Computers and Chemical Engineering 164 (2022) 107873

Contents lists available at ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Multiperiod refinery optimization for mitigating the impact of process unit shutdowns

Ariel Boucheikhchoukh[a], Valentin Berger[a], Christopher L.E. Swartz[a,*], Antoine Deza[b,c], Alexander Nguyen[a], Shaffiq Jaffer[d]

[a] Department of Chemical Engineering, McMaster University, 1280 Main St W, Hamilton, ON, L8S 4L7, Canada
[b] Department of Computing and Software, McMaster University, 1280 Main St W, Hamilton, ON, L8S 4L7, Canada
[c] LIX, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France
[d] TOTAL American Services, Inc., 82 South St, Hopkinton, MA 01748, U.S.A.

ARTICLE INFO

Article history:
Received 30 December 2021
Revised 25 April 2022
Accepted 4 June 2022
Available online 8 June 2022

Keywords:
MILP
Multiperiod planning
Petroleum refinery
Inventory management
Shutdowns
Uncertainty
Two-stage stochastic programming

ABSTRACT

Process unit shutdowns, whether planned or unplanned, disrupt the normal operation of process plants. The negative impact of such disruptions can often be mitigated by the strategic use of material inventories, coupled with redistribution of material flows. In this paper, an optimization framework is proposed for a multiperiod refinery planning model in which a system of inventory tanks is used to mitigate the impact of process unit shutdowns. An economic objective is optimized over a planning horizon, subject to operating, design, and product quality constraints. Two formulation paradigms are considered - an operating problem in which optimal material flow trajectories are determined in response to a shutdown scenario, and a retrofit design problem in which the optimal location and sizing of inventory tanks are determined, the latter formulated as a two-stage stochastic programming problem. The utility of the formulation to shutdown planning is shown through case studies of the operational and design problems applied to a simple refinery scheme.

1. Introduction

The employment of model-based optimization in support of refinery planning has become essential in developing efficient operating policies in an increasingly competitive market. In crude-oil refineries, one or more types of crude oil are distilled, processed, and then pooled together to produce an array of blending products. Separation units split a petroleum feed into cuts of varying densities, whereas reactors catalytically alter the properties of a petroleum stream to improve its quality. Petroleum streams are mixed together in pools and blending tanks in order to achieve final products that meet the prescribed quality specifications. Market conditions and contractual obligations typically impose upper and lower bounds on the production rates of the blends. Refinery models can be used to address a variety of operational problems, ranging from crude-oil scheduling to optimal blending. Overviews of refinery planning models may be found, inter alia, in Pinto et al. (2000); Guerra and Le Roux (2011).

Over the lifetime of a refinery, a given process unit may be brought offline numerous times, either due to planned maintenance or an unplanned failure. Such shutdown events represent a disruption to the baseline operation of the refinery. If the process unit is critical to the overall operation of the refinery, a prolonged shutdown may induce a total shutdown. Total plant shutdowns significantly impact the profitability of the plant, potentially resulting in millions of dollars of lost revenue. In some cases, a refinery may continue operating while a unit is shut down. Under a partial plant shutdown, the process is shifted to an alternate operating policy for the duration of the shutdown, such that the impact on the plant profitability is mitigated. The shutdown policy may involve a number of compensatory actions: rerouting material streams, shifting operational burden to redundant process units, slowing down production, violating quality constraints, or mobilizing inventory. Preventing total plant shutdowns typically requires embedding fail-safes into the design of the refinery itself such that an appropriate response can be undertaken.

A common strategy involves an inventory management system, where inventory capacity is allocated to critical streams in the plant configuration. In its simplest form, an inventory management system comprises two process units connected serially via a material bypass stream and an intermediate inventory tank. The inventory tank may collect material from the upstream unit while it is online to build up a reserve. When the upstream unit is brought

* Corresponding author.
E-mail address: swartzc@mcmaster.ca (C.L.E. Swartz).

https://doi.org/10.1016/j.compchemeng.2022.107873

Nomenclature

Indices/Sets

$s = 1, \ldots, S$	Scenarios
$t = 1, \ldots, T$	Time periods
t_0	Beginning of the first time period
t_f	Last time period
$i \in I$	Material streams
$i \in I^{FG} \subset I$	Fuel gas streams
$i \in I^{BP} \subset I$	Final blending products (PG, RG, DF, and FO)
$u \in U^{Op}$	Unit operations
$u \in U^{SP}$	Split points
$u \in U^{BT}$	Blending tanks
$u \in U^{IT}$	Inventory (buffer) tanks
$i \in I_u$	Input streams for unit u
$j \in J_u$	Output streams from unit u
$q \in Q$	Quality specifications; includes the octane number, the vapor pressure (mmHg), the density (lb/bbl), and the sulfur content (lb/bbl)

Constants

$V_{\max}^{\text{design}}$	Maximum design volume of an inventory tank (bbl)
$N_{\max}^{IT}$	Maximum number of inventory tanks in the refinery scheme
C_i	Cost of the product i ($)
C_u^{Op}	Operational cost of the unit u ($)
C_u^{fixed}	Fixed capital cost of an inventory tank u ($)
C_u^{var}	Variable capital cost of an inventory tank u ($)
$F_u^{\max}$	Maximum capacity of the unit u (bbl/period)
D_i	Minimum demand for the product i (bbl/time period)
$\alpha_{u,t}$	Binary parameter that specifies whether unit u at time period t is shut down (1) or not (0)
μ_i^D	Penalty coefficient for the stream i when violating the demand constraint ($)
$X_{u,i,j}$	Yield coefficient for the incoming product i and the out-going product j for the unit u
$Z_{i,q}$	Quality q of stream i

Positive Variables

$F_{i,t}$	Flow rate of the stream i at the time period t (bbl)
$F_{u,t}^{\text{in}}$	Flow rate of the inlet stream for inventory tank u at the time period t (bbl)
$F_{u,t}^{\text{out}}$	Flow rate of the outlet stream for inventory tank u at the time period t (bbl)
$F_{u,t}^{\text{byp}}$	Flow rate of the bypass stream for inventory tank u at the time period t (bbl)
F_i^{offset}	Offset between the minimum demand and the actual flow rate of the product i (bbl)
$V_{u,t}$	Volume of the buffer tank u at the time period t (bbl)
V_u^{design}	Maximum capacity of the buffer tank u (bbl)
V_u^{start}	Starting and ending volume of the buffer tank u

Binary Variables

δ_u	Binary variable that is set to 1 if a buffer tank $u \in U$ exists, and set to 0 otherwise

offline, the inventory is mobilized to feed the downstream unit over the length of the shutdown. If the critical components of a plant can be sufficiently decoupled from one another such that they can operate independently for a sufficient amount of time, then a unit shutdown need not bring the entire plant offline. This example is trivial in that the ideal operating policy can be determined intuitively. For more complex plant configurations with highly integrated unit operations, the process dynamics may be such that the optimal operating policy is not easily determined through process insight alone. Additional considerations such as quality specifications, product demand, and raw material availability further constrain the elaboration of an appropriate policy. The size of the plant configuration may also significantly complicate the design process of an inventory system. As the number of candidate streams in the process increases, the number of possible tank configurations increases exponentially.

Determining the optimal inventory system design and corresponding operating policies requires balancing competing economic considerations: capital expenditure limits the number and size of tanks that can be installed, whereas a more robust inventory system can better mitigate losses in revenue due to unit shutdowns. A dynamic optimization problem can thus be constructed to address this question, where an economic objective function is maximized (or minimized) over the planning horizon, and an underlying process model simulates the impact of a unit shutdown on the state of the plant and its economics. Additional operational flexibility may be incorporated through the use of economic penalties. These penalties can capture the cost of violating operational constraints, such as a failure to meet client-specified production levels.

Over the lifetime of a plant, operators must contend with a variety of different shutdown events. In this work, we will consider two types of shutdowns. In the case of a pre-emptive shutdown (i.e. planned maintenance), preparatory action can be taken prior to the shutdown in order to transition the plant to the prescribed shutdown policy. Conversely, in the case of a reactive shutdown (i.e. an unplanned failure), the shutdown cannot be anticipated. A shutdown may also occur in different sections of the plant and may last for varying lengths of time. Ideally, the design of the plant should afford the necessary flexibility to the operators to adopt the best policy for a variety of shutdown events. Given that the parameters of a shutdown event are subject to uncertainty, there is a need for a stochastic problem formulation that can hedge the mitigating potential of a possible design for each shutdown scenario.

In two-stage stochastic optimization, the decision variables are partitioned into two sets: first-stage decisions that need to be made prior to the uncertainty realization, and second-stage (or recourse) decisions that are made after the uncertain parameters are revealed. The first-stage decisions consist of design variables, such as the locations of the inventory tanks, their maximum capacities, and their starting levels, whereas the second-stage decisions consist of operating variables, which may include raw material flow rates, routing material at split points, and mobilizing/accruing tank inventory. The stochastic problem is comprised of multiple subproblems, each one representing a different realization of the shutdown uncertainty. The subproblems are coupled to one another via the first-stage decisions, but each may be solved by a unique set of second-stage decisions (provided they are feasible). A thorough treatment of two-stage stochastic programming is given by Birge and Louveaux (2011). Applications to planning and design under uncertainty include Guillén et al. (2005), You et al. (2009), and Gerardi et al. (2013).

Prior works have already examined the use of model-based optimization in determining the optimal operating policy for inventory tanks in continuous processes. Dubé (2000) simulated partial plant shutdowns in a highly-integrated Kraft pulp and paper mill with nonlinear dynamics and a fixed configuration of inventory tanks. A methodology was outlined to determine the maximum duration of a unit shutdown that could be sustained without bringing the entire plant offline. The effect of preparation time on the planning strategy was also explored. Building upon this work, Balthazaar (2006) extended the problem formulation into a design problem. Binary variables were included in the model to capture the decision of whether to retrofit an existing inventory tank with additional storage capacity. The tank capacity expansions and starting levels were optimized over an array of shutdown scenarios to obtain a robust solution. The scenarios included pre-emptive and reactive shutdowns. The duration and location of the shutdown were also considered as sources of uncertainty. Chong and Swartz (2013, 2016) used a model predictive control (MPC) approach combined with a multi-tiered optimization strategy. Uncertainty in the duration of the shutdown was captured by treating it as a feedback parameter and re-estimating it at every time point. In addition to a primary economic objective, additional objective functions were sequentially optimized in order to obtain smoothed trajectories that exhibit minimal deviation from product quality targets.

Numerous refinery models have been proposed to solve a variety of planning and scheduling problems. Joly et al. (2002) developed a nonlinear multiperiod refinery model and embedded it within an economic optimization problem in order to determine the optimal blending and inventory management strategy for a fixed schedule of crude shipments. Due to the volatility of the oil market, Neiro and Pinto (2005) incorporated demand and price uncertainty into a stochastic refinery model with the objective of determining more robust operating policies. Neiro and Pinto (2006) subsequently employed temporal and scenario-wise Lagrangian decomposition to partition a refinery planning problem into smaller subproblems in order to decrease solution times. Similarly, Bengtsson et al. (2013) considered uncertainty in the arrival times of crude oil shipments. Developing optimal shutdown policies could translate to substantial savings for refineries. However, modelling unit shutdowns is a topic generally not considered in the refinery planning literature.

The aim of this work is to develop an optimization-based framework that can guide operational and design decisions in response to process unit shutdowns in crude-oil refineries. The basic framework consists of extending a multiperiod planning model of a given industrial process to include an inventory management system. A deterministic formulation is proposed for determining the optimal distribution of material flows in response to a planned or unplanned unit shutdown. This is followed by a retrofit design formulation in which buffer tank locations, capacities, and starting levels are included within the decision space. The design problem is formulated as a two-stage stochastic programming problem to account for the inherent uncertainty of partial plant shutdowns. The proposed operating and design decision-making paradigms are exemplified through application to case studies involving a simplified crude oil refinery model using linear yield relationships and blending rules.

While the proposed framework draws on well-established areas of refinery modelling and scenario-based stochastic programming, its use for refinery inventory management to mitigate the impact of process unit shutdowns in an operating or design setting is, to our knowledge, new. Other concepts utilized within the proposed formulations are the use of a two-tiered optimization approach to obtain flow and inventory profiles with reduced variation, and use of cumulative compositions to allow for blending strategies that avoid off-spec product. Extension to more complex refinery models would be a useful topic for future research, and is discussed briefly in Section 5.

2. Problem formulation

The refinery under consideration is based on the scheme presented in Pike (1986) and is illustrated in Fig. 1. The process units $u \in U^{Op}$ consist of an atmospheric distillation unit (AD), a catalytic reformer (RF), and a catalytic cracking unit (CC). Each unit possesses inlet streams $i \in I_u$ and outlet streams $j \in J_u$. The refinery includes, in addition, $u \in U^{BT}$ blending tanks, $u \in U^{SP}$ split points, and $u \in U^{IT}$ inventory tanks, each with similarly defined sets of inlet and outlet streams. Additionally, a set I is defined as the set of all material streams. There are a total of four blending products ($i \in I^{\text{BP}}$): premium gasoline (PG), regular gasoline (RG), diesel fuel (DF), and fuel oil (FO). Three fuel gas streams are also produced as byproducts ($i \in I^{\text{FG}}$): FGAD, FGRF and FGCC. The distillation unit is fed by a single crude oil stream ($i \in I^{\text{Cr}}$). The blending products are subject to quality specifications ($q \in Q$): octane number (OCT), vapor pressure (VP), density (DEN), and sulfur content (SULF). The yield relationships and the blending rules are linear, and the properties of the intermediate streams are fixed.

We consider two related paradigms: the operational problem and the design problem. The objective of the former is to determine the optimal flow and inventory profiles over a multiperiod horizon for a given set of storage tank locations and capacities. The design problem, on the other hand, seeks to determine the optimal inventory locations and capacities that mitigate the impact of process unit shutdowns on the operation of the refinery. The problem formulation of the deterministic case is described first, and in Section 2.7 it is extended to a two-stage stochastic formulation.

2.1. Economic objective function

The primary objective is to maximize profit over a multiperiod planning horizon. Sources of revenue include the exportation of intermediate and final products, whereas costs include the cost of crude oil, demand shortfall penalties, process unit operating costs, and the amortized capital cost of the inventory tanks.

$$\max \Phi = \sum_{t=1}^{T}\left[\sum_{i \in I^{BP} \cup I^{FG}} C_i F_{i,t} - C_{cr}F_{cr,t} - \sum_{u \in U^{Op}} \sum_{i \in I_u} C_u^{Op} F_{i,t}\right] - \sum_{i \in I^{BP}} \mu_i^D F_i^{\text{offset}} - T \sum_{u \in U^{IT}} \left(C_u^{\text{fixed}} \delta_u + C_u^{\text{var}} V_u^{\text{design}}\right) \quad (1)$$

In the above equation, Φ is the total profit over the planning horizon. Each term represents either a cost or a source of revenue. The summation over the planning horizon $\{t = 1, \ldots, T\}$ represents the operating profit of the refinery. The per-unit price for each stream i is denoted by C_i, whereas the operating cost of process unit u is denoted by C_u^{Op}. The second term of the objective function applies an economic penalty for production levels that fall below a cumulative minimum demand for the entire planning horizon. F_i^{offset} is the demand shortfall of product i and μ_i^D is the per-unit penalty cost associated with the demand shortfall of product i. The final term of the objective function represents the amortized capital cost of the inventory tank system. The fixed and variables costs of installing an inventory tank are denoted by C_u^{fixed} and C_u^{var} respectively. The cost of a tank scales linearly with its maximum design capacity, and is amortized over a ten-year project lifetime to the length of the planning horizon. Binary variable δ_u indicates the existence of an inventory tank u, while the size of a tank is given by V_u^{design}.

2.2. Flow and process unit constraints

The refinery scheme is modelled via a series of mass balances and capacity constraints indexed with respect to time. Eq. 2 specifies the maximum throughput (denoted by $F_u^{\max}$) of a process unit $u \in U^{Op}$, where I_u is the set of incoming streams for unit u and $\alpha_{u,t}$ is a binary parameter that indicates whether the unit is active (0) or shut down (1). Eq. 3 relates the outlet flow rate(s) of a process unit $u \in U^{Op}$ via a vector of yield coefficients $X_{u,i,j}$ between an incoming feed stream i and an outgoing stream j. Finally, Eq. 4 and 5 balance the inlets and the outlets of each split point $u \in U^{SP}$ and

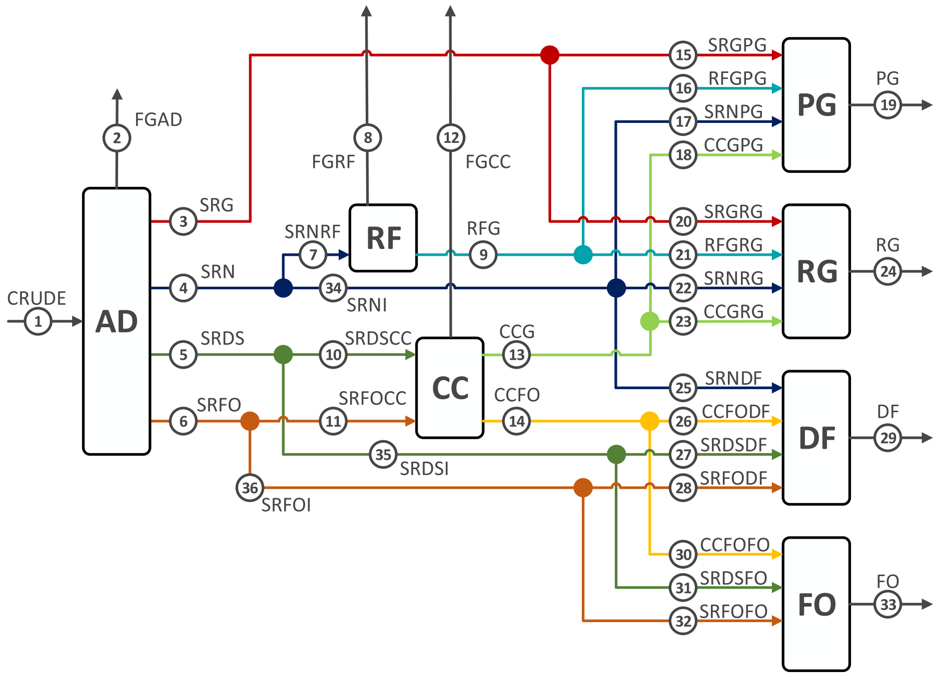

Fig. 1. Refinery scheme.

each blending tank $u \in U^{BT}$ respectively.

$$\sum_{i \in I_u} F_{i,t} \le (1 - \alpha_{u,t}) F_u^{max}, \quad \forall u \in U^{Op}, t = 1, \ldots, T \tag{2}$$

$$(1 - \alpha_{u,t}) \sum_{i \in I_u} X_{u,i,j} F_{i,t} = F_{j,t}, \quad \forall u \in U^{Op}, \forall j \in J_u, t = 1, \ldots, T \tag{3}$$

$$\sum_{i \in I_u} F_{i,t} = \sum_{j \in J_u} F_{j,t}, \quad \forall u \in U^{SP}, t = 1, \ldots, T \tag{4}$$

$$\sum_{i \in I_u} F_{i,t} = \sum_{j \in J_u} F_{j,t}, \quad \forall u \in U^{BT}, t = 1, \ldots, T \tag{5}$$

2.3. Quality constraints

The blending products must each satisfy a set of quality specifications so that they are suitable for sale on the market. During a process unit shutdown, meeting these quality specifications at each time period may become infeasible. This difficulty can be overcome by appropriately blending low-quality material with high-quality material over the course of the planning horizon such that the final mixture meets the specifications. In the proposed formulation, cumulative quality constraints are employed in order to consider blending strategies for off-spec products. Rather than imposing a quality constraint at each time period, a cumulative quality constraint is applied to the average quality over the planning horizon. A linear blending rule is assumed for each quality $q \in Q$. Note that both minimum and maximum quality specifications are in effect. For the present case study, a minimum octane number is specified, whereas maximum limits apply for the vapor pressure, the density, and the sulfur content. The mathematical formulation of the above-described quality constraint scheme is described below.

Minimum quality specification

The minimum cumulative quality specification is expressed as

$$\sum_{t=1}^{T} \sum_{i \in I_u} Z_{i,q} F_{i,t} \ge Z_{j,q}^{spec} \sum_{t=1}^{T} F_{j,t}$$

$$\forall u \in U^{BT}, \forall j \in J_u, q := OCT \tag{6}$$

where I_u and J_u are the inlet and outlet streams of the blending tank u respectively, $Z_{i,q}$ is the quality q of stream i, $Z_{j,q}^{spec}$ is the minimum quality specification for the blended stream, and OCT denotes the octane number specification.

Maximum quality specification

Analogous to Eq. 6, the maximum quality specifications are expressed as

$$\sum_{t=1}^{T} \sum_{i \in I_u} Z_{i,q} F_{i,t} \le Z_{j,q}^{spec} \sum_{t=1}^{T} F_{j,t}$$

$$\forall u \in U^{BT}, \forall j \in J_u, q \in \{VP, DEN, SULF\} \tag{7}$$

where VP, DEN and $SULF$ denote the vapor pressure, density, and sulfur specifications respectively.

2.4. Demand penalty

An economic penalty is applied when production levels fall short of minimum requirements. The proposed formulation assumes that the final products of a refinery are sold in batches at regular intervals. To simulate this, a cumulative demand is considered over the entirety of the planning horizon rather than a daily demand for each time period.

$$\sum_{t=1}^{T} F_{i,t} + F_i^{\text{offset}} \ge D_i, \quad \forall i \in I^{BP} \tag{8}$$

where F_i^{offset} and D_i are the cumulative demand shortfall and the minimum demand for the product i respectively. A penalty term proportional to the demand shortfall is included in the objective function.

2.5. Inventory tanks

A selection of streams $i \in I^{IT} \subset I$ in the refinery scheme are authorized to connect to an inventory tank within a set of potential intermediate storage units, U^{IT}. These streams are decomposed

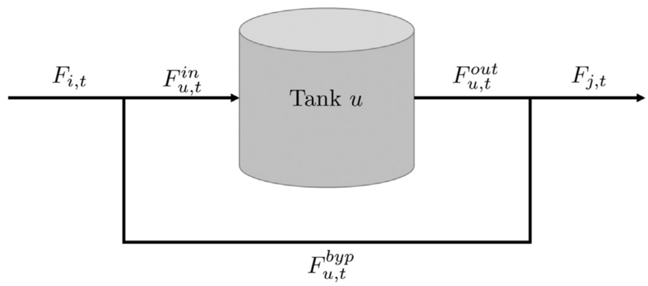

Fig. 2. Model of buffer tank u, with $i \in I_u$, $j \in J_u$.

into five segments, as shown in Fig. 2. Each inventory tank has one inlet, one outlet, and a bypass stream. Eq. 9 and Eq. 10 link the streams in I^{IT} to their corresponding inventory tanks.

$$F_{i,t} = F_{u,t}^{\text{in}} + F_{u,t}^{\text{byp}}, \quad \forall i \in I_u, \forall u \in U^{IT}, t = 1, \dots, T \tag{9}$$

$$F_{j,t} = F_{u,t}^{\text{out}} + F_{u,t}^{\text{byp}}, \quad \forall j \in J_u, \forall u \in U^{IT}, t = 1, \dots, T \tag{10}$$

Additional constraints specify how the inventory tanks are to be operated within the refinery scheme. Eq. 11 balances the incoming stream, the outgoing stream, and the accumulated inventory ($V_{u,t}$) of a tank u. Eq. 12 ensures that the initial and final inventories of a given tank u are equal. Eq. 13 and 14 specify that the inventory of a tank u must be greater than 20% and less than 80% of its design capacity V_u^{design}. Eq. 15 bounds the design capacity of an inventory tank u if it exists, and sets it to 0 otherwise. Finally, Eq. 16 sets an upper limit $N_{\max}^{IT}$ on the number of inventory tanks in the refinery.

$$V_{u,t} - V_{u,t-1} = F_{u,t}^{in} - F_{u,t}^{out}, \quad \forall u \in U^{IT}, t = 1, \dots, T \tag{11}$$

$$V_{u,t_0} = V_{u,t_f} = V_u^{\text{start}}, \quad \forall u \in U^{IT}, t = 1, \dots, T \tag{12}$$

$$0.2V_u^{\text{design}} \le V_{u,t} \le 0.8V_u^{\text{design}}, \quad \forall u \in U^{IT}, t = 1, \dots, T \tag{13}$$

$$0.2V_u^{\text{design}} \le V_u^{\text{start}} \le 0.8V_u^{\text{design}}, \quad \forall u \in U^{IT} \tag{14}$$

$$V_u^{\text{design}} \le V_{\max}^{\text{design}} \delta_u, \quad \forall u \in U^{IT} \tag{15}$$

$$\sum_{u \in U^{IT}} \delta_u \le N_{\max}^{IT} \tag{16}$$

2.6. Process smoothing

Some problems may have multiple solutions that yield equivalent profit values. On the surface, these solutions appear to be equally good, however some solutions may be more practical to implement than others. In general, solutions that deviate little from the base case are more desirable from an operational perspective. In the proposed solution procedure, optimal solutions are determined by solving a second optimization problem with a quadratic objective function that captures the aggregate fluctuation in the production/inventory profiles of the refinery scheme. The sum of the differences between the production levels and inventory tank levels at times t and $t-1$ is minimized over the planning horizon (Eq. 17), resulting in a smoother process. The economic objective function (Eq. 1) is converted into an inequality constraint to ensure that the refinery profit is within a specified tolerance of the maximum profit. The fluctuation minimization problem returns either the solution of the original problem, or returns an alternative solution with a profit within a prescribed error tolerance.

Table 1
Octane rating and vapor pressure specifications.
Source: Pike (1986).

Stream	Octane	Vapor pressure (mmHg)
SRGPG	78.5	18.4
RFGPG	104	2.57
SRNPG	65	6.54
CCGPG	93.7	6.9
PG	≥ 93	≤ 12.7
SRGRG	78.5	18.4
RFGRG	104	2.57
SRNRG	65	6.54
CCGRG	93.7	6.9
RG	≥ 87	≤ 12.7

$$\begin{aligned} \min \quad & \sum_{t=1}^{T-1} \left(\sum_{i \in I^{BP}} (F_{i,t} - F_{i,t+1})^2 + \sum_{u \in U^{IT}} (V_{u,t} - V_{u,t+1})^2 \right) \\ \text{s.t.} \quad & \text{Eq. 1 - 15} \\ & \Phi \ge (1-\epsilon)\Phi_{\max} \end{aligned} \tag{17}$$

where $\Phi_{\max}$ is the optimal value of the first-tier economic maximization problem, and ϵ is the error tolerance for the total profit. The first-tier problem takes the form of a mixed-integer linear program (MILP), and the second tier is a mixed-integer quadratic program (MIQP). A similar two-tiered approach is applied in Cao et al. (2016) in the dynamic economic optimization of an air separation process.

2.7. Optimization under uncertainty

A two-stage stochastic programming formulation is applied to the design problem, where the objective is to determine the optimal locations and capacities of storage tanks to maximize the expected profit subject to uncertainty in process unit shutdowns. In the present application, the first-stage decisions comprise the tank locations and capacities, and the second-stage decisions are the stream flows and tank inventories at each period. The optimization problem can be expressed generically as

$$\begin{aligned} \max \quad & c^T x + \sum_{s=1}^{S} p_s g_s^T y_s \\ \text{s.t.} \quad & Ax \ge b \\ & T_s x + W_s y \ge r_s, \quad \forall s \in S \end{aligned} \tag{18}$$

where x are the first-stage or design decisions; y are the second-stage or operating decisions; p_s is the probability of occurrence of scenario s; S is the number of scenarios; c, A and b are the parameters of the design problem; and g_s, T_s, W_s, and r_s are the parameters subject to uncertainty.

3. Model parameters

The refinery case study outlined in this section is based on the configuration presented in Pike (1986). The unit yields, the unit capacities, the quality specifications, and the material prices are outlined in Table A1, Table A2, Table 1, Table 2, and Table A3 respectively. All quantities are reported in US oil refining units; conversion factors for SI units are listed in Table B1. The prices of crude oil and of the blending products were updated by referring to the refiner prices in the U.S. Energy Information Administration's monthly energy report (U.S. Energy Information Administration (2019)). This report does not include a refiner price for premium gasoline. The price of PG was instead estimated by applying

Table 2
Density and sulfur specifications. Source: Pike (1986).

Stream	Density (lb/bbl)	Sulfur (lb/bbl)
SRNDF	272	0.283
CCFODF	294.4	0.353
SRDSDF	292	0.526
SRFODF	295	0.98
DF	$\leq$ 306	$\leq$ 0.5
CCFOFO	294.4	0.353
SRDSFO	292	0.526
SRFOFO	295	0.98
FO	$\leq$ 352	$\leq$ 3

the PG:RG ratio in retail prices to the refiner price of regular gasoline.

The capital cost of an inventory tank is expressed as a linear function of tank capacity in barrels. The capital cost of a characteristic case size was first determined via the empirical correlations presented in Turton et al. (2012). The material of construction was taken to be carbon steel, and a length:diameter ratio of 3 was selected. The characteristic cost was amortized to a single day in a project lifetime of 10 years. The capital cost of 160 different tank capacities between 500 and 100 000 bbl were estimated by extrapolating the cost of the characteristic case via a heuristic.

$$CAPCOST_2 = CAPCOST_1 \left(\frac{V_2}{V_1} \right)^{0.6} \tag{19}$$

Eq. 19 accounts for economics of scale by decreasing the capital cost per unit volume. The resulting cost function is linearized by determining the line of best fit, where the y-intercept is the fixed tank cost, and the slope is the variable tank cost.

$$CAPCOST = 0.0295V + 666.86 \tag{20}$$

The minimum demands for each of the final products were determined by first solving a deterministic base case problem involving no shutdowns, no inventory tanks, no demand limits, and an upper bound of 100,000 bbl/day for the crude oil feed. The optimal product rates obtained from this base case were halved and then assigned as the minimum production levels for all subsequent case studies. The minimum demands are listed in Table A4.

In accordance with Eqs. 1 and 8, the penalty for not meeting demand increases linearly with the demand shortfall. This is illustrated in Fig. 3 for each of the blending products.

4. Case studies

4.1. Base case

In the base case, all units are active and no inventory tanks are included in the design of the refinery. The solution of the base case represents the optimal operating point of the refinery in which the process remains at a baseline state at each period of the planning horizon. The time periods are decoupled from one another by removing the allowance for inventory and by introducing quality constraints at each period. Consequently, each period becomes an independent subproblem with an identical solution. Intuitively, the optimal operating point will seek to maximize production of PG, RG, and DF, which are all valued above the price of crude oil, while only producing FO from the less valuable intermediate streams in order to recoup costs.

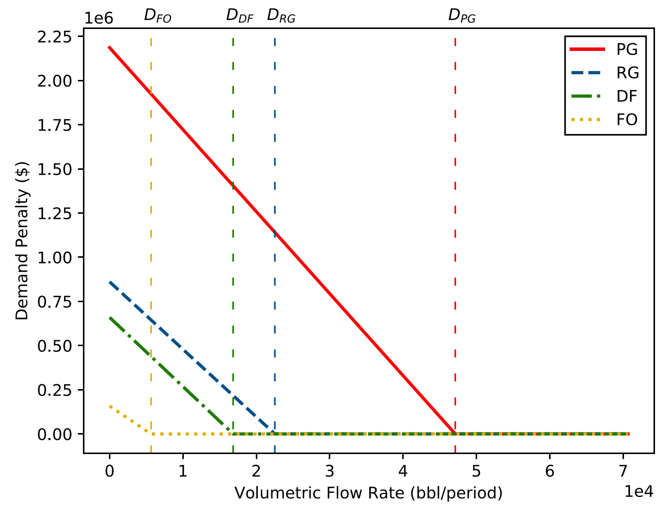

Fig. 3. Demand penalty as a function of product flow rate for a 1-period planning horizon. The color-coded vertical lines denote the minimum demand of each blending product.

An optimal profit value of $18 499 647 was obtained for the entire 10-period planning horizon. As expected, PG was produced in the greatest proportion (47 113 bbl/period), whereas the FO production was comparatively smaller (5653 bbl/period). The throughput of the AD is maximized, with 100 000 bbl of crude being processed each period. The effective throughputs of the RF and the CC are also maximized: the entirety of the straight-run naphtha (SRN) stream is sent to the RF, whereas 30 000 bbl/period of straight-run fuel oil (SRFO) is sent to the CC. The active quality constraints are the octane ratings of the gasoline blends (93 and 87 for PG and RG respectively) and the sulfur content of diesel fuel (0.5 lb/bbl). Due to its higher octane rating of 104, the entirety of the reformed gasoline (RFG) stream produced by the RF is routed to PG blending since it can better satisfy the octane specification of premium gasoline. On the other hand, the cracked gasoline (CCG) stream, with its more modest octane rating of 93.7, is mostly routed to RG blending. Given that the PG and RG blending tanks share the same inlet streams, the optimal solution will match low-value streams (i.e. those with low octane ratings) with high-value streams (i.e. those with high octane ratings) in the appropriate ratio such that the minimum octane ratings of PG and RG are obtained. A similar routing strategy was adopted for DF and FO blending. The low sulfur content of diesel fuel is difficult to satisfy, since only the SRN stream and the cracked fuel oil (CCFO) stream possess sulfur

Table 3
Effect of the inventory tank configuration on the total profit of the operational problem over a 10-period horizon.

Shutdown		Tank Configuration			Profit ($)
Location	Length	Locations	Size (bbl)	Starting Inventory (%)	
None	0	None			18 499 647
CC	2	None			16 425 160
CC	2	CCG	150 000	50	16 536 875
CC	2	CCFO	150 000	50	16 485 225
CC	2	CCG & CCFO	75 000	50	16 589 520

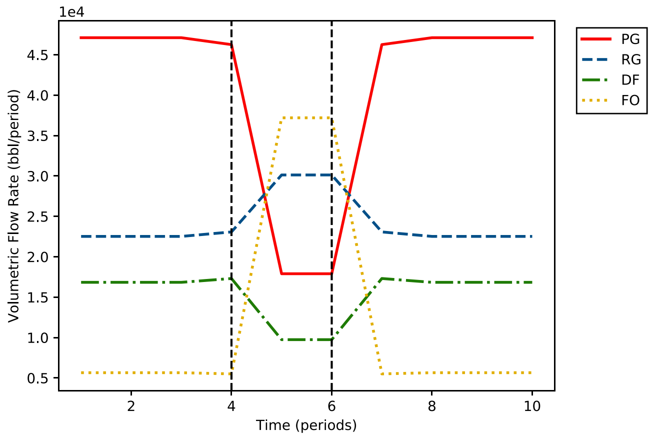

Fig. 4. Optimal production profile for a 2-period pre-emptive shutdown of the catalytic cracker over a 10-period planning horizon with no inventory tanks. The starting and ending points of the shutdown are denoted by dashed lines.

contents below 0.5 lb/bbl. Given that the SRN stream is used entirely as a feed for the reformer, the CCFO stream must be routed in its totality to DF blending in order to maximize the production of diesel fuel. The remaining material streams are less valuable, and are thus used to blend FO as a means of partially recouping the purchasing cost of crude oil.

4.2. Operational problem

When a process unit is shut down, the refinery must deviate from the optimal operating point of the base case. In the operational problem, the objective is to determine the optimal material flow distribution under different partial shutdown scenarios for a prescribed storage tank configuration and design capacity. The aim of the following case studies is to assess how tank inventories may be used to mitigate the impact of partial shutdowns on the profitability of the refinery. Each case study in this section considers a pre-determined buffer tank configuration. To reflect this, the cost of installing a buffer tank was omitted from the objective function (Eq. 1). For simplicity, quality specifications are imposed at each time period.

4.2.1. Pre-Emptive shutdown of the catalytic cracker

The shutdown scenarios presented in this section simulate pre-emptive shutdowns: when a unit is shut down for planned maintenance, the tank inventories may be altered prior to, during, and following the shutdown. The first shutdown scenario that was considered was a 2-period shutdown of the catalytic cracker during periods 5 and 6. Without any buffer tanks, a total profit of \$16 425 160 was obtained, which is significantly less than the total profit of the base case. In general, the response of the model to a shutdown scenario may be divided into three phases: a preparation phase, a shutdown phase, and a restoration phase (Fig. 4). Volumetric flow rates are defined for each time period; connecting lines are drawn so that the trends can be easily followed. This presentation paradigm is used throughout the case studies. Without any inventory tanks, the flow rates of the blending products remain at the baseline during the preparation and restoration phases. During the shutdown of the catalytic cracker, production of CCG and CCFO ceases entirely, resulting in dramatic declines in PG and DF production. The RFG stream is partially rerouted to the RG blending tank in order to meet its octane number specification, resulting in a reduced production of PG during the shutdown. This action is necessary to avoid violating the quality specifications of RG. Similarly, a fraction of the SRN stream is routed to the DF blending tank in order to satisfy its stringent sulfur content specification. Since SRN is the precursor of RFG, there is a trade-off between production of the gasoline blends and diesel fuel.

In the following shutdown scenario, a single buffer tank was placed at the CCG stream with a capacity of 150 000 bbl and an initial level of 75 000 bbl. The catalytic cracker was shut down during periods $t = 5$ and $t = 6$, resulting in a total profit of \$16 536 875 over the 10-period planning horizon. This amount only represents the valuation of crude oil and the blending products, specifically omitting installation cost of the CCG tank. The operational problem assumes that the CCG tank is an existing unit in the refinery scheme and that no retrofit of the plant is necessary. The addition of a CCG tank represents over \$100 000 in savings compared to the case with no tanks. PG is produced at a reduced rate during the preparation phase in order to gradually fill the CCG tank (Fig. 5). Once the shutdown begins, the CCG inventory is mobilized in order to maintain RG production near the base value. During the restoration phase, PG production is gradually returned to its nominal level while the CCG tank is refilled to its initial level.

The main benefit of the CCG tank is that it reduces the diversion of valuable material away from PG blending and towards RG blending. In the case without an inventory tank, this diversion is necessary in order to prevent the violation of the RG quality specifications. The inclusion of CCG inventory allows the refinery to increase its PG production during the shutdown, translating into an increase in total profit. However, the CCG tank has no effect on DF production, and thus cannot avoid the partial redirection of the SRN stream (i.e. the inlet of the reformer) to DF blending.

4.2.2. Effect of tank location

The following cases illustrate how the proposed framework may be used to identify critical streams where inventory tanks would be most beneficial. For example, the inclusion of CCFO inventory may prevent the redirection of SRN during catalytic cracker shutdowns, thereby resulting in greater production of the PG blend. To test this hypothesis, a single 150 000 bbl tank was placed at the CCFO stream, whereas the CCG stream was left without a tank. A maximum profit of \$16 485 225 was obtained over a 10-period planning horizon. This result represents a decline of over \$50 000 in profit compared to the CCG tank scenario. As expected, the SRN stream was fully routed to the reformer. However, the lack of CCG inventory forces the refinery to reroute the RFG stream to RG blending to prevent any quality penalties from being incurred, resulting in an overall decrease in PG production (Fig. 6). Consequently, the substitution of a CCG tank for a CCFO tank results in a drop in total profit.

A logical compromise would be to incorporate both CCG and CCFO inventory in the shutdown mitigation strategy. Two half-sized tanks with maximum capacities of 75 000 bbl were placed at the CCG and CCFO streams; the initial and final levels of both tanks were fixed at 37 500 bbl. The total material inventory is equivalent to that of the previous cases, but it is split into separate CCG and CCFO pools. The smoothed operating point has a profit of \$16 589 520, which represents an increase of over \$50 000 compared to the CCG tank scenario. The CCG inventory is mobilized during the shutdown in order to increase PG production, whereas the CCFO inventory is mobilized to meet the sulfur specification of diesel fuel (Fig. 7). As with the previous case, the SRN stream is used entirely as a feed for RFG production, thereby further increasing the production of PG. By splitting storage capacity between both outlets of the catalytic cracker, the impact of the shutdown on production is further mitigated.

The manner in which inventory is distributed throughout the process has an important effect on the optimal response to a unit shutdown. For a given shutdown scenario, certain streams will take on added importance in mitigating the plant's losses during the

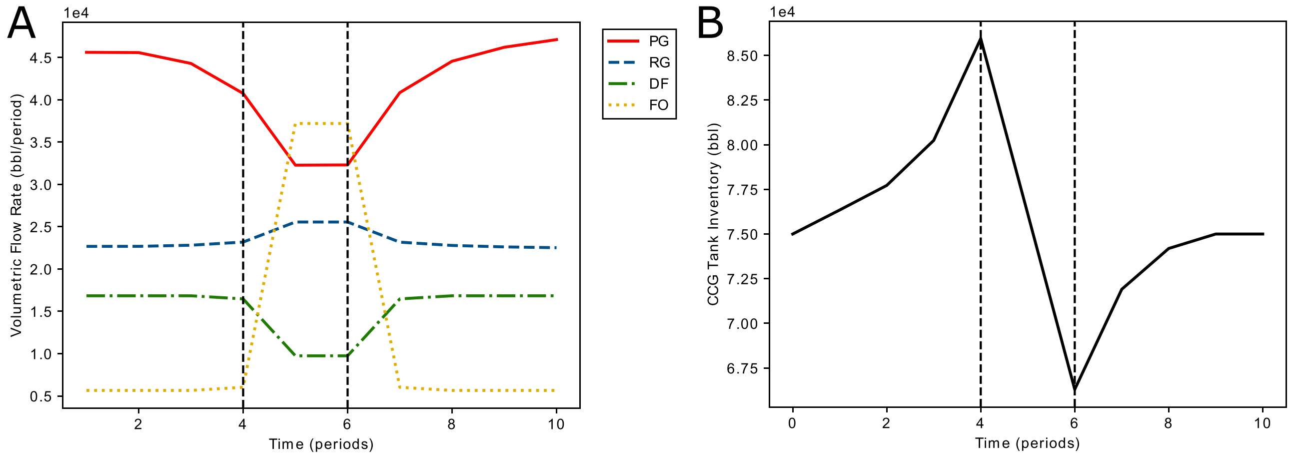

Fig. 5. Optimal operating policy for a 2-period pre-emptive shutdown of the catalytic cracker over a 10-period planning horizon with a fixed CCG tank. The starting and ending points of the shutdown are denoted by dashed lines. A) Flow profiles of the blending products. B) Inventory profile of the CCG tank.

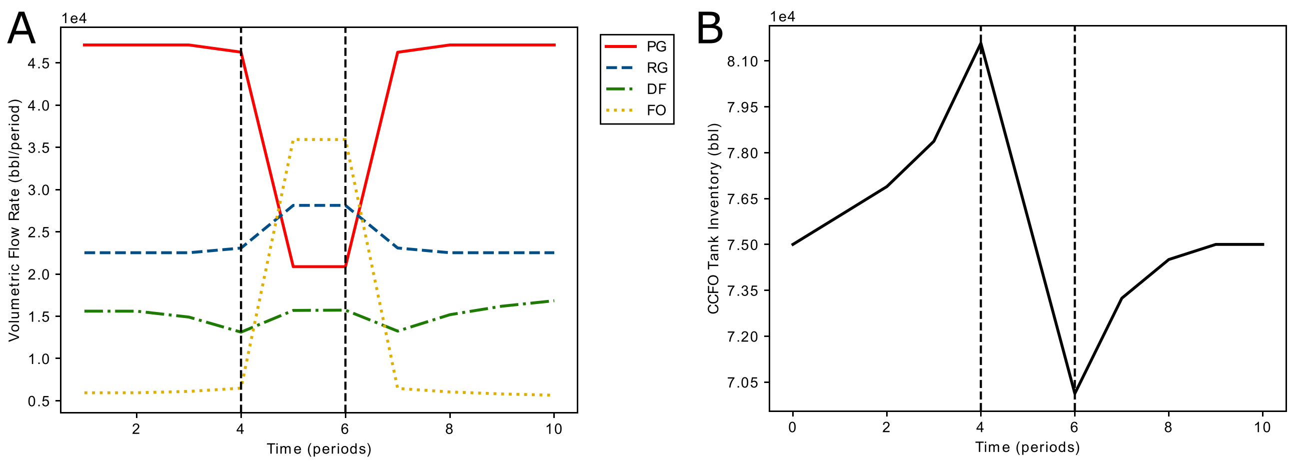

Fig. 6. Smoothed operating policy for a 2-period pre-emptive shutdown of the catalytic cracker over a 10-period planning horizon with a fixed CCFO tank. The starting and ending points of the shutdown are denoted by dashed lines. A) Flow profiles of the blending products. B) Inventory profile of the CCFO tank.

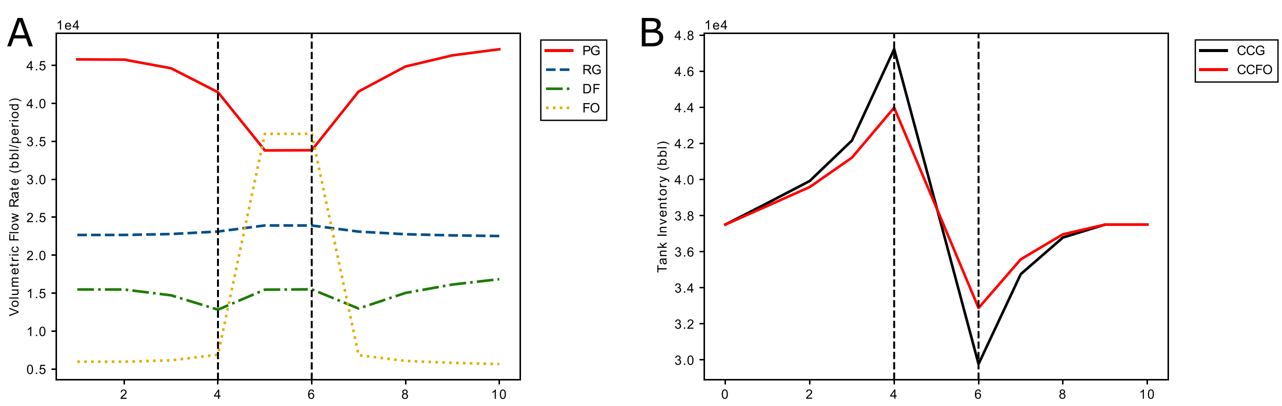

Fig. 7. Smoothed operating policy for a 2-period pre-emptive shutdown of the catalytic cracker over a 10-period planning horizon with fixed CCG and CCFO half-tanks. The starting and ending points of the shutdown are denoted by dashed lines. A) Flow profiles of the blending products. B) Inventory profiles of the CCG and CCFO tanks.

shutdown. Identifying these critical streams and then deciding how to best distribute inventory capacity among them is a non-trivial problem, even for simple processes such as the one under consideration.

4.2.3. Effect of preparation time

The preceding scenarios considered pre-emptive failures where the plant was allowed to deviate from its baseline state prior to the shutdown in order to prepare for the shutdown response. This window of time introduces additional degrees of freedom for the mitigation of the impact of the shutdown. For example, an inventory tank downstream of the unit may stock up on material from an outlet stream in anticipation of a shutdown. A tank upstream of the unit may also deplete its inventory prior to the shutdown in order to make room for additional storage of inlet material while

Table 4
Effect of preparation length on the total profit over a 15-period planning horizon with a 9-period shutdown of the catalytic cracker.

Preparation Time (periods)	Profit ($)	Solution Time (s)	PG:RG
3	18 972 498	0.031	1.257
2	18 960 976	0.047	1.252
1	18 796 600	0.031	1.194
0	18 633 190	0.031	1.139

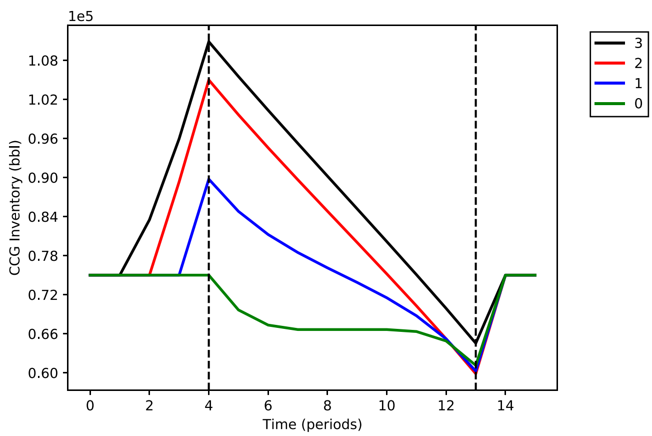

Fig. 8. Effect of the preparation time on the CCG inventory profile over a 15-period planning horizon and with a 9-period shutdown of the catalytic cracker. Each profile corresponds to a different number of periods allotted for preparation, ranging from 0 to 3. The starting and ending points of the shutdown are denoted by dashed lines.

the unit is offline. Shortening the preparation window thus restricts the set of feasible solutions, which may result in a lower objective value. In the extreme case where no preparation time is allowed, the shutdown becomes reactive: the plant is unable to anticipate the unit failure, and thus may only deviate from nominal operation during and after the shutdown. Depending on the dynamics of the process, the preparation length may have no effect on the total profit of the plant. If non-unique solutions exist where the impact of the shutdown can be mitigated once the shutdown begins while maintaining the same total profit, then the preparation time becomes a non-factor. In the case of a 2-period shutdown of the catalytic cracker with a 4-period restoration phase, the same total profit is obtained whether the shutdown is pre-emptive or reactive.

In order to observe the effect of preparation time on profit in the system under consideration, the shutdown was lengthened to 9 periods and the planning horizon was increased to 15. The length of the preparation phase was set by fixing the periods at the beginning of the horizon to the baseline response while maintaining the shutdown phase during periods 5–13. In the case of a 3-period preparation phase, only the very first period was fixed, whereas in the case of a reactive shutdown (i.e. no preparation time), periods 1–4 were fixed. Decreasing the preparation time resulted in progressively lower total profits for a longer shutdown (Table 4). In all four scenarios, the preparation phase was devoted to filling the CCG tank with inventory while producing PG at a decreased rate. A longer preparation phase allowed the refinery to store more material before the shutdown (Fig. 8). The differences in profit arise from the ratio between PG and RG production. Although a total gasoline production of 850 038 bbl is maintained across all four scenarios, the PG:RG ratio (and by extension the total profit) increases with preparation time. A smaller store of CCG inventory results in a greater diversion of the RFG stream away from PG blending and towards RG blending. In other words, the increased inventory allows for a more optimal routing of material to the blending tanks.

4.3. Design problem

The design problem considers the additional decision of how to retrofit the plant with buffer tanks. The buffer tank locations (δ_u) and their design capacities (V_u^{design}) are unfixed, thereby increasing the degrees of freedom of the operational problem. The decision to include an additional tank in the refinery scheme is mediated by economic considerations: increased inventory offers more profitable operational strategies for offsetting the cost of reduced unit throughput, whereas the cost of installation, which scales linearly with design capacity, may at one point supersede any operational benefits. The formulation of the design problem has two principal use cases: the deterministic case and the stochastic case.

4.3.1. Deterministic case

In the following case studies, a shutdown scenario is characterized by four parameters: the location of the shutdown, the start of the shutdown, the length of the shutdown, and the available preparation time. The deterministic case considers a single scenario where each of these parameters is known, with the goal of obtaining the inventory tank configuration that optimally mitigates the impact of the shutdown in the scenario of interest. In practice, a plant may only be retrofitted with a given number of tanks due to limited resources. Insight of the process might also be used to identify the critical streams of the process where inventory would have the greatest mitigating effect. In this work, the following restrictions were applied: the refinery scheme may only be retrofitted with up to two tanks, and these tanks may only be placed at the outlet streams of the distillation unit, the reformer, and the catalytic cracking unit. A maximum tank capacity of 150 000 bbl was chosen. Furthermore, the period-wise quality specifications were relaxed into cumulative specifications in order to expand the possible mitigation strategies.

To illustrate the deterministic case, the optimal buffer tank configuration was determined for a 6-period pre-emptive shutdown of the catalytic cracker over a 15-period planning horizon. Two periods were allocated for preparation and 7 for restoration. The optimal configuration consists of a single SRFO tank, yielding a total profit of $22 046 839 (amortized capital cost included). Prior to the shutdown, the SRFO tank was depleted in preparation for the CC shutdown. While the CC was offline, the SRFO stream (which is normally used as the sole feedstock of the CC unit) was sent either to FO blending (Fig. 9A) or stored in the tank (Fig. 9B). The increased FO production partially compensates for the reduced levels of gasoline and diesel production while the CC is shut down. Once the unit is brought back online, the production levels roughly return to their base case positions. FO production drops substantially to allow for the restocking of SRFO. Compared to the operational cases considered previously, the SRFO tank has a significantly smaller capacity (Table 5). While this decision reduces the necessary capital expenditure, the allowable variation in the inventory level is limited.

4.3.2. Stochastic case

The key limitation of the deterministic formulation is that it optimizes the buffer tank configuration for one set of shutdown parameters. As such, the solution may not be optimal for a different shutdown scenario, and may have no mitigating effect whatsoever. The problem formulation must account for the uncertainty in the shutdown parameters in order to obtain a solution where mitigating action can be taken for more than one scenario. A two-stage stochastic program was formulated to consider an array of shutdown scenarios, each with its own unique set of parameters, where

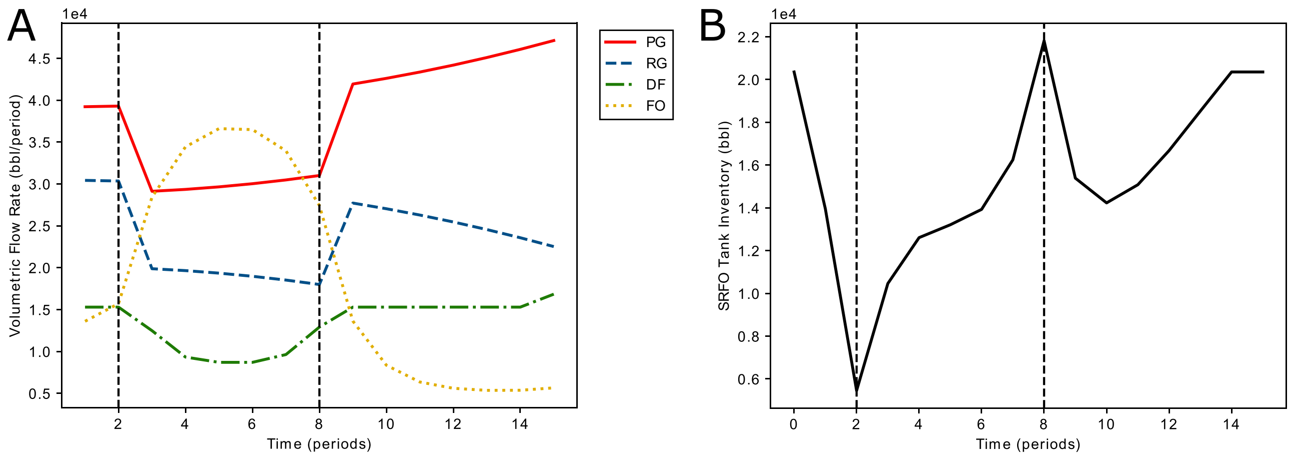

Fig. 9. Flow and inventory profiles of the optimal buffer tank configuration for a 6-period pre-emptive shutdown of the catalytic cracker over a 15-period planning horizon. The starting and ending points of the shutdown are denoted by dashed lines.

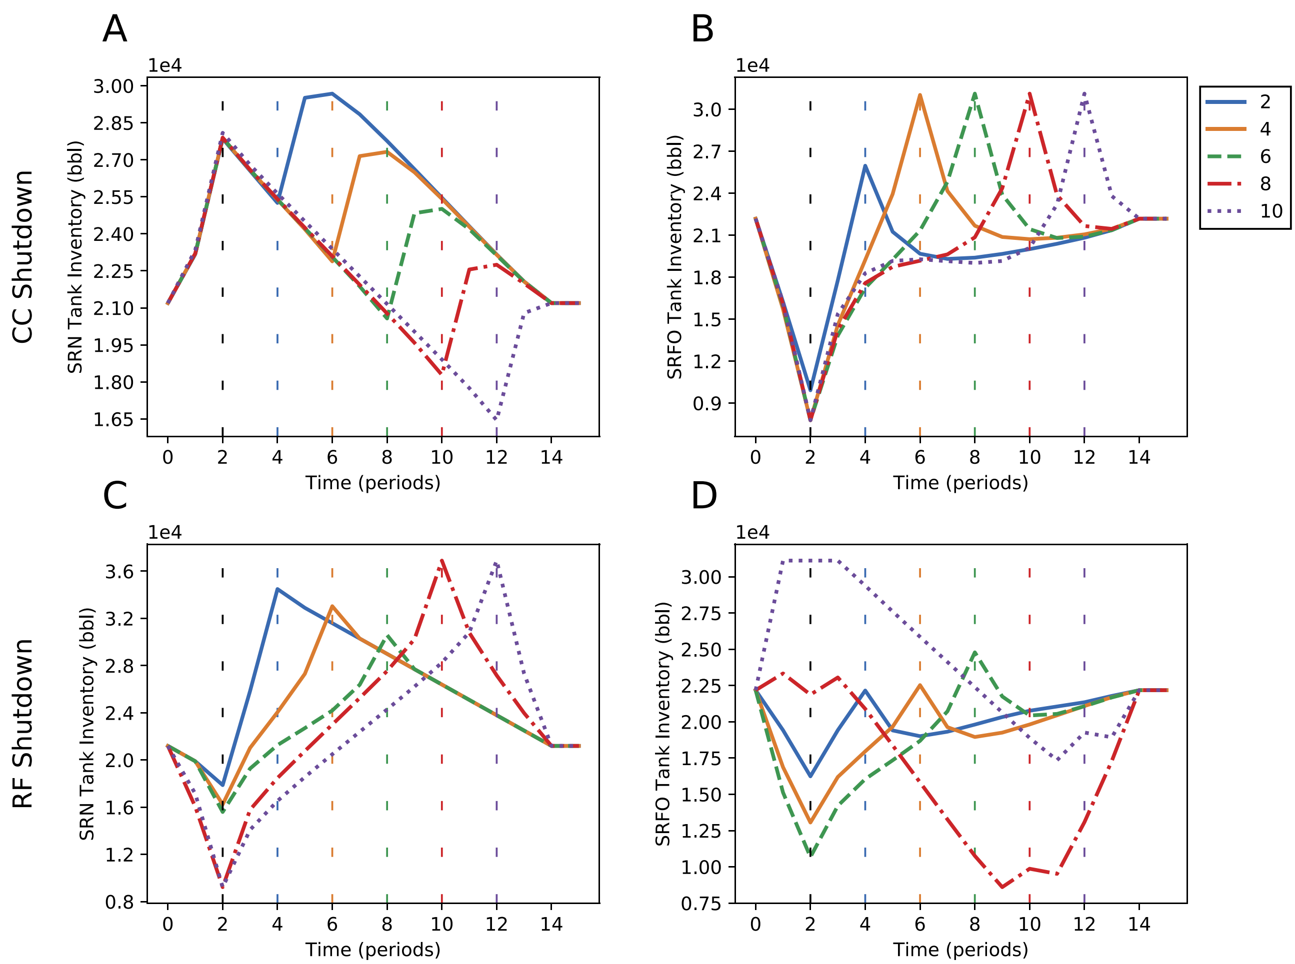

Fig. 10. Inventory profiles of the SRN-SRFO tank configuration over a 15-period planning horizon for on time pre-emptive shutdowns of the CC and the RF. Each profile corresponds to a shutdown length of 2–10 periods. The starting point of the shutdown is designated by a black vertical dashed line. The ending point of the shutdown for each scenario is designated by a color-coded dashed line.

the ultimate goal is to obtain a compromise solution. Each scenario is an operational subproblem with independent operational variables, such as material flow rates and inventories. A common set of design variables, i.e. the tank locations and capacities, link the subproblems together. The probability of occurrence of a given scenario is used to weight the contribution of each scenario to the compromise solution.

The following case consists of 40 scenarios and a planning horizon of 15 periods. Each scenario is a combined realization of each of four uncertain shutdown parameters: (1) the shutdown location, which may either be the catalytic cracker or the reformer, is characterized by a uniform distribution; (2) the start time of the shutdown, ranging from $t = 2$ to $t = 4$, is characterized by a normal distribution centered at 3; (3) the shutdown duration, rang-

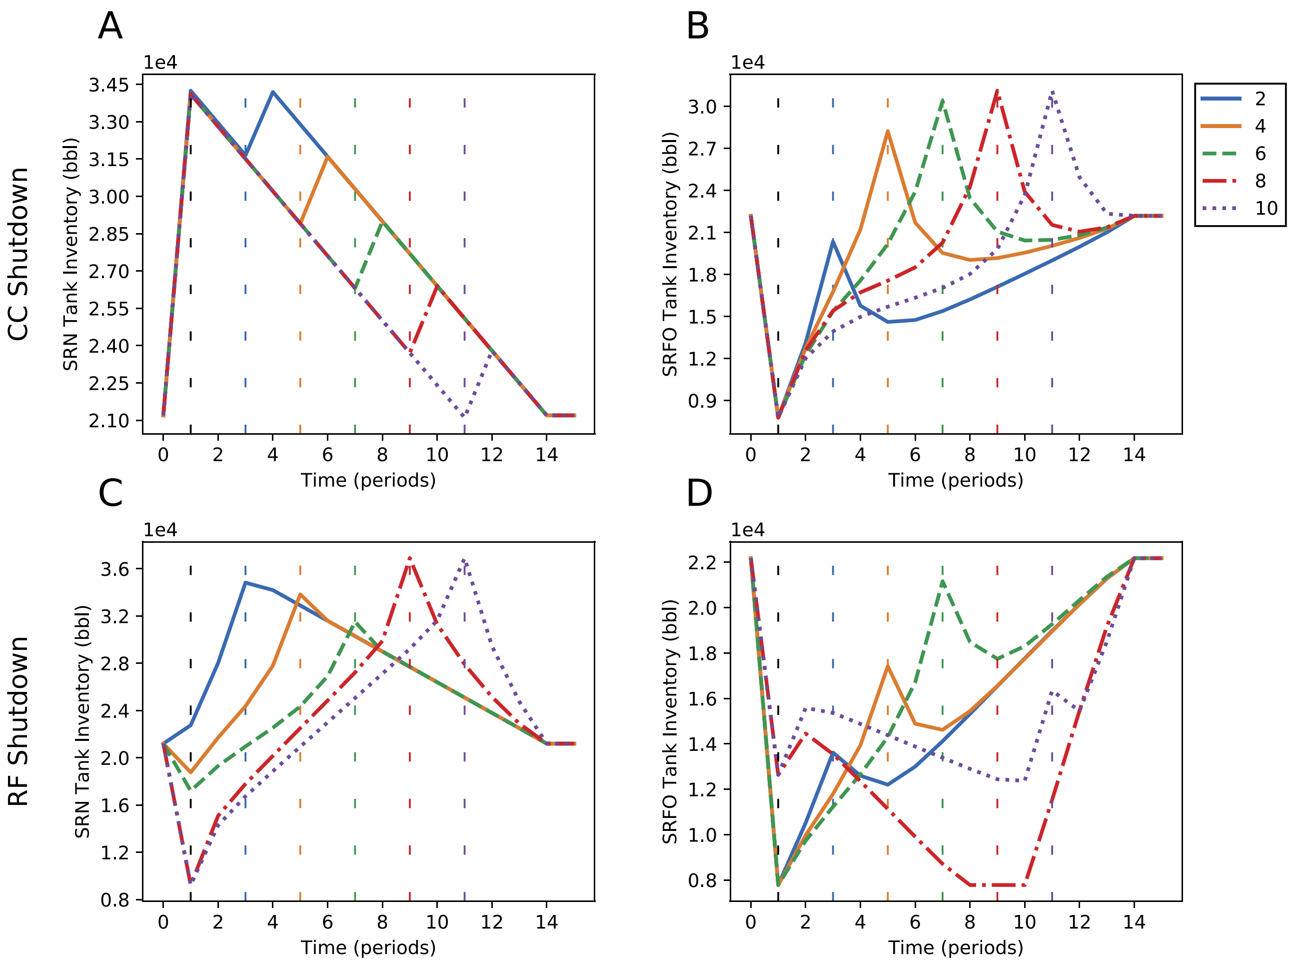

Fig. 11. Inventory profiles of the SRN-SRFO tank configuration over a 15-period planning horizon for early pre-emptive shutdowns of the CC and the RF. Each profile corresponds to a shutdown length of 2–10 periods. The starting point of the shutdown is designated by a black dashed line. The ending point of the shutdown for each scenario is designated by a color-coded dashed line.

Table 5
Optimal tank configuration for the deterministic and stochastic cases.

Tank	Starting Level (%) / Design Capacity (bbl)	
Location	Deterministic	Stochastic
SRG	0 / 0	0 / 0
SRN	0 / 0	45.9/46124
SRDS	0 / 0	0 / 0
SRFO	74.7/27 268	57.0/38892
RFG	0 / 0	0 / 0
CCG	0 / 0	0 / 0
CCFO	0 / 0	0 / 0

ing from 2 to 10 periods, is characterized by a normal distribution centered at 6; (4) whether the shutdown can be anticipated, with the pre-emptive and reactive cases sharing a probability ratio of 4:1. In the pre-emptive case (scenarios 1–30), the unit shutdowns are planned long in advance for the purpose of regular maintenance. The shutdown is expected to start at $t = 3$ and to last for 6 periods. However the exact starting points and durations of the shutdowns only become known at the beginning of the planning horizon and may differ from their expected values. In the reactive case (scenarios 31–40), the unit shutdowns are the result of unexpected failures and cannot be anticipated. To reflect this, the refinery operates according to its nominal base case policy prior to the start of the shutdown. For simplicity, it is assumed that all reactive shutdowns start at $t = 3$.

The weighted average of the profits of the compromise solution is \$22 297 889. The optimal tank configuration of the stochastic case is similar to that of the deterministic case, consisting of a larger SRFO tank and an additional tank at the SRN stream (Table 5). The inventory profiles of both tanks are illustrated in Figs. 10–13. For each scenario, a distinct strategy is adopted wherein the tanks act in concert to mitigate the impact of the unit shutdown on the economic objective.

The optimal strategy is greatly affected by the location and the duration of the shutdown. For example, in a pre-emptive shutdown of the CC, the stock of the SRN tank is depleted during the shutdown in order to maximize the capacity of the RF and to produce additional reformate (RFG) (Fig. 10A). The shutdown duration has a monotonic effect on the amount of SRN mobilized from the tank. Conversely, the SRFO tank fills up over the length of the shutdown (Fig. 10B). Recall that SRFO is a feedstock for the CC unit. Once the CC comes back online, the SRFO tank begins to mobilize its surplus material in order to produce high-octane gasoline feedstock. This action supplements the nominal RFG production of the RF unit, thus allowing the SRN tank to return to its initial level. Apart from the 2-period shutdown, the SRFO tank accumulates a fixed amount of material regardless of shutdown duration.

In the case of the analogous RF shutdowns, a different set of strategies is adopted. While the RF unit is offline, the SRN feed-

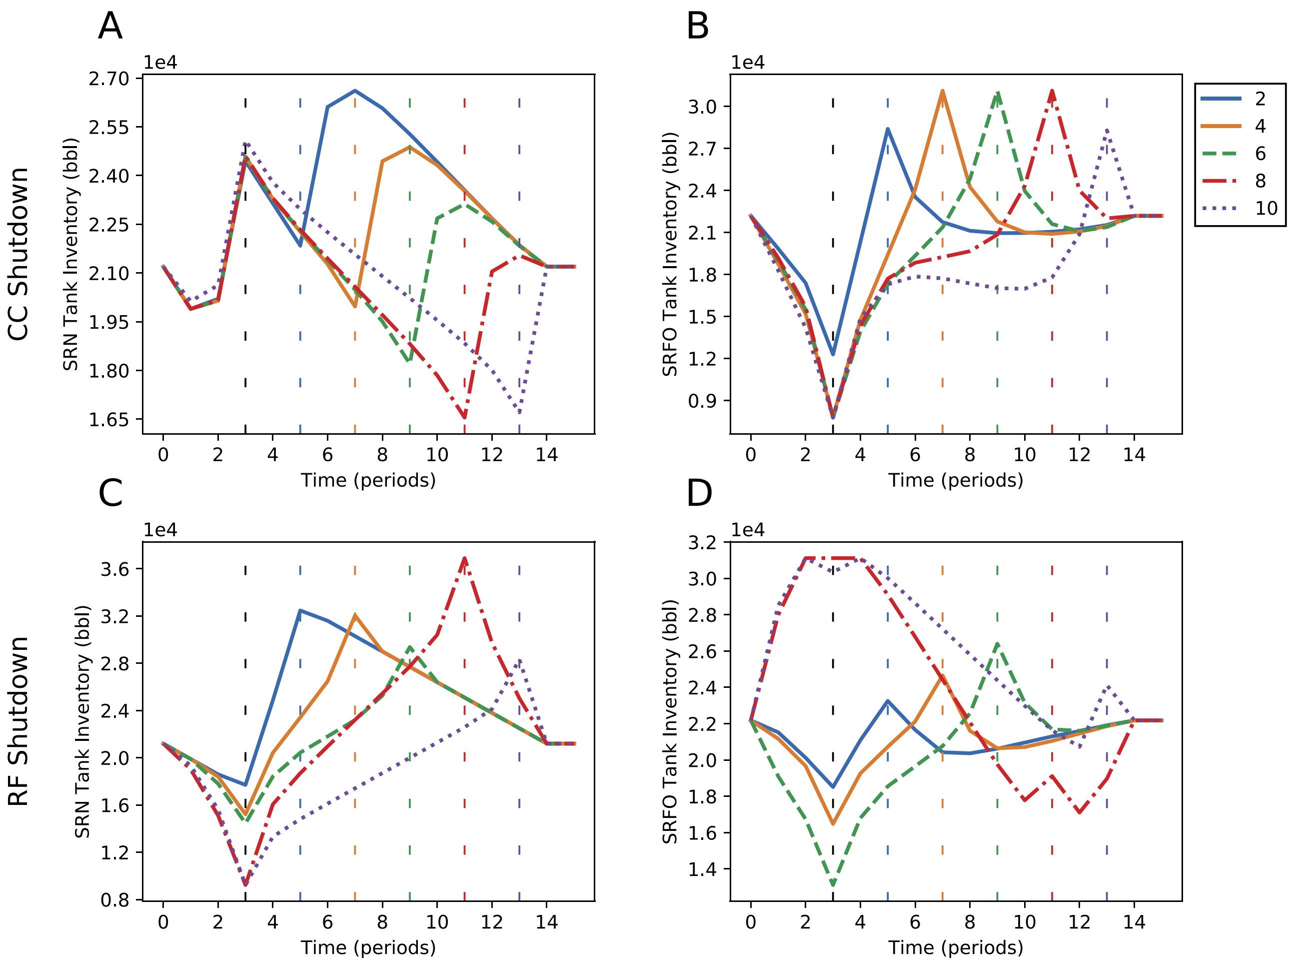

Fig. 12. Inventory profiles of the SRN-SRFO tank configuration over a 15-period planning horizon for late pre-emptive shutdowns of the CC and the RF. Each profile corresponds to a shutdown length of 2–10 periods. The starting point of the shutdown is designated by a black dashed line. The ending point of the shutdown for each scenario is designated by a color-coded dashed line.

stock is redirected to DF blending while the excess material is accumulated in the corresponding inventory tank (Fig. 10C). In parallel, the SRFO stock is either increased in the case of a short-duration shutdown (< 8) or depleted in the case of a long one (≥ 8) (Fig. 10D). During a short RF shutdown, the SRFO material that would otherwise be used in DF blending is stocked in the SRFO tank. Once the RF unit returns online, the SRN inventory is mobilized to maximize its output and make up for lost PG production, whereas the SRFO stock is used to maintain the production of DF. For a longer shutdown of the RF unit, the quantity of excess SRN is such that it is either unprofitable or infeasible to run the distillation unit at full capacity. The throughput is thus lowered to as much as 60% of its capacity, which results in insufficient production of SRFO to saturate the CC unit. To compensate for this, the SRFO stock is mobilized, thereby maintaining the production of CCG and CCFO.

The effect of the start time of the shutdown has a less pronounced effect on the mitigation strategies. The inventory profiles of the early-start (Fig. 11) and late-start shutdowns (Fig. 12) follow the same pattern as those of the on-time shutdowns (Fig. 10). The early-start shutdowns only afford a single period of preparation, but benefit from a longer restoration phase. Conversely, the late-start shutdowns allow for additional preparatory action at the expense of a shorter restoration window. The constraints of a shorter preparation phase can effectively be removed by designing sufficiently large tanks with appropriate starting levels such that the refinery has the necessary flexibility to immediately react to a unit shutdown. On the other hand, a shorter restoration phase limits the deviation of a tank level from its initial value, which may necessitate an adjustment to the operating policy.

The reactive shutdown scenarios represent a special case where no preparation time is afforded. These scenarios function as a stress test on the design of the inventory system. The tank configuration must ultimately be able to absorb sudden increases or decreases in stock in response to an unplanned shutdown. The SRN tank nearly reaches its minimum allowable level during the end point of the 10-period reactive shutdown of the CC (Fig. 13A). Likewise, the SRFO tank reaches its maximum allowable level at $t = 4$ and, for durations greater than 2 periods, once more during the end point of the shutdown (Fig. 13B). During a reactive RF shutdown, the SRN tank is maximized at the end of the shutdown for almost every duration (Fig. 13C), while the SRFO tank reaches it allowable minimum towards the end of an 8-period shutdown (Fig. 13D). These worst-case scenarios effectively impose minimum tank sizes on the compromise solution such that feasible mitigation strategies may be adopted in response to them.

When only a single scenario is considered, as with the deterministic case, the design of the inventory tank system may prove to be too rigid to appropriately mitigate the shutdown impacts of other scenarios that may arise over the lifetime of the plant. On the other hand, the stochastic formulation produces a versatile solution that allows the refinery to react to a multitude of different scenarios. The tank locations are selected such that the system can respond to unit shutdowns in different sections of the plant. Moreover, the optimal starting levels and tank sizes allow the tanks to suddenly stock or destock material during the worst-case reactive

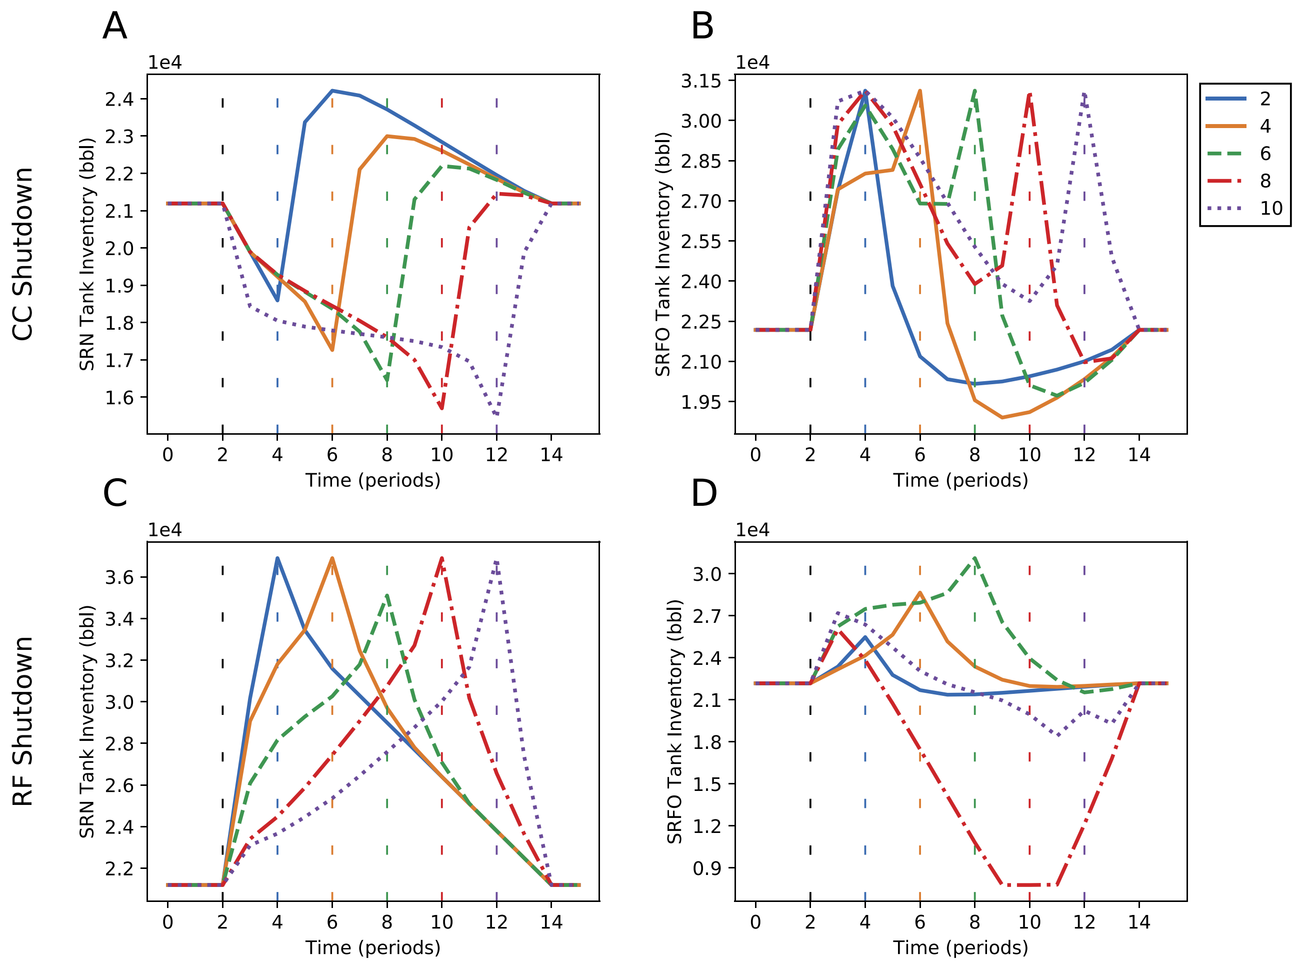

Fig. 13. Inventory profiles of the SRN-SRFO tank configuration over a 15-period planning horizon for reactive shutdowns of the CC and the RF. Each profile corresponds to a shutdown length of 2–10 periods. The starting point of the shutdown is designated by a black dashed line. The ending point of the shutdown for each scenario is designated by a color-coded dashed line.

Table 6
Computational results of the unit shutdown case studies.

Case Study	Planning Horizon	Problem Size			Solution Time (s)
		Cont. Variables	Binaries	Constraints	
Base Case	10	1028	0	1125	0.031
No tanks	10	1028	0	1125	0.016
CCG tank	10	1028	0	1125	0.031
CCFO tank	10	1028	0	1125	0.031
CCG-CCFO tanks	10	1028	0	1125	0.031
3P prep	15	1528	0	1665	0.031
2P prep	15	1528	0	1665	0.047
1P prep	15	1528	0	1665	0.031
Reactive	15	1528	0	1693	0.031
Design (S=1)	15	1528	7	1665	0.328
Design (S=40)	15	63 491	7	78 165	375.8

shutdowns. Although the stochastic formulation results in significantly larger problem sizes, it has the potential to yield robust solutions that are significantly more flexible.

4.4. Computational results

The CPLEX (IBM) and Gurobi optimization solvers were used to solve the case studies on an Intel Core i7-9700 CPU (8 cores) with 16 GB of RAM. The optimality criterion of both solvers was set to 0.1%. The problem was formulated in the GAMS modelling language as an MILP/MIQP. The error tolerance (ϵ) of the profit when solving the second MIQP problem was set to 0.1%.

The solution times and problem sizes of each case study are summarized in Table 6. The solution time remains relatively stable for the operational problem, even when the planning horizon is increased to 15 periods. The addition of binary variables in the design problem translates to a marginal increase in solution time. The stochastic case study, on the other hand, required substantially more time to solve. The secondary MIQP problem is responsible for the bulk of this runtime.

5. Conclusion

This work proposes an optimization framework for multiperiod refinery planning in which inventories are used to mitigate the impact of process unit shutdowns. A two-tier multi-objective optimization problem is described where the profit of the refinery is maximized and the production/inventory profiles are smoothed during a process unit shutdown event. Both operational and design problems are considered. The former is posed as a deterministic optimization problem to determine the operating policy that optimally mitigates the impact of a specific process unit shutdown for a given tank configuration. The design problem, on the other hand, is formulated as a two-stage stochastic programming problem where the optimal configuration of an inventory tank system is determined for a set of different shutdown scenarios. The stochastic problem hedges against a multitude of distinct shutdown events in order to determine the best compromise design for the inventory system to allow the refinery to react to unit shutdowns of different durations and at different locations with varying degrees of preparation time.

The application of the proposed formulation to a simplified refinery case study serves as a proof-of-concept for how model-based optimization may be employed to mitigate the impact of process unit shutdowns in real-world refineries. The operational problem reveals which streams are critical for maintaining the operation of the refinery under partial shutdown conditions, as well as what actions might be taken prior to, during, and after the shutdown in order to reduce losses of revenue. In addition, the use of cumulative quality constraints allow for the application of compensatory blending policies. The design problem allows engineers to account for process unit shutdowns and response policies early in the design of a refinery, while minimizing the capital cost.

The utility of the proposed problem formulation may be further investigated by applying it to a more realistic refinery model. Such models typically feature nonlinear yield relationships and blending rules, resulting in a large-scale nonconvex mixed-integer nonlinear program (MINLP). Given the difficulty of solving this class of problems, a more sophisticated solution strategy may be necessary to yield solutions within reasonable computation times. This may include decomposition approaches, possibly coupled with the use of surrogate models to reduce the extent of nonlinearity and/or nonconvexity.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Ariel Boucheikhchoukh: Conceptualization, Methodology, Software, Investigation, Writing – original draft, Visualization. **Valentin Berger:** Conceptualization, Methodology, Software, Investigation, Writing – review & editing, Visualization. **Christopher L.E. Swartz:** Supervision, Conceptualization, Methodology, Writing – review & editing. **Antoine Deza:** Supervision, Conceptualization, Methodology, Writing – review & editing. **Alexander Nguyen:** Conceptualization, Methodology, Software. **Shaffiq Jaffer:** Conceptualization, Funding acquisition.

Acknowledgements

The support of TotalEnergies SE and the McMaster Advanced Control Consortium (MACC) is gratefully acknowledged.

Appendix A.

Table A1
Volumetric yield coefficients ($X_{u,i,j}$). Source: Pike (1986).

Input	Output									
	FGAD	SRG	SRN	SRDS	SRFO	FGRF	RFG	FGCC	CCG	CCFO
CRUDE	35.42	0.27	0.237	0.087	0.372	-	-	-	-	-
SRNRF	-	-	-	-	-	158.7	0.928	-	-	-
SRDSCC	-	-	-	-	-	-	-	336.9	0.619	0.189
SRFOCC	-	-	-	-	-	-	-	386.4	0.688	0.2197

Table A2
Unit capacities (F_u^{max}) and operating costs (C_u^{op}). Source: Pike (1986).

Unit	Capacity (bbl)	Operating cost ($/bbl)
AD	100,000	1.00
RF	25,000	2.50
CC	30,000	2.20

Table A3
Crude cost and product prices. Sources: Pike (1986) and U.S. Energy Information Administration (2019).

Stream	Cost/Price (US$/bbl)
CRUDE	58.38
FGAD	0.02
FGRF	0.02
FGCC	0.02
PG	92.80
RG	76.47
DF	78.36
FO	55.92

Table A4
Minimum demands for the final products.

Final Product	Minimum Demand (bbl/day)
PG	23 556.55
RG	11 260.15
DF	8 419.05
FO	2826.30

Appendix B

Table B1
Unit conversions.

Quantity	Unit	Equivalent SI Unit
Volume	bbl	0.1590 m^3
Density	lb/bbl	2.8528 kg/m^3
Vapour Pressure	mmHg	133.322 Pa
Sulfur Content	lb/bbl	2.853 kg/m^3

References

Balthazaar, A., 2006. Dynamic Optimization of Multi-unit Systems Under Failure Conditions. M.A.Sc. Thesis, McMaster University, Hamilton, Ontario, Canada.

Bengtsson, J., Bredström, D., Flisberg, P., Rönnqvist, M., 2013. Robust planning of blending activities at refineries. Journal of the Operational Research Society 64 (6), 848–863. doi:10.1057/jors.2012.86.

Birge, J.R., Louveaux, F., 2011. Introduction to Stochastic Programming, Vol. 150, 2nd Springer, New York doi:10.1007/978-1-4419-1642-6_1.

Cao, Y., Swartz, C.L.E., Flores-Cerrillo, J., 2016. Optimal dynamic operation of a high-purity air separation plant under varying market conditions. Ind. Eng. Chem. Res. 55 (37), 9956–9970. doi:10.1021/acs.iecr.6b02090.
Chong, Z., Swartz, C.L.E., 2013. Optimal operation of process plants under partial shutdown conditions. AIChE J. 59 (11), 4151–4168. doi:10.1002/aic.14154.
Chong, Z., Swartz, C.L.E., 2016. Optimal response under partial plant shutdown with discontinuous dynamic models. Comput. Chem. Eng. 86, 120–135. doi:10.1016/j.compchemeng.2015.12.011.
Dubé, J.-F. H., 2000. Pulp Mill Scheduling: Optimal use of Storage Volumes to Maximize Production. M.Eng. Thesis, McMaster University, Hamilton, Ontario, Canada.
Gerardi, D., Marlin, T.E., Swartz, C.L.E., 2013. Optimization of primary steelmaking purchasing and operation under raw material uncertainty. Ind. Eng. Chem. Res. 52 (35), 12383–12398. doi:10.1021/ie3035543.
Guerra, O.J., Le Roux, G.A.C., 2011. Improvements in petroleum refinery planning: 1. formulation of process models. Industrial & Engineering Chemistry Research 50 (23), 13403–13418.
Guillén, G., Mele, F.D., Bagajewicz, M.J., Espuña, A., Puigjaner, L., 2005. Multiobjective supply chain design under uncertainty. Chem Eng Sci 60 (6), 1535–1553. doi:10.1016/j.ces.2004.10.023.
Joly, M., Moro, L.F., Pinto, J.M., 2002. Planning and scheduling for petroleum refineries using mathematical programming. Braz. J. Chem. Eng. 19 (2), 207–228. doi:10.1590/S0104-66322002000200008.
Neiro, S.M., Pinto, J.M., 2005. Multiperiod optimization for production planning of petroleum refineries. Chem Eng Commun 192 (1–3), 62–88. doi:10.1080/00986440590473155.
Neiro, S.M.S., Pinto, J., 2006. Lagrangean decomposition applied to multiperiod planning of petroleum refineries under uncertainty. Latin American Applied Research 36, 213–220.
Pike, R.W., 1986. Optimization for Engineering Systems. Van Nostrand Reinhold, New York.
Pinto, J.M., Joly, M., Moro, L.F., 2000. Planning and scheduling models for refinery operations. Comput. Chem. Eng. 24 (9–10), 2259–2276.
Turton, R., Bailie, R.C., Whiting, W.B., Shaeiwitz, J.A., Bhattacharyya, D., 2012. Analysis, Synthesis, and Design of Chemical Processes, 4 Prentice Hall, Upper Saddle River.
U.S. Energy Information Administration, 2019. November 2019 Monthly Energy Review. Technical Report. U.S. Energy Information Administration, Washington D.C. doi:10.1044/leader.ppl.24062019.28.
You, F., Wassick, J.M., Grossmann, I.E., 2009. Risk management for a global supply chain planning under uncertainty: models and algorithms. AIChE J 55 (4), 931–946. doi:10.1002/aic.11721.