Structured Approach to Storage Allocation for Improved Process Controllability

J. D. Kelly and J. F. Forbes
Dept. of Chemical and Materials Engineering, University of Alberta, Edmonton, AB T6G 2G6, Canada

An optimization-based approach is presented that determines how plant feedstocks should be allocated to storage when there are fewer storage vessels than feedstocks. It is assumed here that material from the storage vessels will be subsequently blended for processing in downstream processes. The objective of the feedstock allocation strategy is chosen to ensure maximum flexibility for downstream process operation. Given the stated objective for feedstock allocation and the physical constraints, the feedstock storage allocation problem is posed in optimization form. The solution of the resulting singular-value optimization problem is discussed in terms of semidefinite programming techniques. The ideas presented are illustrated using a crude-oil storage case study. Finally, a number of observations regarding useful extensions to the proposed methods are presented.

Introduction

Many industrial plants have been specifically designed to process a range of feedstocks into a variety of products (e.g., oil refining, petrochemicals, pharmaceuticals, food, and beverage). The need to reduce plant capital and operating expenditures has driven many manufacturers to a "just-in-time" processing philosophy that minimizes, as much as is possible, both feedstock inventory and available storage (or ullage). Thus, such plants must address two competing objectives: maintaining (or improving) process operating flexibility while keeping the feed storage facilities to a bare minimum. These competing objectives often create the situation where a limited number of storage vessels are used to contain a larger number of process chargestocks.

When these plants produce a range of products, care must be taken in allocating the different feedstocks into the appropriate storage vessels. One of the objectives for any feedstock allocation strategy should be to ensure that operating flexibility is properly maintained for the downstream units to which these materials are blended and then fed. To motivate discussions, an oil refining example is used throughout this article. The example problem investigates a crude-oil storage allocation problem and is used to clarify the ideas presented here. Despite the focus of the illustrative example, the proposed methods apply to a wide range of storage allocation problems.

Then for illustrative purposes, consider an oil refinery that purchases the crude-oil slate given in Table 1 for a certain processing period, in amounts w_i. Figure 1 illustrates the crude-oil storage problem. In Figure 1, there are 12 incoming crude oils (i.e., $m = 12$) to be stored in p storage vessels. Perhaps the simplest solution to the feedstock storage allocation problem is to maintain separate storage facilities for each crude oil. This would then require 12 storage vessels (i.e., $p = 12$), all suitably sized to ensure sufficient storage capacity for the expected range of crude-oil slate variation. Although

Table 1. An Example Crude-Oil Slate

Crude Oil	Vol. (m^3)
1	814.2
2	1,576.1
3	2,393.5
4	837.2
5	1,989.9
6	1,928.5
7	1,475.8
8	1,687.1
9	828.6
10	568.4
11	221.2
12	1,579.0

Correspondence concerning this article should be addressed to J. F. Forbes.
Current address of J. D. Kelly: Comcept Canada Inc., 275 Renfrew Drive, Suite 101, Markham, Ont., Canada L3R 0C8.

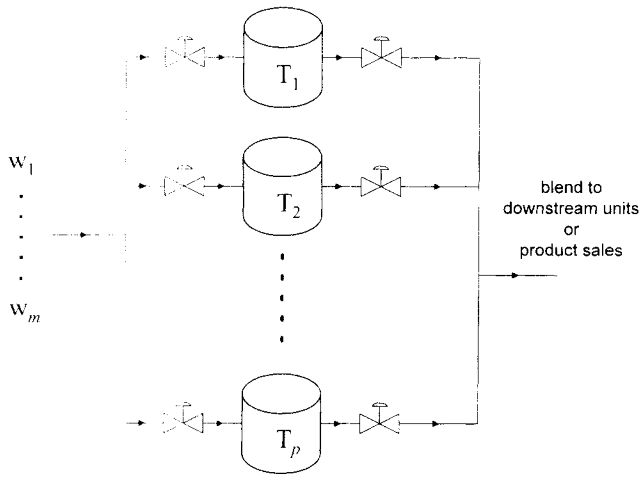

Figure 1. Typical plant storage facility.

this solution offers significant flexibility of operation to produce suitable crude-oil blends for the downstream process units, it also has the maximum possible capital cost associated with it. As a result, oil refiners are choosing to find solutions to the storage allocation problem that require fewer storage vessels than feedstocks (i.e., they are choosing the situation where $p < m$).

In Figure 1, when there are fewer available feed storage vessels than incoming crude oils (i.e., $p < m$), the incoming crude oils must be blended into the available tanks. In this situation, one solution to the storage allocation problem is to blend the 12 crude oils equally into the available storage vessels. This solution, while having the potential for minimum capital and maintenance costs, also minimizes the flexibility of the blend controller to prepare a wide range of crude-oil blends, which may adversely affect the operation of downstream processes. Rather than dividing the incoming crude oils equally between the available storage vessels, these feedstocks can be allocated to storage such that the flexibility to produce a wide range of blends is maximized while meeting any operating constraints (e.g., sulfur content).

The intent of this article is to introduce a structured approach for allocating feedstocks to storage, when there are more feedstocks than available storage vessels. The approach taken in this work is to pose the storage allocation problem in a manner that ensures operating flexibility of the downstream process units, while satisfying specified operating constraints. Relevant control concepts and analysis tools are discussed, as well as a constrained maximization problem, which is formulated to reflect the key objectives of the allocation problem and whose solution is discussed in terms of current optimization methods. The techniques in this article are illustrated with a detailed crude-oil storage case study. The article also discusses possible extensions of the current work and future research directions.

Background and Literature Review

This section of the article provides the necessary background for developing the proposed formulation of the storage allocation problem. The first subsection reviews the available control literature to provide a useful metric for determining the "best" storage allocation strategy from among a set of alternative strategies. This section concludes with a discussion of how the steady-state gain matrix can be developed for blending control problems.

Storage allocation objectives

Morari (1983) introduced the term "resiliency" to describe the ability of a controlled process to move quickly and smoothly from one operating condition to another, and to deal effectively with disturbances. This resiliency is measured in terms of the size of the minimum singular value of the process steady-state gain matrix (Morari, 1983). Mathematically, a resilient process is one that attains a large minimum singular value for the steady-state gain matrix (i.e., at zero frequency). Processes with large minimum singular values are less susceptible to manipulated variable saturation than are processes with small minimum singular values. Further, it has also been shown that processes with large minimum singular values are more robust (or insensitive) to process/model mismatch than those with smaller minimum singular values (Johnston and Barton, 1984; Grosdidier et al., 1985; Koung and MacGregor, 1992). Then, it would seem that a key objective in the feedstock storage allocation problem should be to provide the blending operation with as resilient a process as possible (i.e., one that has the largest possible minimum singular value).

To illustrate the importance of the minimum singular value of the steady-state gain matrix, consider the steady-state process model:

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}, \tag{1}$$

where $\boldsymbol{A}$ is the steady-state process gain matrix, $\boldsymbol{x}$ is the vector of manipulated inputs, and $\boldsymbol{b}$ is the vector of controlled outputs. (This is a departure from standard control nomenclature, but it is done here to facilitate later discussions of the storage allocation problem.) Upper and lower bounds on the size of required control actions for a desired output can be determined using well-known results from linear algebra (Golub and Van Loan, 1989):

$$\sigma_n(\boldsymbol{A}) \| \boldsymbol{x} \|_2 \le \| \boldsymbol{A}\boldsymbol{x} \|_2 = \| \boldsymbol{b} \|_2 \le \| \boldsymbol{A} \|_2 \| \boldsymbol{x} \|_2 \le \sigma_1(\boldsymbol{A}) \| \boldsymbol{x} \|_2, \tag{2}$$

where σ_1 and σ_n denote the maximum and minimum singular values of the steady-state gain matrix $\boldsymbol{A}$, respectively. Then, the relationships in the inequality (Eq. 2) can be rearranged to yield an upper bound on the size of control action required for any given output:

$$\| \boldsymbol{x} \|_2^{\max} = \frac{\| \boldsymbol{b} \|_2}{\sigma_n(\boldsymbol{A})}. \tag{3}$$

Equation 3 clearly shows that a small minimum singular value of the steady-state gain matrix will produce a large upper bound on required control action $\boldsymbol{x}$ for any desired output $\boldsymbol{b}$. Equation 3 does not state that for a specific value of the desired output vector $\boldsymbol{b}$, the required input vector $\boldsymbol{x}$ will necessarily take on the upper-bound value. This, of course, depends upon the degree to which the required input vector x is aligned with the right singular vector associated with the minimum singular value σ_n. However, Eq. 3 illustrates the importance of working to ensure that the process control system is designed such that the control objectives can be met with as small a control effort as possible (i.e., minimize the upper bound on the input vector x by maximizing the minimum singular value).

A further consideration is the ability of the controller to be able to achieve offset free control of its target specifications subject to any manipulated input limits. The difficulties associated with physical limits on the manipulated variables can be somewhat mitigated by ensuring that the required control action to achieve the control objectives is generally as small as possible (i.e., minimize the upper bound on the input vector x). Thus from a controllability standpoint, there is a clear incentive to maximize to the fullest extent possible the minimum singular value of A in the design of any process control scheme.

Control system designs that ensure a large minimum singular value of the steady-state gain matrix have also been shown to tolerate more additive uncertainty between the process model and the plant, while still maintaining the stability of the closed-loop system (Koung and MacGregor, 1992). In their work, Koung and MacGregor (1992) have shown that closed-loop integral stabilizability is guaranteed if

$$\| A_P - A \|_2 < \sigma_n(A), \tag{4}$$

where A_P is the true plant steady-state gain matrix. Hence, the larger the value of the minimum singular value, the more model mismatch that can be tolerated by the controller. In addition to robustness considerations, a number of researchers have stressed the importance of ensuring the largest possible minimum singular value of the steady-state gain matrix for accurate process identification (Skogestad and Morari, 1987; Koung and MacGregor, 1993; Li and Lee, 1996).

Based on controllability, robustness, and identifiability arguments, a key objective of the storage allocation problem should be to ensure that the blend controller is presented with a process that has the largest singular value of the steady-state gain matrix. Unlike many process control situations, selecting how incoming feedstocks will be allocated to storage provides the opportunity to directly manipulate the steady-state gain matrix for the subsequent blend control problem.

Steady-state gain matrix construction

Currently there exists a considerable literature on blend control. Much of this literature has concentrated on control of blending in the production of automotive fuels (e.g., Agrawal, 1995, and DeWitt et al., 1989); however, the general concepts presented within the literature apply to a broad spectrum of blending problems. A common approach in blending is to treat the problem as a steady-state control problem, where the primary objective is to satisfy process constraints (blend quality specifications, component availability, and product demand), and the secondary objective is to produce an economically attractive blend recipe. Figure 2 illustrates the general blending control problem.

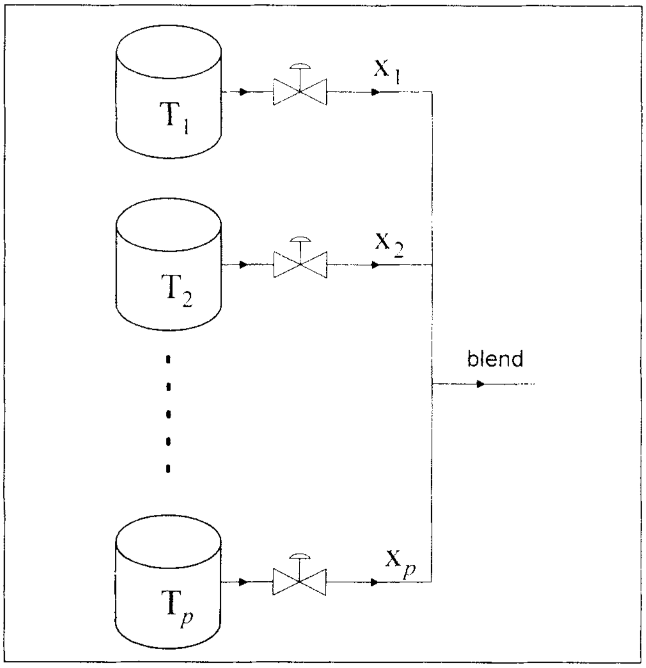

Figure 2. Blending control problem.

As indicated in Figure 2, the blending control problem is to determine the flow rates x_i from each tank such that the blended qualities are satisfied. This control problem is often formulated as

$$Ax \leq b, \tag{5}$$

where A is a matrix containing the linear blending indices for each of the component streams, x is a vector containing the flow rates from each tank, and b is a vector containing the blended quality specifications for the product stream. In this formulation, each column of the matrix A represents the linear blending indices for every component quality in a specific tank.

This article deals with the situation where there are more feedstocks than available storage vessels. These feedstocks are often purchased using some optimal planning strategy, which is commonly based on linear programming. In this strategy, physical storage limitations are often ignored (or the available storage volumes are lumped together) and the decision variables for optimization purposes are the amounts of each feedstock to be purchased during the planning period (Bodington, 1995). In a well-posed linear program the solution lies at a vertex of the feasible region and as a result, the number of active constraints at the solution is identical to the number of decision variables. Thus, when there are more purchased feedstocks than available storage vessels, the blend controller may not be capable of enforcing all of the active constraints in the optimal purchase plan. (Recall that the decision variables for the blend control problem are the flows from each storage tank.) Mathematically this can be stated as

$$\{\dim(x) = p\} < \{\dim(b) = m\}. \tag{6}$$

Then, the blend control problem may be solved by selecting some key subset of the active constraints that are considered

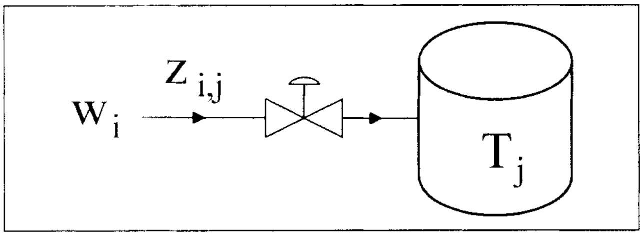

Figure 3. Storage allocation problem.

most vital to enforce. In this situation, the reduced constraint set is chosen to provide a square A matrix. Thus, the process model for blending control purposes is often reduced to the form:

$$Ax = b. \tag{7}$$

In this formulation, elements of the blending index matrix A depend upon the qualities and amounts of each feedstock that have been stored in a particular vessel. Consider the jth storage vessel shown in Figure 3. Determining the properties of the mixture in storage is analogous to the blending problem previously discussed. Assuming a linear blending index approach is used to determine the mixture properties in a specific storage vessel, the blended properties of any specific mixture (given by the flow rates in x) can be determined from the feedstock properties as follows:

$$QZV^{-1}x = b \tag{8}$$

$$V = \text{diag}(v), \tag{9}$$

where the matrix Q contains the linear blending indices for each of the feedstocks, the matrix Z contains the allocation of feedstocks to storage vessels, and the vector v contains the inventory in each storage vessel. In Eq. 8, the (i, j) elements of the matrix Z represents the allocation of the ith feedstock to the jth tank. The matrix Q is organized such that the (i, j) element represents the blending index of the ith property for the jth feedstock. Finally, it is important to note that the blend control problem must usually ensure that the inventory in each tank is consumed during the plan period. Thus the nominal flow set point from each tank can be set as

$$x = \frac{v}{\Delta t}, \tag{10}$$

where Δt is the plan period.

In Eq. 8 the decision variables for the blending control problem are the flows from the individual tanks, represented by the elements of the vector x. The decision variables for the storage allocation problem are the elements of the allocation matrix Z and inventory vector v. It is clear from Eq. 8 that the performance of the blend controller will depend on the allocation strategy employed in storing the feedstocks. As discussed in the preceding subsection, the key objective for the allocation strategy (i.e., in determining the allocation matrix Z and inventory vector v) should be to allow the maximum possible flexibility for blend controller, while meeting all of the required constraints.

Allocation Problem Formulation and Solution

The previous section of this article dealt with the key issues that should be considered in formulating any storage allocation strategy. As previously discussed, the objective of such a strategy should be to ensure the maximum flexibility available for the operation of subsequent processes. Using this objective and an understanding of the blending control problem, the following two subsections will develop the full formulation of the optimization-based storage allocation problem and discuss its solution in terms of available optimization algorithms.

Allocation problem formulation

As previously stated, any storage allocation strategy must satisfy a number of constraints. A constraint is required to ensure that all of the purchased feedstocks are allocated to storage vessels and to determine the amount of material stored in each vessel. This can be expressed as

$$Z\mathbf{1} = w \tag{11}$$

$$Z^T\mathbf{1} = v, \tag{12}$$

where w is a vector containing the amounts of each feedstock purchased, v is a vector containing the amount of material stored in each vessel, and $\mathbf{1}$ is a vector of appropriate length containing only ones (i.e., $\mathbf{1} = [1\ 1 \ldots 1]^T$). Note that the rows of matrix Z contain the amount of a given feedstock allocated to each storage vessel.

As discussed in the previous section, one objective for any storage allocation strategy is to maximize the minimum singular value of the matrix product QZV^{-1}, which provides the subsequent blending operation with the largest possible resilience. Then, the storage allocation problem can be formulated as an optimization problem:

$$\begin{aligned} &\max_{Z,v} \sigma_n(QZV^{-1}) \\ &\text{subject to:} \\ &Z\mathbf{1} = w \\ &Z^T\mathbf{1} = v \\ &V = \text{diag}(v) \\ &Z_{\text{lower}} \le Z \le Z_{\text{upper}} \\ &v_{\text{lower}} \le v \le v_{\text{upper}}, \end{aligned} \tag{13}$$

where Z_{lower} and Z_{upper} denote the limits on the amount of each feedstock to be allocated to individual vessels, and v_{lower} and v_{upper} denote the storage limits for each vessel. Note that the storage limits of any individual vessel may be determined by either the vessel capacity or by pumping constraints to (or from) the vessel. Limits for elements of the allocation matrix arise in situations where certain feedstocks may not be suitable for storage in some of the vessels, for example, if the feeds are in different phases.

This problem (Eq. 13) is a nontraditional nonlinear programming problem (NLP) since the objective function depends on a matrix property. Solution strategies for this problem are discussed in the next section of this article. In Eq. 13 the decision variables are the values contained in the feedstock allocation matrix Z (where $Z \in \Re^{m \times p}$) and the inventory amounts for each vessel contained in the elements of v (where $v \in \Re^{p}$). The availability constraints represented by Eq. 11 consume m degrees of freedom from the optimization problem and the equality constraints used to determine the physical inventory in each vessel consume a further p degrees of freedom. Then, before considering the storage limitations and the simple bounds on the elements of Z, there are at most $m \times (p-1)$ degrees of freedom for optimization (assuming that the allocation constraints are linearly independent). It is these extra degrees of freedom that provide the opportunity to construct a storage allocation strategy that yields a resilient subsequent blending operation. Finally, as expected there are no degrees of freedom available for optimization purposes in the single storage vessel situation.

Allocation problem solution

The objective in Eq. 13 is to maximize the minimum singular value of the matrix product QZV^{-1}, subject to linear constraints on the elements of the allocation matrix Z and the tank inventories v. Such problems belong to the broad class of eigenvalue optimization problems (Lewis and Overton, 1996). Recall that the singular values of a given matrix can be determined from the eigenvalues of a symmetric matrix product (of the given matrix with itself) (Golub and Van Loan, 1989).

Although eigenvalue optimization problems are nondifferentiable (Shapiro and Fan, 1995; Vandenberghe and Boyd, 1996), they are convex and a variety of algorithms are available for their solution. Solution methods for eigenvalue optimization problems, and the more general class of optimization problems that can be reformulated as linear matrix inequality (LMI) constrained optimization problems, are called semidefinite programming (SDP) methods (Nesterov and Nemirovskii, 1994; Boyd et al., 1994; Vandenberghe and Boyd, 1996; Lewis and Overton, 1996). Although Eq. 13 does not appear to be a conventional SDP formulation (since the decision variables are organized in the matrix Z and the vector v), it can be transformed to yield a structure that more closely resembles the required SDP form. However, as shown in the Appendix, the transformed problem still is not in the traditional SDP format since the decision variables do not enter the objective function linearly. This did not present insurmountable difficulties for the selected solver.

There are a large and rapidly growing number of algorithms available for the solution of SDP problems (e.g., Alizadeh et al., 1997; Gahinet et al., 1995; Luo et al., 1998). In this study the EA3 ellipsoid algorithm (Kupferschmid, 1984; Ecker and Kupferschmid, 1985; Warren et al., 1987) was used. This is an older algorithm, but it was chosen based on a number of factors: (1) previous experience with the algorithm, (2) speed and problem size were not a factor in this study; and (3) it was recognized that the allocation problem would be solved infrequently. It is not the intent of this article to recommend any particular solution strategy for Eq. 13 or to compare available methods; the main objective of this article is to formulate the storage allocation problem and to discuss possible means of solving it. Selection of a solver for a given application is left to the reader.

Solution of Eq. 13 required that a number of implementation issues of the ellipsoid algorithm be addressed. First, the ellipsoid algorithm does not handle the equality constraints in Eq. 13 explicitly. Consequently, the implementation of the algorithm required transformation of the equality constraints into narrowly bounded inequality constraints using appropriately chosen relaxation tolerances. Second, the ellipsoid volume that is successively shrunk at each iteration must be defined. For this study the ellipsoid volume was initialized and stored as a vector u representing a symmetric matrix in EA3 algorithm:

$$[u]_{[i \times (i+1)]/2} = \frac{m \times p}{4} \epsilon^2 \;\; \forall\, i = 1, \ldots, m \times p, \qquad (14)$$

where ϵ is an adjustable positive number determining the ellipsoid's initial size (Kupferschmid, 1984).

Crude-Oil Storage Case Study

The general methods presented in the preceding section are applicable to a wide range of storage allocation problems. One such method is the allocation of purchased crude oils to storage in an oil refinery. The trend toward minimizing both inventory and storage costs has caused many oil refiners to reduce both the number of storage tanks and the total amount of crude-oil storage capacity within their refinery. This section will illustrate the proposed approach using a crude-oil storage case study.

Case study description

In most oil refineries, individual crude oils are purchased and placed in storage. Often oil refiners are faced with the situation where the number of crude oils purchased during a planning period exceeds the number of available storage tanks. In this situation, some of the crude oils must be mixed in storage tanks. Eventually the crude oil mixtures in storage are blended to form a feed stream for the refinery's atmospheric and vacuum distillation units. This process is illustrated in Figure 1.

In this case study, the allocation of the twelve purchased crude oils (as shown in Table 2) to three crude-oil storage tanks is discussed (i.e., $m = 12$, $p = 3$). Given that there are only three manipulated variables for blending from storage, only three of the active constraints in the purchase plan can be enforced without offset by the crude-oil blending controller. For the purposes of this study, the three constraints that have been deemed most important (and that can be easily influenced by the flow of crude-oil mixtures from each tank) are (1) a volume constraint on the amount of naphtha that can be produced by the atmospheric and vacuum distillation units of 1,272 m^3; (2) a constraint on the concentration of sulfur contained in the diesel stream of 0.3000 wt.%; and (3) a volume constraint to ensure that all of the purchased crude oil (15,898 m^3) is used during the plan period.

The three given constraints were considered essential to the effective operation of the refinery. Further, they were selected for enforcement by the blend controller since downstream unit operating conditions such as internal stream temperatures, pressures, flows, and so forth have negligible abil-

Table 2. Crude-Oil Storage Case Study Data

Crude Oil No.	Napththa Stream Yield (vol. %)	Diesel Stream Density (SG)	Diesel Stream Yield (vol. %)	Diesel Stream Sulfur (wt. %)	Purchased Crude-Oil Vol. (m^3)
1	6.29	0.8805	16.42	1.0510	814.2
2	8.71	0.8526	18.19	0.3140	1,576.1
3	2.96	0.8582	38.46	0.0950	2,393.5
4	8.12	0.8657	17.64	0.9290	837.2
5	6.52	0.9014	23.37	0.1250	1,988.9
6	8.18	0.8734	19.56	0.0760	1,928.5
7	7.68	0.8824	20.56	0.4100	1,475.8
8	8.30	0.8551	18.36	0.2410	1,687.1
9	25.47	0.8735	11.91	0.3480	828.6
10	5.81	0.8761	15.02	1.3740	568.4
11	9.52	0.8962	11.88	1.8430	221.2
12	8.78	0.8605	18.69	0.3090	1,579.0

ity to enforce them. Thus, these three constraints cannot be effectively enforced except though crude-oil purchases and storage allocation. As an example, the amount (or yield) of naphtha leaving the crude unit can be marginally affected by the operation of the atmospheric distillation column; however, the amount of naphtha contained in the crude-oil mixture charged to the crude distillation unit provides a wider range of controllability.

Table 2 contains the information required to build the naphtha and diesel constraints for the blending controller and Q matrix used in the objective function of the storage allocation problem (Eq. 13). The inequality constraints in Eq. 13, as defined for this crude-oil allocation case study, were storage-capacity limits of each tank and bounds for the amount of each crude oil that can be allocated to specific storage tanks. Since the three tanks were assumed to be identical in size and pumping capacity, the upper and lower limits for each element of the inventory vector v were set at 9,539 and 1,590 m^3, respectively. In this case study, the only bounds placed directly on the elements of the allocation matrix Z were nonnegativity restrictions. No other allocation restrictions were placed on the problem; however, these could have been easily added with no significant computational effect on the problem.

Case study solution and results

Before discussing the results obtained from the proposed storage allocation approach, consider the heuristic approach that many oil refineries use in which crude oils are segregated into available storage tanks based on an overall bulk property (e.g., sulfur content). A common approach of allocating crude oils to storage is based on the total amount of reactive and nonreactive sulfur compounds found in the oil. A typical refinery might designate each tank as either being a high-medium or low-sulfur tank. Then as a crude oil is received, it is allocated based on these predefined designations and without considering what constraints the controller must enforce. Table 3 gives the crude-oil segregation information required to heuristically segregate the crude oils used in this case study. For discussion purposes, this allocation strategy is termed the three-tank heuristic allocation.

Table 4 gives the calculated results for the optimization-based allocation strategy when three storage tanks are available. Table 5 presents the minimum singular values of the blending constraint matrix for each of the allocation strategies and cases.

As can be seen in Table 5, the minimum singular value for the blending control matrix using the optimization-based allocation strategy given in Eq. 13 for the three tanks is almost double that for the heuristic allocation strategy based on bulk sulfur content. This clearly shows that the optimization-based allocation strategy provides the blend controller with a more resilient process, which implies better blending control performance in terms of simultaneously meeting the naphtha yield, diesel stream sulfur content, and throughput constraints.

Table 3. Bulk Sulfur Segregation Information

Crude Oil No.	Bulk Sulfur (wt. %)
1	1.7100 (H)
2	0.5300 (M)
3	0.2000 (L)
4	1.3500 (H)
5	0.2100 (L)
6	0.2400 (L)
7	0.3600 (M)
8	0.3900 (M)
9	0.2600 (L)
10	2.3600 (H)
11	3.2000 (H)
12	0.5500 (M)

H, M, L—high, medium, and low, respectively.

Table 4. Optimization-Based Allocation Strategy Results (3 Tanks)

Crude Oil No.	Tank 1 (m^3)	Tank 2 (m^3)	Tank 3 (m^3)	Allocated Crude (m^3)
1	748.6	47.1	18.4	814.2
2	1.3	1,484	91.1	1,576.1
3	8.2	57.2	2,328	2,393.5
4	69.3	742.3	25.6	837.2
5	0.8	1,540	447.7	1,988.9
6	0	224.2	1,704	1,928.5
7	3.2	1,454	18.8	1,475.8
8	5.7	1,515	166.1	1,687.1
9	4.7	764.2	59.7	828.6
10	551.1	4.4	12.8	568.4
11	201.2	14.5	5.4	221.2
12	1.6	1,510	67.7	1,579.0
v(m^3)	1,596	9,357	4,946	

Table 5. Minimum Singular Values of QZV^{-1}

No. of Tanks	Heuristic	Optimization-Based
3	6.480×10^{-4}	1.244×10^{-3}
12	2.518×10^{-3}	2.816×10^{-3}

The success of the optimization-based allocation strategy in the three-tank problem raised the question of its success when faced with the situation where there were as many storage vessels (of appropriate size) as there were incoming crude oils. In this situation, a "natural" heuristic solution to the storage allocation problem would be to place each incoming crude oil in a separate storage tank. For discussion purposes this will be termed the twelve-tank heuristic solution. To investigate the performance of the proposed optimization-based allocation strategy, the minimum inventory constraints for each tank were reduced to zero (i.e., $v_{\text{lower}} = \mathbf{0}$). This was required to ensure a feasible solution since the lower bounds used in the three-tank allocation problem could not be met given the availability of seven of the crude oils and the total amount of crude oil available. For comparison, Table 6 presents the solution to the optimization-based allocation strategy.

As can be seen in Table 5, the optimization-based allocation strategy produced a better subsequent blend control problem than the 12-tank heuristic approach. This is not surprising since the optimization-based strategy takes advantage of all of the available degrees of freedom in the allocation problem to improve the resilience of the blending control, whereas the blending control problem was not considered in the twelve-tank heuristic allocation strategy. Table 6 shows that the optimization-based allocation strategy is much more complex than the heuristic approach. Such a trade-off between allocation strategy complexity and resilience of the blending control problem should be expected. However, it is worth noting that implementation of the more complex allocation strategy is not difficult since it only occurs as individual crude oils arrive at the refinery. (Typically, crude-oil deliveries are no more frequent than twice per day.)

Finally, to emphasize that these allocation problems are not very computationally intensive, all of the calculations for the crude oil case study were performed on an IBM compatible computer (Intel 486, 50 MHz) using Fortran. Typical solution times were approximately 90 s for the three-tank problems in this case study. These are not onerous computational requirements considering that such calculations will be performed infrequently, either at the crude purchase plan frequency or perhaps more regularly as new crude oils arrive at the refinery.

Summary and Conclusions

This article has presented a structured, optimization-based approach to the allocation of plant feedstocks into storage vessels. The objective of the proposed allocation strategy is to provide a subsequent blending control problem that is as resilient as possible. This storage allocation objective can be interpreted mathematically as maximizing the minimum singular value of the blending constraint matrix. The optimization-based allocation problem can be formulated as given in the storage allocation problem (Eq. 13), which is easily recognized as belonging to the broad class of eigenvalue optimization problems. Such eigenvalue optimization problems can be solved using a variety of semidefinite programming algorithms.

The proposed optimization-based storage allocation strategy was illustrated using a crude-oil case study. In the case study, initially the situation where there were more types of crude oil than vessels available in which to store them was examined. A conventional crude-oil segregation method based on crude-oil sulfur content was compared to the proposed optimization-based approach. The proposed approach did a significantly better job in ensuring the maximum resilience of the subsequent blending control problem. The case study concluded by investigating the situation where there were as many storage vessels available as crude oils. The heuristic approach in this case would have placed each individual crude oil in its own storage vessel. The optimization-based strategy took advantage of all available degrees of freedom in the allocation problem to yield a more resilient blending control problem.

Although the storage allocation strategy presented in this article is based on a semidefinite programming problem that is nonlinear in the decision variables, the case study problems were solved using an ellipsoid algorithm. The problems in the case study had a reasonably large number of degrees of free-

Table 6. Optimization-Based Allocation Strategy Results (12 Tanks)

Crude Oil No.	T1 (m^3)	T2 (m^3)	T3 (m^3)	T4 (m^3)	T5 (m^3)	T6 (m^3)	T7 (m^3)	T8 (m^3)	T9 (m^3)	T10 (m^3)	T11 (m^3)	T12 (m^3)
1	6.1	5.3	5.8	717.7	7.4	8.8	44.3	4.6	8.6	1.4	0.3	3.8
2	2.9	77.6	19.0	4.0	32.4	2.3	1,383	32.2	17.7	0.4	1.1	3.4
3	2.3	61.1	446.1	0.5	544.3	435.5	44.0	222.8	557.6	1.0	0.7	77.5
4	0.8	18.1	1.8	747.2	10.0	3.0	33.9	2.5	10.5	4.5	0.8	4.2
5	0.3	43.6	61.8	10.3	14.3	109.8	1,317	121.1	129.1	0.2	0.2	181.5
6	1.3	724.6	53.4	6.4	42.7	19.1	190.2	5.9	68.1	1.4	0.4	814.9
7	1.4	12.3	5.5	14.2	7.1	4.2	1,336	26.0	16.2	0.9	0.3	52.4
8	3.4	86.4	34.8	0.4	5.0	26.3	1,416	61.5	45.0	0.1	0.0	8.8
9	1.0	193.6	73.1	0.0	54.4	57.6	26.1	380.4	15.2	0.3	0.2	26.7
10	370.1	4.5	0.2	127.2	4.2	0.6	16.6	2.0	4.3	27.6	7.6	3.4
11	10.8	3.1	3.9	0.7	4.0	1.8	1.0	0.1	2.3	101.2	90.7	1.6
12	2.0	62.8	27.6	12.8	6.3	15.0	1,318	48.8	30.7	0.1	1.0	53.6
v (m^3)	402.4	1,293	733.1	1,641	732.2	683.8	7,125	907.8	905	139.2	103.3	1,232

dom in comparison to what would be expected of an industrial-scale problem, yet proved to be easily solved and not computationally expensive. All computations were performed on a personal computer and solve times were less than 2 min for the three-tank problems.

The method presented in this work provides a structured approach to making storage allocation decisions and focuses on determining allocation strategy where incoming feedstocks are to be stored in standing vessels. The proposed method may also provide the basis for addressing a number of issues, such as (1) allocating incoming feedstocks to running storage vessels, and (2) design of storage facilities for some given design-basis feedstock slate. Although the approach presented in this article may provide a useful starting point for developments in these two areas, further work is required to properly address critical issues.

Finally, it should be stressed that although many of the discussions contained in this article have revolved around a specific oil-refining application, there are many other opportunities to use this technology both within and beyond oil-refining operations.

Notation

A = blending control constraint matrix ($\Re^{n \times n}$)
b = blending specifications vector ($\Re^{n}$)
n = number of chosen blending constraint specifications
x = vector of volume flows from each vessel ($\Re^{p}$)

Operators

$\| \cdot \|_2$ = spectral norm of a matrix of Euclidean norm of a vector
vec = convert a matrix to a vector

Literature Cited

Alizadeh, F., J.-P. A. Haeberly, M. V. Nayakkankuppam, M. L. Overton, and S. Schmieta, *SDPPACK User's Guide*, NYU Computer Science Dept. Tech. Rep., New York Univ., New York (1997).
Agrawal, S. S., "Integrate Blending Control, Optimization, and Planning," *Hydrocarbon Process.*, 129 (1995).
Bodington, C. E., *Planning, Scheduling and Control Integration in the Process Industries*, McGraw-Hill, New York (1995).
Boyd, S., L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*, SIAM, Philadelphia (1994).
DeWitt, C. W., L. S. Ladson, A. D. Warren, D. A. Brenner, and S. A. Melhem, "OMEGA: An Improved Gasoline Blending System for Texaco," *Interfaces*, **19**, 85 (1989).
Ecker, J. G., and M. Kupferschmid, "A Computational Comparison of the Ellipsoid Algorithm with Several Nonlinear Programming Algorithms," *SIAM J. Control Optim.*, **23**, 5 (1985).
Gahinet, P., A. Nemirovskii, A. J. Laub, and M. Chilahi, *LMI Control Toolbox User's Guide*, The MathWorks, Natick, MA (1995).
Golub, G. H., and C. F. Van Loan, *Matrix Computations*, The Johns Hopkins Univ. Press, Baltimore, MD (1989).
Grosdidier, P., M. Morari, and B. R. Holt, "Closed-Loop Properties from Steady-State Gain Information," *Ind. Eng. Chem. Fundam.*, **24**, 2 (1985).
Johnston, R. D., and G. W. Barton, "Improved Process Conditioning using Internal Control Loops," *Int. J. Control*, **40**, 6 (1984).
Koung, C.-W., and J. F. MacGregor, "Robustness of Multivariable Linear Controllers to Process Nonlinearities," *Ind. Eng. Chem. Res.*, **31**, 4 (1992).
Koung, C.-W., and J. F. MacGregor, "Design of Identification Experiments for Robust Control. A Geometric Approach for Bivariate Processes," *Ind. Eng. Chem. Res.*, **32**, 8 (1993).
Kupferschmid, M., "EA3: A Practical Implementation of an Ellipsoid Algorithm for Nonlinear Programming," Dept. of Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, NY (1984).
Lewis, A. S., and M. L. Overton, "Eigenvalue Optimization," *Acta Numer.*, **5**, 149 (1996).
Li, W., and J. H. Lee, "Control Relevant Identification of Ill-Conditioned Systems: Estimation of Gain Directionality," *Comput. Chem. Eng.*, **20**, 8 (1996).
Luo, Z.-Q., J. F. Sturm, and S. Zhang, "Superlinear Convergence of a Symmetric Priaml-Dual Path Following Algorithm for Semidefinite Programming," *SIAM J. Optim.*, **8**(1) (1998).
Morari, M., "Flexibility and Resiliency of Process Systems," *Comput. Chem. Eng.*, **7**, 4 (1983).
Nesterov, Y., and A. Nemirovskii, *Interior Point Polynomial Algorithms in Convex Programming*, SIAM, Philadelphia (1994).
Shapiro, A., and M. K. H. Fan, "On Eigenvalue Optimization," *SIAM J. Optim.*, **5**, 3 (1995).
Skogestad, S., and M. Morari, "Implications of Large RGA Elements on Control Performance," *Ind. Eng. Chem. Res.*, **26**, 11 (1987).
Vandenberghe, L., and S. Boyd, "Semidefinite Programming," *SIAM Rev.*, **38**(1), (1996).
Waren, A. D., M. S. Hung, and L. S. Lasdon, "The Status of Nonlinear Programming Software: An Update," *Oper. Res.*, **35**, 4 (1987).

Appendix

The optimization problem (Eq. 13) is not in conventional semidefinite programming form. Vandenberghe and Boyd (1996) present the conventional SDP formulation as:

$$\max_{y} c^T y$$

subject to

$$F(y) \succeq 0$$

$$F(y) \equiv F_0 + \sum_{i=1}^{m} F_i y_i, \tag{A1}$$

where F_i are $m+1$ symmetric matrices, and the symbol $\succeq$ is called the Loewner partial ordering, which indicates that the matrix $F(y)$ is positive semidefinite. Conversion of eigenvalue optimization problems into this conventional SDP form can be found in a number of sources (e.g., Vandenberghe and Boyd, 1996; Lewis and Overton, 1996). The key features to note in the problem (Eq. A1) are that both the objective function and the constraints are linear in the decision variables. Also note that the decision variables are organized into a vector.

The optimization problem (Eq. 13) presents two difficulties. The first is that some of the decision variables (i.e., the elements of the allocation matrix Z) are organized very naturally into a matrix. The second is that the blending constraint matrix (i.e., the matrix product QZV^{-1}) is nonlinear in the elements of the inventory vector v. In reformulating Eq. 13 into SDP form, a choice must be made as to whether the nonlinearity (in terms of the inventory vector v) is placed in the objective function or in the constraints. Experience with the ellipsoid algorithm used in this study showed that it was preferable to retain the nonlinearity in the objective function. It should be noted that the choice of where the nonlinearity should be placed is solver dependent.

Assuming that nonlinearity is retained in the objective function, the storage allocation problem can be reformulated to convert the allocation matrix Z into vector form as

$$\max_{\zeta, v} \sigma_n \left(\sum_{i=1}^{mp} \bar{Q}_i \frac{\zeta_i}{v_i} \right)$$

subject to:

$$[I_{m\times m}\ I_{m\times m}\ \cdots\ I_{m\times m}]\zeta = w$$

$$\begin{bmatrix} 1_m^T & 0_m^T & \cdots & 0_m^T \\ 0_m^T & 1_p^T & \cdots & 0_m^T \\ \vdots & \vdots & \ddots & 0_m^T \\ 0_m^T & 0_m^T & 0_m^T & 1_m^T \end{bmatrix} \zeta = [I_{p\times p}\ 0_{p\times(m-p)}]\boldsymbol{v} \quad \text{(A2)}$$

$$\{[I_{(m-1)p\times(m-1)p}\ 0_{(m-1)p\times p}] - [0_{(m-1)p\times p}\ I_{(m-1)p\times(m-1)p}]\}\boldsymbol{v} = 0$$

$$\zeta_{\text{lower}} \le \zeta \le \zeta_{\text{upper}}$$

$$\boldsymbol{v}_{\text{lower}} \le \boldsymbol{v} \le \boldsymbol{v}_{\text{upper}},$$

where

$$\zeta = \text{vec}(Z) \quad \text{(A3)}$$

$$\boldsymbol{v} = \begin{bmatrix} I_{p\times p} \\ \vdots \\ I_{p\times p} \end{bmatrix} v = \begin{bmatrix} v \\ \vdots \\ v \end{bmatrix} \quad \text{(A4)}$$

$$Q = \sum_{i=1}^{mp} \bar{Q}_i. \quad \text{(A5)}$$

In the preceding problem (Eq. A2) the allocation matrix Z is unfolded and the columns of the matrix "stacked" on top of each other to form the allocation vector ζ. As a result of the conversion of the allocation matrix to vector form, the inventory vector v must be modified by "stacking" m copies of it in the vector $\boldsymbol{v}$. These modifications result in an objective function that depends solely upon the new allocation vector ζ and inventory vector $\boldsymbol{v}$. (Note that the nonlinearity in terms of the tank inventories is retained in the objective function.) The blending indices of the feedstocks are now contained in a family of matrices $\bar{Q}_i$, which sum to the original blending index matrix Q. The first two equality constraints in Eq. A2 represent a reformulation of the feedstock availability and inventory constraints in Eq. 13. Notice that in Eq. A2 there is an additional $(m-1)\times p$ decision variable, due to the extra copies of the inventory vector v contained in $\boldsymbol{v}$. The third equality constraint is used to eliminate the extra degrees of freedom by insisting that each copy of the inventory vector v in $\boldsymbol{v}$ is identical.

Manuscript received Aug. 16, 1997, and revision received Apr. 24, 1998.