{i1})$.

ACKNOWLEDGMENTS

The authors thank David Challenor, the international manufacturing director, and Franz Behlau, manufacturing director, Cerestar Benelux B.V. for their intellectual and administrative support during the conception and implementation of this project. They are also grateful to all the process-development managers and engineers, plant superintendents, and operators at the seven refinery sites who accepted and tested their ideas. They thank Andreas Robotis for programming the algorithms and conducting the trials described in the computational section of this paper. Several excellent suggestions were made by the reviewers.

REFERENCES

Bertsekas, D. P. 1995. *Non Linear Programming*. Athena Scientific, Belmont, MA.

Bollapragada, S., T. E. Morton. 1999. Myopic heuristics for the random yield problem. *Oper. Res.* **47**(5) 713–722.

Bourland, K. E., M. J. Brown, C. A. Yano. 1997. The product cycling problem under increasing yield rates. *IIE Trans.* **29**(5) 417–422.

Dobson, G. 1987. The economic lot-scheduling problem. *Oper. Res.* **15**(5) 764–771.

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

Fleischmann, B. 1990. The discrete lot-sizing and scheduling problem. *Eur. J. Oper. Res.* **44**(3) 347–348.

Graves, S. 1980. The multi-product production cycling problem. *AIEE Trans.* **12** 233–240.

Karmarkar, U. S. 1987. The multilocation multiperiod inventory problem: Bounds and approximations. *Management Sci.* **33**(1) 86–94.

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

——, Jin-Sung Yoo. 1994. The stochastic dynamic product cycling problem. *Eur. J. Oper. Res.* **73** 360–373.

Kleindorfer, P. R. 1973. Linear convex problems with partially observed states. *Proc. Fifth Annual Sympos. on System Theory*. Duke University Press, Durham, NC, 1–6.

——, K. Glover. 1973. Linear convex stochastic optimal control with applications in production planning. *IEEE Trans. Automatic Control* **14**(2) 56–58.

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

Pochet, Y., L. A. Wolsey. 1991. Solving multi-item lot-sizing problems using strong cutting planes. *Management Sci.* **71**(1) 53–67.

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. Special Issue: Franz Edelman Award Papers. *Interfaces* **29**(1) 30–48.

Salomon, M. 1990. Deterministic lotsizing models for production planning. Unpublished Ph.D. thesis, Erasmus University, Rotterdam, The Netherlands.

——, L. Kroon, R. Kuik, L. N. VanWassenhove. 1989. The discrete lotsizing and scheduling problem. Management Report Series 30. Rotterdam School of Management, Rotterdam, The Netherlands.

Yano, C. A., H. L. Lee. 1995. Lot sizing with random yields: A review. *Oper. Res.* **43**(2) 311–334.