applied sciences

MDPI

*Article*

# Risk-Aware Crude Oil Scheduling in Petrochemical Supply Chains: A CVaR-Driven Reactive GRASP Simheuristic

**Antonio Giallanza [1] and Giuseppe Marannano [2,*]**

[1] Department of Engineering, University of Messina, Contrada di Dio, 98166 Messina, Italy; antonio.giallanza@unime.it
[2] Department of Engineering, University of Palermo, Viale delle Scienze, 90128 Palermo, Italy
* Correspondence: giuseppe.marannano@unipa.it

**Abstract**

The scheduling of crude oil operations in marine refineries is a complex combinatorial problem, exacerbated by stochastic disruptions like vessel delays and port congestion. Traditional deterministic and expected-value approaches fail to mitigate high-impact tail events, causing severe demurrage and production bottlenecks. To address this, we propose a novel CVaR-Driven Reactive GRASP Simheuristic. This framework hybridizes GRASP with Monte Carlo simulation, embedding Conditional Value-at-Risk (CVaR) into the adaptive memory to actively steer the search away from catastrophic logistical gridlocks. Overcoming standard “unlimited port capacity” assumptions, the model endogenously calculates demurrage dynamics and introduces an automated Failure Taxonomy for explainable insights. Evaluated on a 30-day industrial case study, representing a standard short-term operational scheduling horizon, under baseline conditions and severe dynamic disruptions (vessel delays, unit maintenance), the diagnostic reveals that over 80% of scheduling failures stem from endogenous port congestion rather than internal dead-ends. Furthermore, a comprehensive ablation study mathematically validates the superiority of the CVaR-driven memory over standard expected-cost optimization in preventing catastrophic tail-risk scenarios. Results demonstrate that this CVaR-driven approach effectively absorbs stochastic shocks, prevents stockouts, and minimizes worst-case costs, generating highly robust schedules in under three minutes. Ultimately, it provides a robust, risk-aware Decision Support System (DSS) for supply chain and operations managers.

**Keywords:** crude oil scheduling; simheuristics; Conditional Value at Risk; supply chain resilience; Reactive GRASP; decision support systems

## 1. Introduction

The petrochemical industry is a cornerstone of the global economy, and the coordinated scheduling of crude oil procurement, unloading, storage, blending, and processing is essential to sustaining profitable refinery operations [1–3]. Crude oil scheduling constitutes the critical operational interface between upstream supply logistics and the uninterrupted execution of downstream crude distillation units (CDUs). In globally interconnected markets, this interface is highly vulnerable to stochastic disruptions, including maritime shipping delays, supply shortages, and fluctuations in refinery operating conditions [4–6]. These vulnerabilities are particularly severe in marine access refineries. In such settings, the coordination of port channels, berths, tidal windows, storage resources, and transfer operations can generate severe congestion and demurrage costs that propagate directly



Received: 21 May 2026
Revised: 30 June 2026
Accepted: 2 July 2026
Published: 5 July 2026

into refinery scheduling decisions [7–10]. When deterministic schedules confront these dynamic real-world conditions, the resulting shocks frequently trigger bottlenecks, stockout risks, and significant economic losses, motivating the urgent need for adaptive, risk-aware scheduling strategies [11–13].

The physical execution of crude oil scheduling is governed by tightly coupled operational constraints, making its modeling and optimization a highly combinatorial and NP-hard problem [14,15]. A feasible schedule requires the synchronized management of vessel unloading, intermediate inventory levels, mandatory brine-settling times, safety-stock policies, single-crude tank allocation rules, and precise blending recipes to satisfy strict CDU feed specifications [16–19]. Designing optimization models that respect these rigid physical boundaries without generating infeasible schedules remains a major operational challenge.

Historically, the scientific literature has tackled this challenge predominantly through deterministic mathematical programming. Foundational discrete- and continuous-time MILP/MINLP formulations successfully modeled industrially important features such as inventory dynamics, settling times, and nonlinear blending constraints to optimize costs under perfect information [20–23]. To improve operational realism, subsequent research broadened the perspective from isolated short-term scheduling toward integrated planning. Studies demonstrated that procurement, terminal logistics, and refinery-wide decisions must be optimized simultaneously to preserve operational feasibility [24–26]. This integrated logic was further reinforced at the terminal–refinery interface, emphasizing that port-side routing, berth allocation, and unloading operations act as complex scheduling layers whose dynamics dictate downstream refinery performance [2,9,27]. Concurrently, to cope with the scale and combinatorial difficulty of these integrated models, simulation-based and metaheuristic approaches, such as structure-adapted genetic algorithms and evolutionary multiobjective optimization, emerged as valuable tools when exact models became computationally intractable for practical deployment [26,28–30].

Recognizing the limitations of deterministic assumptions, a major shift in the literature has been the transition toward uncertainty-aware and disruption-oriented formulations. Chance-constrained, robust, and data-driven optimization models have been widely adopted to hedge against volatility in supply, crude quality, and demand [13,31–33]. In parallel, reactive and preventive scheduling methodologies, including model predictive control (MPC) and safe reinforcement learning, have been proposed to manage supply-chain disruptions, VLCC delays, and dynamic rescheduling [11,13,19,34]. Recent literature has also begun to incorporate formal risk management into these mathematical programming frameworks. Notably, Garcia-Verdier et al. [35] developed a two-stage stochastic mixed-integer nonlinear programming (MINLP) model to handle uncertain vessel arrival dates, actively minimizing tail risk by utilizing the Conditional Value-at-Risk (CVaR) measure as the objective function. While such exact mathematical approaches successfully integrate CVaR, they suffer from well-known scalability limitations when the number of uncertainty scenarios increases. Furthermore, the integration of CVaR within adaptive metaheuristic frameworks for this domain remains largely unexplored.

Despite these significant advances, critical operational and methodological gaps remain. First, apart from recent exact formulations, most existing uncertainty-aware simheuristics and metaheuristics guide the optimization search primarily through expected performance, scenario feasibility, or robust admissibility. They fundamentally lack a mechanism to explicitly steer the metaheuristic search away from low-probability, high-impact "tail events" (e.g., catastrophic port gridlocks or total refinery stockouts). Second, many models still rely on simplified representations of terminal operations, such as "unlimited port" assumptions or static, exogenous delay penalties [36]. These simplifications fail to capture the endogenous, compounding dynamics of port congestion and the nonlinear growth of demurrage costs. Finally, while robust and reactive methods improve schedule resilience, they operate largely as algorithmic "black boxes". As noted in schedulability studies [37,38], many mathematically attractive schedules fail in practice due to deeply hidden hybrid dynamic constraints. Existing models provide limited diagnostic insight into why candidate schedules fail, thereby diminishing their practical value for human schedulers who must make real-time, explainable decisions during a crisis.

To bridge these gaps, this paper proposes a CVaR-Driven Reactive GRASP Simheuristic for the risk-aware crude oil scheduling problem. The proposed approach seamlessly hybridizes a Reactive Greedy Randomized Adaptive Search Procedure (GRASP) with a Monte Carlo simulation engine. Unlike traditional risk-neutral simheuristics or robust optimization methods, our framework explicitly evaluates the tail risk of candidate schedules and uses CVaR estimates to continuously update the metaheuristic's adaptive memory. By feeding back the severity of adverse disruption scenarios into the constructive search, the algorithm actively learns to avoid configurations susceptible to extreme delay penalties and operational bottlenecks.

While risk measures have recently been integrated into exact mathematical models for crude oil scheduling (e.g., [35]), their application has primarily been restricted to static objective function evaluations. The primary methodological contribution of this paper is the development of a simheuristic Reactive GRASP where a tail risk measure (CVaR) acts as the internal driver to dynamically update the adaptive memory of the algorithm. To the best of our knowledge, this is the first work to utilize CVaR iteratively within a reactive metaheuristic for crude oil scheduling. Consequently, the specific novelties of this study are:

- Algorithmic Innovation (Tail-Risk Steering): We introduce a novel simheuristic where the adaptive memory (e.g., the $\alpha$ probabilities) of a Reactive GRASP is directly driven by CVaR feedback. This differentiates our method from prior reactive and robust approaches [13,32,35] by embedding risk aversion directly into the learning mechanism of the constructive search, rather than treating uncertainty merely as a passive evaluation metric.
- Modeling Realism (Endogenous Port Congestion): We explicitly model dynamic port congestion, capturing the nonlinear, cascading effects of cumulative vessel waiting and demurrage costs. This overcomes the widespread "unlimited port" or fixed-delay assumptions, fully endogenizing the terminal–refinery risk interface highlighted by recent literature [9,10].
- Explainable Optimization (Automated Failure Taxonomy): We move beyond traditional black-box optimization by developing an automated diagnostic framework. The algorithm systematically classifies the physical root causes of discarded solutions (e.g., settling blocks, safety-stock violations, port congestion deadlocks), transforming the optimizer into a transparent decision-support tool that provides actionable insights for industrial practitioners.

The remainder of this paper is organized as follows. Section 2 formally defines the physical problem and its deterministic operational constraints. Section 3 details the architecture of the proposed CVaR-driven Reactive GRASP simheuristic. Section 4 presents the industrial case study, the baseline validation, and the computational results of dynamic rescheduling under disruption scenarios. Finally, Section 5 concludes the paper and outlines future research directions.

https://doi.org/10.3390/app16136733

## 2. Problem Description and Deterministic Constraints

### *2.1. Problem Statement and Assumptions*

The crude oil scheduling problem addressed in this study is modeled on a representative marine-access petrochemical refinery. The conceptual layout of the system, illustrating the continuous flow from the maritime port interface through the storage tank farm and into the Crude Distillation Unit (CDU), is depicted in Figure 1.

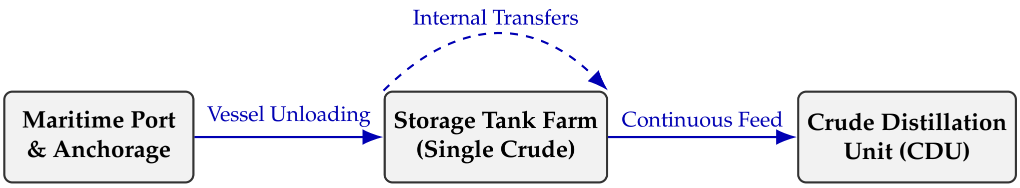


**Figure 1.** Representative diagram of the petrochemical supply chain considered in the problem statement.

The scheduling environment is governed by the following theoretical assumptions: (1) the planning horizon is discretized into daily operational periods; (2) demurrage costs accrue linearly per day for vessels waiting in the port queue; (3) the continuous topping rate must be strictly maintained within the specific physical bounds of the CDU; and (4) stochastic disruptions manifest primarily as delays in vessel arrival dates or unexpected CDU maintenance shutdowns. The given items include the deterministic arrival schedule (under baseline conditions), the maximum volumetric capacities of the tanks, the required blending recipes, and the economic penalty weights for operational violations. The items to be dynamically determined by the optimizer include: the daily unloading volume from each vessel to specific tanks, the exact daily blend composition fed to the CDU, the execution of internal tank-to-tank transfers, and the management of the port waiting queue. To formalize the problem mathematically, we define a discrete-time planning horizon $\mathcal{T} = \{1, 2, \ldots, |\mathcal{T}|\}$. Let $\mathcal{C} = \{\text{Azeri}, \text{WTI}, \text{CPC}\}$ be the set of crude oil types, $\mathcal{K} = \{T_1, T_2, T_3\}$ be the set of storage tanks, and $\mathcal{V}$ be the set of arriving vessels. The deterministic operational rules defining the feasible solution space are translated into the mathematical constraints presented below.

### *2.2. System Architecture and Inventory Balances*

The refinery receives crude oil via maritime vessels that anchor at the port. Crude is discharged into the tank farm, which acts as a buffer before being fed to the CDU. The mass balance for each tank $k \in \mathcal{K}$ and crude type $c \in \mathcal{C}$ at the end of day $t$ is defined as:

$$I_{k,c,t} = I_{k,c,t-1} + \sum_{v \in \mathcal{V}} U_{v,k,c,t} + \sum_{k' \neq k} X_{k',k,c,t} - \sum_{k' \neq k} X_{k,k',c,t} - F_{k,c,t} \quad \forall k, c, t \tag{1}$$

where $I_{k,c,t}$ is the inventory level, $U_{v,k,c,t}$ is the amount of crude unloaded from vessel $v$ to tank $k$, $X_{k',k,c,t}$ represents internal transfers, and $F_{k,c,t}$ is the volume fed to the CDU.

Furthermore, the total volume discharged from a specific vessel into the tank farm on day $t$ cannot exceed the available crude volume on that vessel from the previous day:

$$\sum_{k \in \mathcal{K}} U_{v,k,c,t} \leq Q_{v,c,t-1} \quad \forall v, c, t \tag{2}$$

Each tank has a maximum volumetric capacity $Cap_k$. In our baseline configuration, $Cap_{T1} = 65$ kt, and $Cap_{T2} = Cap_{T3} = 120$ kt. The total inventory inside a tank cannot exceed its capacity:

$$\sum_{c \in \mathcal{C}} I_{k,c,t} \leq Cap_k \quad \forall k, t \tag{3}$$

To absorb unpredicted delays and prevent overflow, the refinery enforces strict global safety margins. A global Safety Stock ($SS = 120$ kt) must be maintained to avoid starvation, and a global Safety Space ($SF = 30$ kt) must be kept free to accommodate emergency discharges:

$$SS \leq \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} I_{k,c,t} \leq \left( \sum_{k \in \mathcal{K}} Cap_k \right) - SF \quad \forall t \tag{4}$$

*2.3. Single-Crude Allocation and Settling Dynamics*

To prevent cross-contamination and unpredictable blend properties, the tank farm operates under a strict single-crude policy. A tank can only hold one type of crude oil at any given time. Let $Y_{k,c,t} \in \{0, 1\}$ be a binary variable equal to 1 if tank $k$ stores crude $c$ on day $t$:

$$I_{k,c,t} \leq Cap_k \cdot Y_{k,c,t} \quad \forall k, c, t \tag{5}$$

$$\sum_{c \in \mathcal{C}} Y_{k,c,t} \leq 1 \quad \forall k, t \tag{6}$$

Furthermore, physical operations on the tanks are mutually exclusive. A tank cannot simultaneously receive crude (from a ship or another tank) and feed the CDU on the same day. To formulate this in standard Mixed-Integer Programming (MIP) canonical form, let $Z_{k,t} \in \{0, 1\}$ be an auxiliary binary variable, and $M$ be a sufficiently large positive constant (set to $\sum_{k \in \mathcal{K}} Cap_k$, representing the maximum theoretical farm capacity):

$$\sum_{c \in \mathcal{C}} F_{k,c,t} \leq M \cdot (1 - Z_{k,t}) \quad \forall k, t \tag{7}$$

$$\sum_{v \in \mathcal{V}} \sum_{c \in \mathcal{C}} U_{v,k,c,t} + \sum_{k' \neq k} \sum_{c \in \mathcal{C}} X_{k',k,c,t} \leq M \cdot Z_{k,t} \quad \forall k, t \tag{8}$$

Additionally, after receiving crude, the mixture must undergo a mandatory brine-settling period ($\tau = 1$ day) for water separation before it can become eligible to feed the topping unit. Let $S_{k,t}$ be an integer variable tracking the consecutive days a tank has been settling. Mirroring the receiving conditions formalized above, $S_{k,t}$ is updated as:

$$S_{k,t} = \begin{cases} 0 & \text{if } Z_{k,t} = 1 \\ S_{k,t-1} + 1 & \text{otherwise} \end{cases} \tag{9}$$

A tank $k$ is eligible to feed the CDU ($F_{k,c,t} > 0$) only if the settling time constraint is satisfied:

$$F_{k,c,t} \leq Cap_k \cdot \mathbb{I}(S_{k,t} \geq \tau) \quad \forall k, c, t \tag{10}$$

Note that while the bound $Cap_k$ in the indicator function constraint is intentionally loose and is further tightened implicitly by the mass balance constraint Equation (1), it serves exclusively as a logical activation gate. Furthermore, it is important to explicitly state that Equations (8) and (9) are not formulated as strict linear mathematical programming constraints for an exact solver. Rather, they are defined as procedural algorithmic constraints strictly evaluated during the constructive phase of our heuristic to enforce physical eligibility without artificially over-constraining the continuous mass balance variables.

*2.4. Topping Feed and Blending Recipes*

The crude distillation unit (CDU) requires a continuous and stable feed. The total daily processing rate, $R_t$, is constrained by the technical limits of the topping unit ($R^{min} = 16$ kt/day, $R^{max} = 22$ kt/day):

$$R_t = \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} F_{k,c,t} \quad \forall t \tag{11}$$

$$R^{min} \leq R_t \leq R^{max} \quad \forall t \tag{12}$$

The total monthly production target is set to 540 kt. In addition to volumetric limits, the CDU feed must respect specific chemical blending recipes dictated by shareholders' economic strategies. While Azeri acts as the base crude, the inclusion of WTI and CPC is strictly bounded to prevent distillation anomalies. The blend fractions are limited to a maximum of 50% for WTI and 25% for CPC:

$$\sum_{k \in \mathcal{K}} F_{k,\text{WTI},t} \leq 0.50 \cdot R_t \quad \forall t \tag{13}$$

$$\sum_{k \in \mathcal{K}} F_{k,\text{CPC},t} \leq 0.25 \cdot R_t \quad \forall t \tag{14}$$

### 2.5. Port Congestion and Demurrage Accumulation

A critical aspect of marine-access refineries is the management of port logistics. Arriving vessels (in sizes of 80 kt or 130 kt) anchor in the port. If the tank farm lacks sufficient compatible free space to discharge the vessel completely, the ship remains in the waiting queue, generating cumulative demurrage costs and blocking subsequent arrivals.

Let $Q_{v,c,t}$ be the residual volume of crude $c$ on vessel $v$ waiting in the port at the end of day $t$. The port queue is updated as:

$$Q_{v,c,t} = Q_{v,c,t-1} - \sum_{k \in \mathcal{K}} U_{v,k,c,t} \quad \forall v, c, t \tag{15}$$

The total backlog in the port, $B_t = \sum_{v,c} Q_{v,c,t}$, represents the severity of the congestion. Avoiding the uncontrolled growth of $B_t$ is a primary logistical challenge, as high backlog levels exponentially increase demurrage penalties and threaten the feasibility of the entire supply chain. Finally, all physical volume and inventory tracking variables must strictly adhere to non-negativity constraints:

$$I_{k,c,t}, U_{v,k,c,t}, X_{k',k,c,t}, F_{k,c,t}, Q_{v,c,t} \geq 0 \quad \forall k, k', c, v, t \tag{16}$$

The deterministic constraints outlined in this section define the boundaries of the physical solution space. However, due to the stochastic nature of vessel arrival delays (characterized by a known probability distribution ranging from −1 to +6 days), a schedule that is mathematically feasible under nominal conditions may become severely bottlenecked in practice. While solving these equations as an exact stochastic Mixed-Integer Nonlinear Programming (MINLP) model would guarantee mathematical optimality, such formulations are notoriously intractable under severe uncertainty due to the exponential explosion of scenarios. To overcome this limitation and solve the problem efficiently, the deterministic boundaries formalized above act strictly as procedural constraints governing the constructive phase of the risk-aware metaheuristic framework described in the following section.

## 3. The CVaR-Driven Reactive GRASP Simheuristic

To solve the highly constrained, stochastic problem described in Section 2, we propose a novel simheuristic framework. Building upon the foundational GRASP methodology established by Feo and Resende [39] and its reactive extension by Prais and Ribeiro [40], our framework couples a semi-greedy constructive phase with a Monte Carlo simulation engine. Unlike traditional reactive metaheuristics that update their search memory based on deterministic costs or expected simulation values, our approach introduces a tail-risk steering mechanism. To the best of our knowledge, this represents a first-of-its-kind integration within simheuristics explicitly tailored for crude oil scheduling. The adaptive memory is driven by the Conditional Value-at-Risk (CVaR) of the objective function, dynamically guiding the constructive heuristic away from schedules vulnerable to severe port congestion and stockout cascading effects. The complete architectural flowchart of the proposed framework, illustrating the interaction between the constructive phase, local search, simulation engine, and risk-aware memory update, is depicted in Figure 2.

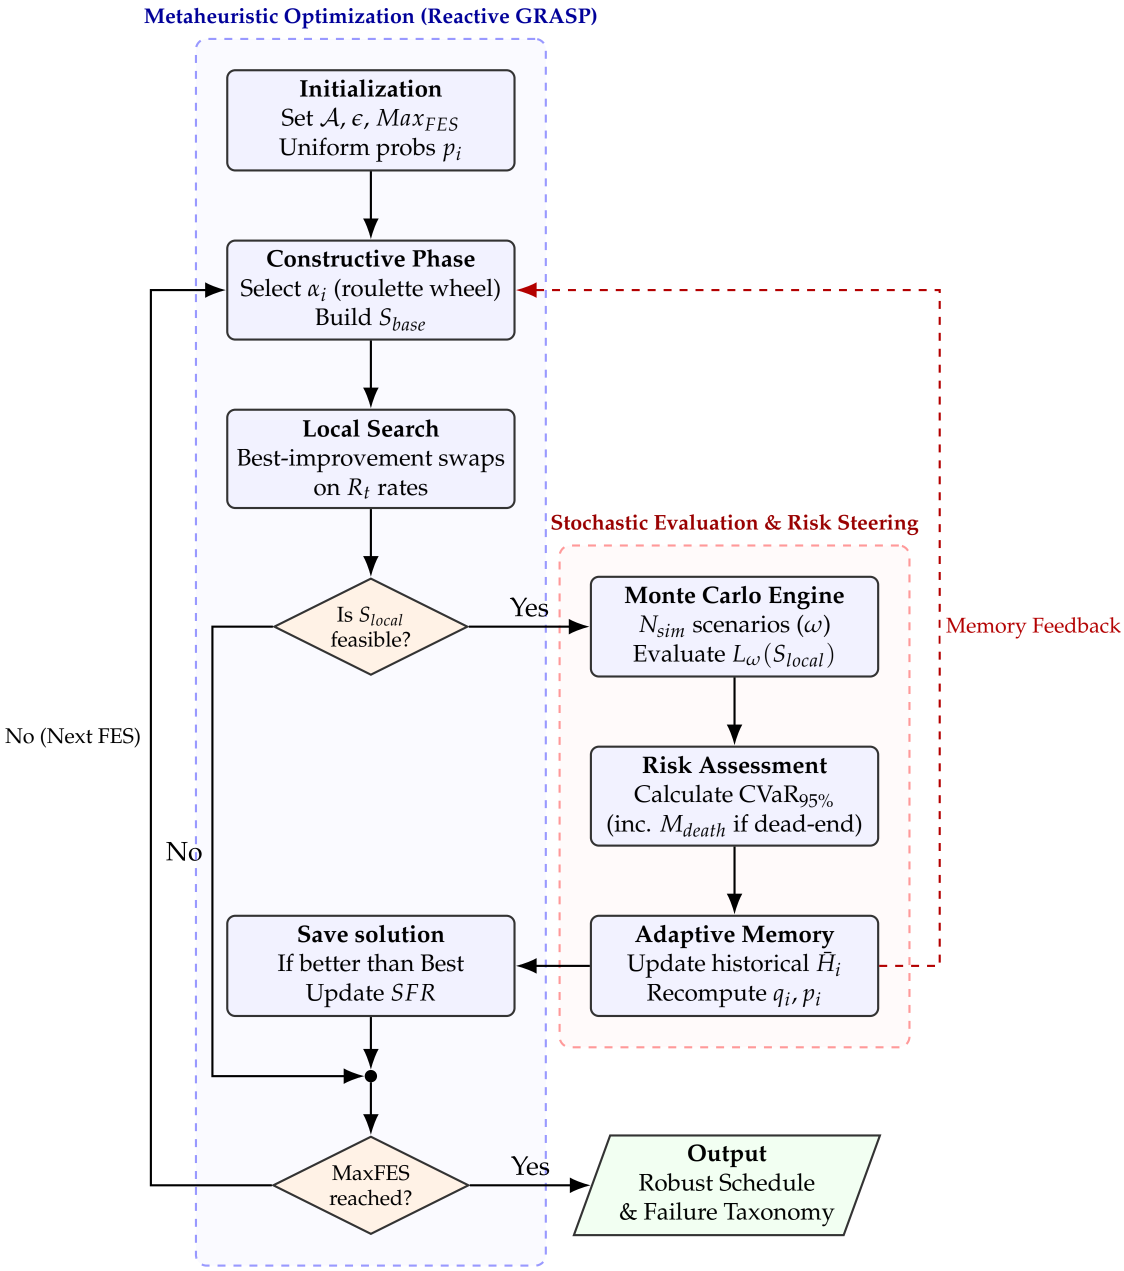


**Figure 2.** Architectural flowchart of the CVaR-Driven Reactive GRASP Simheuristic.

### 3.1. Framework Overview

The standard GRASP methodology consists of two iterative phases: a semi-greedy constructive phase and a local search phase. Our simheuristic extends this by embedding an uncertainty-evaluation module. Let $\mathcal{A} = \{\alpha_1, \alpha_2, \ldots, \alpha_m\}$ be a discrete set of relaxation parameters, where $\alpha \in [0, 1]$ controls the greediness of the constructive phase. The framework operates as follows:

1. Selection: An $\alpha$ value is selected from $\mathcal{A}$ based on a dynamically updated probability distribution $\mathcal{P} = \{p_1, \ldots, p_m\}$.
2. Construction: A baseline schedule is generated step by step. At each day, all physically feasible operational moves are evaluated, and one is selected randomly from a Restricted Candidate List (RCL) bounded by $\alpha$.
3. Local Search: If a valid baseline schedule is completed, a deterministic local search is applied to improve blending yields and minimize nominal demurrage.
4. Simulation & Evaluation: The locally optimal schedule is subjected to $N_{sim}$ Monte Carlo scenarios of stochastic vessel delays. The overall performance distribution is recorded, and the CVaR is calculated.
5. Reactive Memory Update: The CVaR score is used as feedback to update the selection probabilities $\mathcal{P}$, rewarding $\alpha$ values that consistently generate risk-averse schedules.

### *3.2. Constructive Phase and Failure Taxonomy*

The constructive phase builds a schedule $S$ day by day. At day $t$, the system identifies the set of all physically feasible moves $\mathcal{M}_t$, encompassing topping rates, blend compositions, tank-to-tank transfers, and vessel discharges. Each move $m \in \mathcal{M}_t$ is assigned a greedy score $g(m)$ balancing the economic reward of processing high-value crudes (WTI, CPC) against the immediate penalty of generated port demurrage:

$$g(m) = \left(V_m^{\mathrm{WTI}} \cdot \pi_{\mathrm{WTI}} + V_m^{\mathrm{CPC}} \cdot \pi_{\mathrm{CPC}}\right) - \left(\Delta B_m \cdot \lambda_{\mathrm{demurrage}}\right) \tag{17}$$

where $V_m$ represents the volume of specific crudes fed to the CDU, and $\pi$ denotes their respective premium economic rewards. Note that the Azeri crude acts as the baseline feedstock for the refinery's nominal margins; hence, its relative marginal reward is set to zero, and the algorithm seeks to maximize the inclusion of higher-value WTI and CPC blends. $\Delta B_m$ is the marginal increase in port backlog caused by the move (estimated heuristically during the day's construction based on the immediate discharge volume versus the waiting queue), and $\lambda_{\mathrm{demurrage}}$ is the daily demurrage unit cost.

Let $g^{max} = \max_{m \in \mathcal{M}_t} g(m)$ and $g^{min} = \min_{m \in \mathcal{M}_t} g(m)$. The Restricted Candidate List is constructed using the selected $\alpha$:

$$\mathrm{RCL}_\alpha = \{m \in \mathcal{M}_t \mid g(m) \geq g^{min} + \alpha(g^{max} - g^{min})\} \tag{18}$$

A move is then chosen uniformly at random from $\mathrm{RCL}_\alpha$.

To provide explainability, if $\mathcal{M}_t = \emptyset$ (i.e., a combinatorial dead-end is reached), the algorithm does not merely discard the incomplete schedule. Instead, it systematically logs the binding physical constraint responsible for the failure via an automated Failure Taxonomy (e.g., settling-time blockage, safety-stock violation, or port congestion deadlock), providing actionable diagnostic data for industrial planners.

### *3.3. Regenerative Local Search*

Once a complete baseline schedule $S_{base}$ is constructed, a regenerative local search explores its neighborhood $\mathcal{N}(S_{base})$ to improve the deterministic cost function $f(S)$ (formally defined as the nominal baseline evaluation of the operational loss without stochastic disruption penalties). The neighborhood is defined by swap operators acting on the topping feed rates $R_t$. For any two adjacent days $t$ and $t + 1$ where $R_t \neq R_{t+1}$, the local search evaluates the feasibility and cost impact of swapping the processing rates. While restricting the neighborhood to adjacent days represents a narrow local search space, it is an intentional design limitation chosen to keep computational overhead strictly negligible while effectively resolving high-frequency, immediate bottlenecks. Furthermore, by definition, this localized swap operation intrinsically preserves the total monthly production target (e.g., 540 kt), focusing strictly on localized intra-month stability. If the swap allows for a more profitable blending profile (recalculated via a rule-based deterministic procedure that maximizes WTI and CPC fractions under the limits defined in Section 2) without violating settling or capacity constraints, the schedule is updated. The local search employs a best-improvement termination criterion, iteratively applying swaps until no further deterministic cost reduction is achievable. To maintain a strict comparative benchmark, the computational budget ($Max_{FES}$) focuses exclusively on the highly combinatorial constructive phase; the deterministic local search evaluations are considered computationally negligible and are intentionally excluded from the FES counter.

### *3.4. Monte Carlo Simulation and Tail-Risk (CVaR) Evaluation*

The refined schedule $S_{local}$ is evaluated under uncertainty to assess its logistical robustness. We execute $N_{sim}$ independent Monte Carlo scenarios. In each scenario $\omega$, actual vessel arrival days are perturbed according to the historical probability distribution. The schedule $S_{local}$ is applied to the perturbed environment.

For each scenario $\omega$, the total operational loss $L_\omega(S_{local})$ is calculated as the sum of multiple penalized physical violations:

$$L_\omega(S_{local}) = \lambda_{\mathrm{SF}} \cdot \mathrm{SF}_\omega + \lambda_{\mathrm{SS}} \cdot \mathrm{SSV}_\omega + \lambda_{\mathrm{demurrage}} \cdot \mathrm{CumDem}_\omega - \mathrm{YieldReward}_\omega \tag{19}$$

where $\mathrm{SF}_\omega$ is the production shortfall, $\mathrm{SSV}_\omega$ is the total safety stock violation, and $\mathrm{CumDem}_\omega$ is the cumulative demurrage incurred in scenario $\omega$. The parameters $\lambda_{\mathrm{SF}}$, $\lambda_{\mathrm{SS}}$, and $\lambda_{\mathrm{demurrage}}$ are strictly positive unit penalty weights assigned to these physical violations. To ensure dimensional homogeneity across the loss function, these parameters are strictly calibrated as economic conversion factors (e.g., \$/kt), rendering the physical violation penalties perfectly commensurable with $\mathrm{YieldReward}_\omega$, which represents the total economic premium generated by the specific blending recipe. Crucially, if a scenario encounters an unrecoverable combinatorial dead-end during simulation (i.e., absolute physical infeasibility), it is assigned a massive finite penalty ($M_{death}$). Because $M_{death}$ is orders of magnitude larger than the regularizer $\epsilon$, any $\alpha_i$ generating even a single infeasible tail scenario will yield an updated quality score $q_i$ that is negligibly small relative to well-performing $\alpha$ values. This extreme penalty-based approach ensures that infeasible scenarios dominate the worst-case tail of the loss distribution, driving an aggressive CVaR steering mechanism that strictly avoids schedules with low structural resilience. In addition to economic losses, the simulation module independently tracks the Stochastic Feasibility Rate ($SFR$). This critical KPI is defined as the percentage of Monte Carlo scenarios in which the schedule is successfully executed without encountering combinatorial dead-ends or unrecoverable constraint violations.

The algorithm sorts the resulting losses $L_{(1)} \le L_{(2)} \le \cdots \le L_{(N_{sim})}$. For a given confidence level $\beta \in (0,1)$ (e.g., $\beta = 0.95$), the Value-at-Risk ($\mathrm{VaR}_\beta$) is the $\lceil \beta \cdot N_{sim} \rceil$-th worst loss. Following the canonical simheuristic paradigm [41] and the standard empirical foundations for discrete tail risk by Rockafellar and Uryasev [42,43], the Conditional Value-at-Risk ($\mathrm{CVaR}_\beta$), representing the expected loss in the worst $(1-\beta)$% of cases, is evaluated. For the scope of our implementation and subsequent computational experiments, we configure the confidence level at $\beta = 0.95$ and the Monte Carlo sampling size at $N_{sim} = 1000$. This sample size ensures an adequate representation of the tail distribution. Specifically, at a 95% confidence level, the tail encompasses exactly 50 scenarios, providing a statistically stable evaluation of the expected worst-case loss without imposing the prohibitive computational overhead associated with larger sample sizes in an iterative simheuristic framework. Note that while the classic discrete formulation often denotes the tail cardinality as $\lfloor (1-\beta) \cdot N_{sim} \rfloor$, we adopt the inclusive index bound of cardinality $(N_{sim} - \lceil \beta \cdot N_{sim} \rceil + 1)$, which is numerically equivalent to the standard formulation for all practical values of $N_{sim}$. The metric is therefore computed as:

$$\mathrm{CVaR}_\beta(S_{local}) = \frac{1}{N_{sim} - \lceil \beta \cdot N_{sim} \rceil + 1} \sum_{i=\lceil \beta \cdot N_{sim} \rceil}^{N_{sim}} L_{(i)} \tag{20}$$

*3.5. CVaR-Driven Memory Update*

In a standard Reactive GRASP, the probability $p_i$ of selecting $\alpha_i \in \mathcal{A}$ is inversely proportional to the average deterministic cost. In our risk-aware framework, we replace this with the CVaR, explicitly avoiding configurations prone to catastrophic tail events.

Let $\mathcal{H}_i$ be the historical set of all $\mathrm{CVaR}_\beta$ scores obtained using $\alpha_i$. At every newly evaluated feasible schedule (the update condition), the empirical average of the tail risk for $\alpha_i$ is calculated as $\bar{H}_i = \frac{1}{|\mathcal{H}_i|} \sum_{v \in \mathcal{H}_i} v$. Subsequently, the quality score $q_i$ is updated:

$$q_i = \left(\max(0, \bar{H}_i) + \epsilon\right)^{-1} \tag{21}$$

where $\epsilon > 0$ is a regularization constant acting as a positive shift. The formal inclusion of the $\max(0, \cdot)$ operator acts as a mathematical guard: it ensures strict positivity of the denominator and structural stability of the probability distribution even in highly lucrative scenarios where extreme yield rewards might otherwise drive the absolute operational loss $\bar{H}_i$ into negative territory. This ensures numerical stability and maintains strict positivity even if highly profitable blends render the absolute loss $v$ negative. In our computational experiments, $\epsilon$ was empirically calibrated to $10^4$, a value commensurate with the typical order of magnitude of the objective function. This calibration prevents the flattening of the probability distribution while maintaining the algorithm's sensitivity to tail-risk variations across different $\alpha$ levels. Furthermore, the overall execution is strictly bounded by a predefined computational budget ($Max_{FES}$) to guarantee objective comparisons without requiring secondary convergence stagnation criteria. The selection probabilities are then updated as:

$$p_i = \frac{q_i}{\sum_{j=1}^{m} q_j} \quad \forall i \in \{1, \ldots, m\} \tag{22}$$

This mechanism dynamically steers the search. If a highly greedy $\alpha$ generates schedules that perform well on average but fail disastrously under severe vessel delays (producing high CVaR), its selection probability $p_i$ will drastically decrease, favoring more conservative and robust $\alpha$ configurations.

The complete workflow of the proposed framework is summarized in Algorithm 1.

**Algorithm 1** CVaR-Driven Reactive GRASP Simheuristic

**Require:** Initial state $S_0$, Max function evaluations $Max_{FES}$, Alpha pool $\mathcal{A}$, confidence $\beta$, scenarios $N_{sim}$

```
Initialize p_i ← 1/|A| for all α_i ∈ A
Initialize histories H_i ← ∅ for all α_i ∈ A
BestSchedule ← null, BestCVaR ← ∞
CurrentFES ← 0
while CurrentFES < Max_FES do
    Select α* from A using roulette-wheel selection on P
    S_base, FES_used ← ConstructivePhase(S_0, α*)
    CurrentFES ← CurrentFES + FES_used
```

**Algorithm 1** *Cont.*

    **if** $S_{base}$ is feasible **then**
        $S_{local} \leftarrow$ RegenerativeLocalSearch$(S_{base})$
        $LossArray \leftarrow [\,]$
        **for** $\omega = 1$ to $N_{sim}$ **do**
            $L_\omega \leftarrow$ ApplyMonteCarloScenario$(S_{local}, \omega)$
            Append $L_\omega$ to $LossArray$
        **end for**
        $CurrentCVaR \leftarrow$ CalculateCVaR$(LossArray, \beta)$
        Append $CurrentCVaR$ to $\mathcal{H}_{\alpha^*}$
        **if** $CurrentCVaR < BestCVaR$ **then**
            $BestCVaR \leftarrow CurrentCVaR$
            $BestSchedule \leftarrow S_{local}$
        **end if**
        ▷ Reactive memory update (executed for every feasible schedule, see Section 3.5)
        Recompute quality scores $q_i$ using Equation (21)
        Update selection probabilities $p_i$ using Equation (22)
    **else**
        UpdateFailureTaxonomy$(S_{base})$ ▷ Explainability log
    **end if**
**end while**
**return** $BestSchedule$, $BestCVaR$, FailureTaxonomy

## 4. Results and Discussion

In this section, we evaluate the performance of the proposed CVaR-driven Reactive GRASP simheuristic. The algorithm was coded in Python 3.14 and executed on a workstation equipped with an Intel Core i7-3940XM CPU (3.00 GHz), 16 GB of RAM, and an NVIDIA Quadro K2000M graphics card. To ensure statistical robustness and mitigate the intrinsic randomness of the metaheuristic, the baseline optimization was performed over 25 independent runs. The selection of a 30-day time horizon ($|T| = 30$ days) for the case study is not arbitrary; rather, it aligns perfectly with the standard framework of short-term operational scheduling defined in Operations Management literature. This timeframe represents the typical industrial window within which marine refineries translate tactical procurement plans into frozen, daily executable processing operations. To strictly evaluate the out-of-sample tail risk of the locally optimal schedules, the final evaluation phase subjected the best configurations to a Monte Carlo simulation with $N_{sim} = 1000$ stochastic delay scenarios.

### 4.1. Computational Statistics and Methodological Positioning

Before detailing the operational results, it is imperative to contextualize the computational footprint of the proposed framework (Table 1). As detailed in Section 3, the methodology is a simheuristic rather than an exact mathematical programming model evaluated via branch-and-bound algorithms. Therefore, the framework does not explicitly instantiate a rigid sparse matrix of binary and continuous variables in a commercial solver. Instead, the computational complexity is structurally governed by the discrete scheduling horizon, the maximum function evaluations budget, and the stochastic simulation scenarios.

**Table 1.** Computational footprint and execution statistics of the Simheuristic framework.

| Computational Parameter | Value |
|---|---|
| Scheduling Horizon Dimension ($\lvert\mathcal{T}\rvert$) | 30 days |
| Optimization Budget ($Max_{FES}$) | 200,000 evaluations |
| Risk Evaluation Sampling ($N_{sim}$) | 1000 scenarios per optimal state |
| Average Solving Time (per independent run) | ≈165 s |

Compared to existing deterministic exact models in the literature, the primary contribution of this work does not lie in merely reducing variable or constraint counts. Instead, the advancement is methodological: by transitioning from exact deterministic programming to a CVaR-driven metaheuristic, the framework overcomes the computational intractability typical of stochastic MINLP models. This allows the algorithm to generate highly robust, risk-aware operational schedules in under three minutes, rendering it suitable for real-time reactive rescheduling in industrial settings.

### *4.2. Baseline Schedule Optimization and Explainability*

Across all 25 independent optimization runs, the proposed framework maintained a 100% success rate in identifying strictly feasible baseline schedules, demonstrating the high structural reliability of the constructive heuristic in navigating the heavily constrained solution space. Among these, the algorithm successfully identified the most resilient production schedule (Run 15), balancing high-value yield recipes with strict physical constraints. Figure 3 illustrates the baseline operational dashboard derived from this optimal Run 15 over the 30-day horizon. Panel A demonstrates the continuous satisfaction of the 540 kt monthly production target, maintaining the topping rate strictly within the 16–22 kt operational window while maximizing the inclusion of the WTI premium crude. Notably, the baseline schedule completely avoids scheduling the CPC premium crude over the 30-day horizon. While CPC yields higher economic rewards, its specific processing requirements introduce severe logistical friction and elevate the risk of port congestion. The CVaR-driven memory successfully recognized that, under nominal penalty settings, the marginal profit of the CPC mix is insufficient to justify the exponential tail-risk of terminal gridlock, intelligently opting for a structurally safer crude mix. Panel B highlights the inventory dynamics: the total inventory safely oscillates within the bounds, successfully absorbing vessel arrivals without triggering demurrage or violating the 120 kt safety stock. To ensure complete reproducibility and provide granular insights, the full day-by-day numerical dataset corresponding to Figure 3, including daily topping rates, exact blend compositions, and discrete port queue levels, has been made available as Table S1 in the Supplementary Materials.

To rigorously assess the out-of-sample robustness of the optimal baseline schedule (Run 15), the configuration was subjected to a Monte Carlo validation encompassing 1000 stochastic delay scenarios. The statistical summary and the resulting tail-risk metrics are presented in Table 2.

The necessity of the risk-aware simheuristic is starkly revealed when examining these stochastic validation metrics. The best baseline schedule yielded an Expected Operational Cost of 10,855.68. Yet, under the 95% confidence level, the $CVaR_{95\%}$ surged to 22,666.51. This divergence mathematically demonstrates that standard expected-value optimization underestimates the financial impact of low-probability tail events (such as consecutive severe vessel delays). By explicitly steering the search via CVaR, the algorithm prioritized a schedule that bounded the physical damages in the worst-case scenarios: the 95th-percentile port backlog was restricted to 86.0 kt, while safety stock violations and production shortfalls were capped at 171.0 kt and 78.0 kt, respectively.

It is crucial to contextualize the Stochastic Feasibility Rate (SFR) of 1.1% observed in Table 2. We formally define SFR as the percentage of Monte Carlo scenarios executed without incurring any operational penalty or constraint violation. Therefore, in this highly constrained stochastic environment, a low SFR does not imply algorithmic failure or an unreliable schedule; rather, it indicates that 98.9% of the simulated scenarios incurred at least minor operational friction (e.g., slight demurrage hours or marginal safety stock dips) to survive. The CVaR-driven algorithm's strength lies precisely in its ability to navigate this overwhelmingly penalized solution space, absorbing shocks by accepting calculated operational penalties rather than succumbing to unrecoverable combinatorial dead-ends.

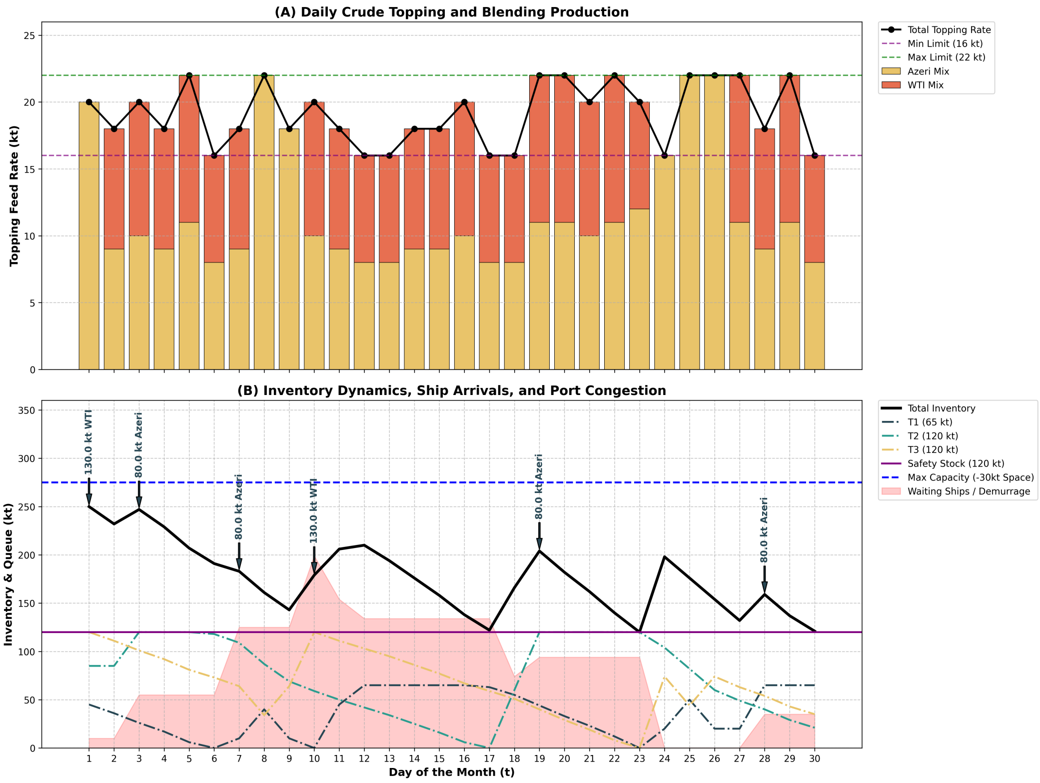


**Figure 3.** Baseline schedule execution under nominal conditions. (**A**) details the daily blending profile and topping rates. (**B**) displays tank inventory dynamics and port congestion levels.

**Table 2.** Statistical summary and P95 risk metrics derived from the Monte Carlo validation on the optimal Run 15 baseline schedule.

| Performance Metric | Validation Value |
|---|---|
| Constructive Phase Success Rate (25 runs) | 100.0% |
| Stochastic Feasibility Rate (SFR) | 1.1% |
| Expected Operational Cost | 10,855.68 |
| Conditional Value-at-Risk ($CVaR_{95\%}$) | 22,666.51 |
| P95 Port Backlog | 86.0 kt |
| P95 Safety Stock Violation | 171.0 kt |
| P95 Production Shortfall | 78.0 kt |

Furthermore, the automated Failure Taxonomy provided critical explainability insights regarding the combinatorial dead-ends encountered during the search (Table 3).

**Table 3.** Automated failure taxonomy during the optimization search.

| Failure Mode (Active Constraint) | Occurrences | Frequency (%) |
|---|---|---|
| Port Congestion (No compatible tank for discharge) | 13,092 | 80.6% |
| Combinatorial Dead-end (No feasible production move) | 2407 | 14.8% |
| Safety Stock Violation (<120 kt) | 739 | 4.6% |
| Total Discarded Branches | 16,238 | 100.0% |

Over 16,000 discarded schedule branches were analyzed. In detail, 80.6% of algorithmic failures were caused by Port Congestion (i.e., no compatible free tank space available for discharge, leading to terminal gridlock). Pure Combinatorial Dead-ends accounted for

14.8%, while Safety Stock Violations triggered only 4.6% of the structural deaths. This empirical finding aggressively challenges the standard "infinite port capacity" assumption prevalent in deterministic literature, mathematically proving that endogenous maritime logistics are the primary bottleneck in petrochemical scheduling.

### *4.3. Sensitivity Analysis and Risk-Shifting Behavior*

To validate the algorithmic responsiveness to shifting economic priorities, a sensitivity analysis was conducted by perturbing the primary penalty weights ($\lambda$) and economic rewards by $\pm$25%. To ensure rigorous statistical significance, each perturbed configuration was independently evaluated over a dedicated set of 1000 Monte Carlo validation scenarios. The comprehensive results are summarized in Table 4.

**Table 4.** Sensitivity analysis of the simheuristic under varying penalties and rewards.

| Scenario | Exp. Cost | Δ% vs. Base | CVaR (95%) | SFR (%) | P95 Backlog | P95 SS Viol. |
|---|---|---|---|---|---|---|
| Baseline | 10,855.68 | - | 22,666.51 | 1.1 | 86.0 | 171.0 |
| Demurrage Penalty (+25%) | 18,275.97 | +68.3% | 30,445.93 | 0.2 | 138.0 | 259.2 |
| Demurrage Penalty (−25%) | 14,675.71 | +35.2% | 24,827.03 | 1.6 | 221.1 | 294.0 |
| Shortfall Penalty (+25%) | 12,998.33 | +19.7% | 22,395.16 | 6.0 | 135.0 | 247.0 |
| Shortfall Penalty (−25%) | 11,085.11 | +2.1% | 23,827.96 | 2.2 | 158.0 | 250.0 |
| CPC Reward (+25%) | 13,397.75 | +23.4% | 21,929.61 | 2.8 | 135.0 | 223.0 |
| CPC Reward (−25%) | 14,186.20 | +30.7% | 22,748.60 | 1.6 | 166.0 | 260.0 |

The sensitivity results demonstrate a sophisticated, nonlinear risk-shifting behavior. Notably, when the Shortfall Penalty is increased by 25%, the algorithm strictly prioritizes meeting production targets, driving the SFR up to 6.0%. This demonstrates that risk can be reallocated across dimensions: to guarantee production, the algorithm finds inherently more stable physical configurations. Conversely, when the Demurrage Penalty is artificially reduced by 25%, the algorithm relaxes its port-clearing urgency, allowing the 95th-percentile backlog to soar from 86.0 kt to 221.1 kt as it favors internal stability.

Furthermore, given the algorithm's conservative avoidance of the highly constrained CPC crude in the baseline schedule, the CPC Reward parameter was specifically selected for perturbation to test the framework's economic elasticity. Interestingly, modifying this economic reward shifts the logistical risk profile: increasing the CPC reward by 25% tightens the expected CVaR to 21,929.61, proving that the algorithm actively renegotiates the physical constraints to process more lucrative blends only when the economic incentive outweighs the logistical friction. This confirms that the CVaR-driven memory is not statically fitting a single path, but dynamically adapting the Pareto tradeoff between maritime congestion and internal production based strictly on the user-defined economic utility function.

### *4.4. Dynamic Disruption Mitigation: Scenarios A and B*

To evaluate the proposed framework's proactive resilience and dynamic reactive capabilities, we subject the optimal robust schedule identified in Section 4 to two severe, unpredicted logistical and process disruptions occurring after the month has started. The algorithm must automatically calculate dynamic rescheduling actions starting from the "Rescheduling Day" when the disruption becomes known, adhering strictly to the operational constraints defined in Section 2. The numerical results, extracted from the detailed simulation logs, are summarized below.

#### 4.4.1. Scenario A: Severe Vessel Arrival Delay

This scenario simulates a major maritime disruption: the Azeri base crude vessel (Ship 3 expected this month), originally scheduled to arrive on Day 7 (see Figure 3), is delayed by five full days, anchoring only on Day 12. The refinery becomes aware of this delay on Day 4, which is set as the Rescheduling Day. The algorithm must manage a looming stockout of base crude while subsequent WTI and CPC ships are already en route.

Figure 4 illustrates the reformulated operational plan to mitigate the delay. The algorithm prioritized stockout avoidance over production targets. Since Ship 3 carried the base crude, the total monthly target of 540 kt could not be met. The simulation results confirm that the dynamic reschedule achieved a total processed volume of only 516.38 kt, resulting in a strict production shortfall ($SF$) of 23.62 kt.

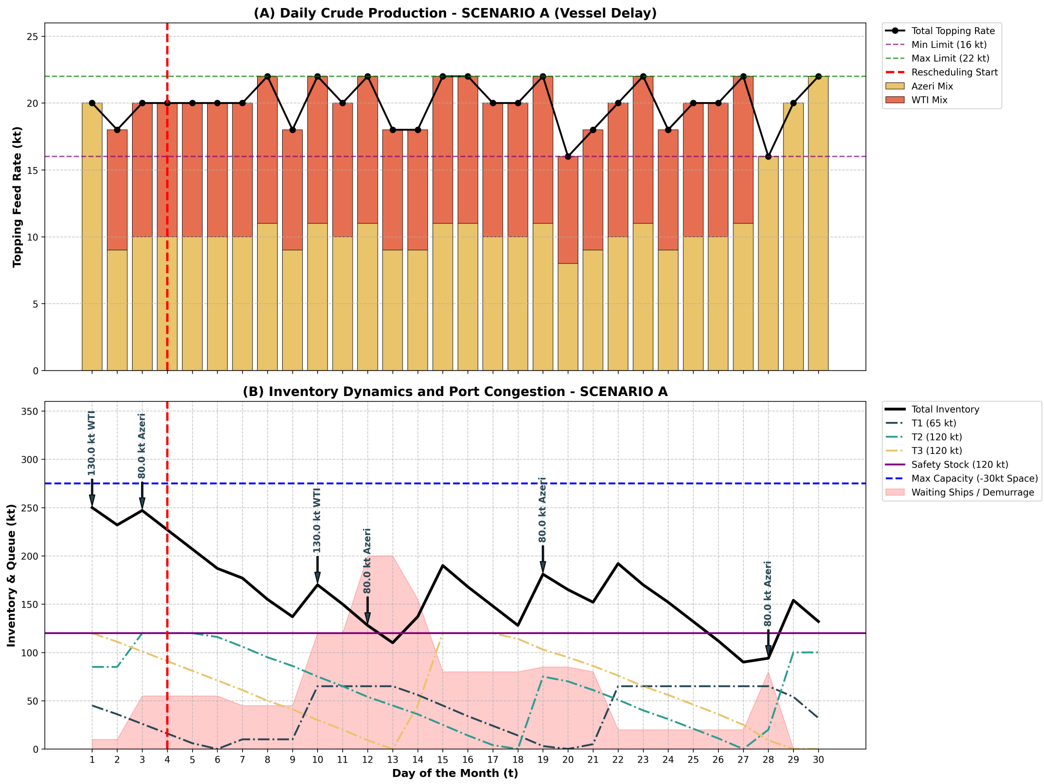


**Figure 4.** Dynamic rescheduling under Scenario A (6-day base crude vessel delay).

Crucially, as shown in the simulation logs (Panel B of Figure 4), the delay of Ship 3 to Day 8 was initially absorbed without violating the 120 kt safety stock limit, as the algorithm successfully leveraged existing internal buffers.

However, the structural unbalance caused by this major delay generated a severe cascading effect over the remainder of the month. To continuously feed the Crude Distillation Unit (CDU) and avoid a total shutdown, the algorithm was forced to draw heavily from the reserves, causing the total inventory to plummet below the 120 kt safety stock limit on Day 13 (dropping to 110 kt). A more severe and sustained violation occurred towards the end of the horizon, from Day 26 to Day 28, where the inventory reached a critical monthly minimum of 90 kt. Concurrently, the altered tank availability impacted subsequent arrivals, generating significant port backlogs that peaked at 200 kt around Days 12 and 13.

This dynamic response perfectly demonstrates the framework's ability to negotiate complex Pareto tradeoffs: it intelligently absorbed the immediate shock to prevent an early stockout cascade, strategically confining the unavoidable safety stock violations and extreme logistical friction to later periods of the scheduling horizon where they could be managed without stopping production. The complete daily numerical dataset for this dynamic rescheduling response, corresponding to Figure 4, has been made available as Table S2 in the Supplementary Materials.

#### 4.4.2. Scenario B: Unscheduled Topping Unit Maintenance

This scenario simulates an unexpected process-side disruption: a mandatory emergency shutdown of the Topping unit (CDU) for two full days. According to the detailed simulation logs, this disruption is executed on Days 19 and 20. The requirement becomes known on Day 15 (Rescheduling Day). The algorithm must immediately calculate dynamic rescheduling actions to manage incoming vessels and surging inventories while the primary process sink is disabled.

Figure 5 displays the dynamic response to the process shutdown. Contrary to a catastrophic system failure, the algorithm successfully recovered the lost processing time by dynamically boosting the daily CDU utilization to its maximum operational limit (22 kt/day) in the active days surrounding the disruption. The total processed volume reached 548 kt by the end of the month, successfully meeting the 540 kt monthly target with zero production shortfall.

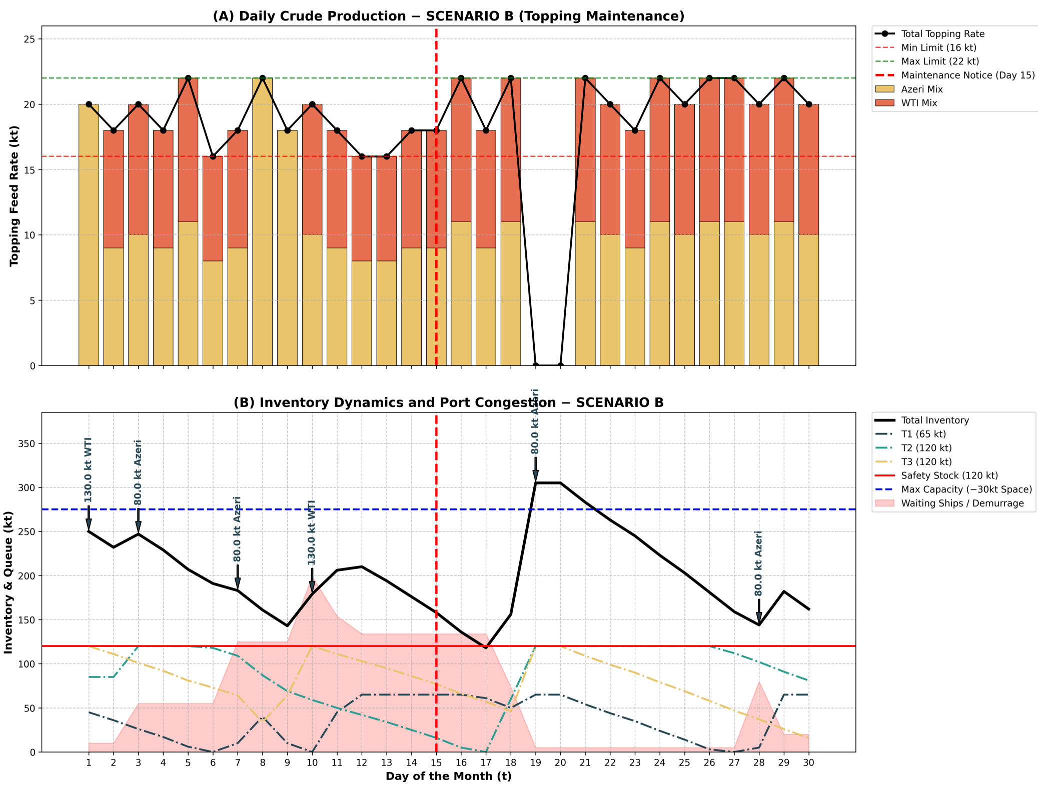


**Figure 5.** Dynamicrescheduling under Scenario B (2-day Topping unit shutdown on Days 19–20).

The analysis reveals that the primary challenge was the massive cascading logistical bottleneck caused by the disabled CDU. Knowing the shutdown was imminent, the algorithm initiated a pre-emptive tank loading strategy to clear the port. The detailed logs highlight a remarkable physical phenomenon: to empty the port queue before the CDU stopped, the algorithm discharged both the waiting backlog and a newly arrived 80 kt Azeri vessel precisely on Day 19. Consequently, on Days 19 and 20, the global tank inventory reached exactly 305 kt, representing the absolute 100% maximum physical capacity of the entire terminal ($T_1$65 kt + $T_2$120 kt + $T_3$120 kt).

Comparing the two scenarios provides a profound operational insight. In Scenario A, the pre-emptive tank clearing strategy succeeded smoothly because the CDU remained active, providing a continuous "sink" for the inventory. In contrast, Scenario B demonstrates that when the process sink itself is disabled, the terminal inherently loses its degrees of freedom. The algorithm was forced to utilize the tanks as a static maximum-capacity buffer (305 kt) to prevent vessels from accumulating infinite demurrage during the dead days. This dynamic response successfully contained the port queue, transforming a severe internal process disruption into a managed logistical shock, and ultimately protecting the global monthly production target. However, it is important to acknowledge a specific operational limitation in the current formulation regarding this sharp load increase. In real-world refining, instantly boosting the CDU to its maximum capacity after a cold shutdown could induce thermal shocks or fractionation instability. While our mathematical model strictly enforces the daily volumetric limits, it currently lacks a gradient constraint to govern the ramp-up rate between adjacent days, an aspect that warrants integration in future developments. The detailed numerical dataset logging this recovery strategy and inventory saturation, corresponding to Figure 5, is provided as Table S3 in the Supplementary Materials.

### *4.5. Ablation Study and Baseline Comparison*

To rigorously isolate and validate the algorithmic contribution of the CVaR-driven memory and the endogenous modeling of port operations, an ablation study was conducted in strict accordance with the experimental framework of the baseline model. The proposed framework was benchmarked against three algorithmic variants: a standard Reactive GRASP driven solely by the deterministic expected cost (Neutral Risk), a baseline GRASP operating without the automated Failure Taxonomy, and a simplified variant evaluating terminal operations through static exogenous delay penalties rather than endogenous, cumulative port congestion. The computational results, evaluated over 1000 Monte Carlo validation scenarios, are summarized in Table 5.

**Table 5.** Ablation study comparing the proposed framework against baseline variants under 1000 validation scenarios.

| Ablation Variant | Exp. Cost | CVaR (95%) | SFR (%) | P95 Backlog | Failure Taxonomy Status |
|---|---|---|---|---|---|
| Proposed CVaR GRASP (Baseline) | 10,855.68 | 22,666.51 | 1.1 | 86.0 kt | Active (Full Tracking) |
| GRASP with Expected Cost | 14,168.60 | 27,733.15 | 2.4 | 175.0 kt | Active (Partial Tracking) |
| GRASP without Failure Taxonomy | 12,715.89 | 25,943.18 | 1.9 | 280.0 kt | Absent (Black-box) |
| GRASP without Endogenous Congestion | 9851.98 | 18,081.76 | 6.2 | 0.0 kt | Active (Partial Tracking) |

The empirical results confirm the structural superiority of the proposed architecture. The expected-cost GRASP exhibited a severe degradation in worst-case resilience, producing a Conditional Value-at-Risk of 27,733.15 (an increase of over 22 percent compared to the baseline) and doubling the 95th-percentile port backlog to 175.0 kt. This mathematically demonstrates that standard expected-cost optimization fails to protect against catastrophic tail events and terminal gridlocks. Furthermore, the variant lacking endogenous congestion consistently underestimated the compounding nature of demurrage, yielding an artificially low CVaR of 18,081.76 and a false zero-kiloton backlog. This proves that static exogenous penalties are insufficient to capture the cascading terminal dead-ends observed in real-world severe delay scenarios. Finally, the removal of the Failure Taxonomy did not improve operational metrics but reduced the framework to a black-box optimizer. By eliminating the comprehensive diagnostic capabilities of the algorithm, this ablation severely limited its practical utility as an explainable industrial Decision Support System.

https://doi.org/10.3390/app16136733

## 5. Conclusions and Future Work

This paper addressed the highly constrained, stochastic crude oil scheduling problem in marine access petrochemical refineries. While recent exact mathematical models have begun to incorporate risk measures to protect against low-probability, high-impact tail events, such as severe vessel delays or unscheduled unit shutdowns, their application often remains computationally prohibitive for large-scale, highly combinatorial operations. Conversely, standard expected value heuristics scale well but systematically fail to proactively mitigate these extreme disruption risks. To bridge this methodological gap, we proposed a novel Reactive GRASP simheuristic driven by CVaR. By explicitly embedding the Conditional Value-at-Risk into the adaptive memory of the metaheuristic, the framework dynamically steers the search away from schedules susceptible to catastrophic logistical gridlocks, actively negotiating the Pareto tradeoff between economic yield and structural resilience.

Extensive computational experiments conducted on a realistic 30-days industrial scheduling horizon demonstrated the framework's operational superiority and transparency. The automated Failure Taxonomy yielded a critical industrial insight: 80.6 percent of all algorithmic scheduling failures were driven by terminal congestion rather than pure physical dead ends or safety stock violations. This finding empirically challenges the infinite port capacity assumption prevalent in classical optimization literature, proving that endogenous maritime logistics are the primary bottleneck in supply chain continuity. Crucially, these findings were definitively validated through a comprehensive ablation study, which mathematically demonstrated that standard expected cost optimization suffers a severe degradation in worst-case resilience (increasing CVaR by over 22 percent), and that ignoring endogenous port congestion fundamentally fails to trigger the preemptive tank clearing strategies required to survive dynamic disruptions.

Furthermore, the dynamic disruption analysis validated the algorithm's proactive resilience. Under a severe maritime delay of 6 days (Scenario A), the framework successfully absorbed the immediate shock to guarantee the zero demurrage discharge of the critical base crude, strategically confining the unavoidable safety stock violations and logistical friction to later periods of the scheduling horizon. Even more remarkably, during an unscheduled distillation unit shutdown of 2 days (Scenario B), the algorithm demonstrated advanced recovery capabilities: it dynamically boosted daily processing rates to completely eliminate production shortfalls and preemptively cleared the port queue by utilizing the terminal absolute maximum physical capacity (100 percent inventory saturation) as a static buffer during the dead days. From an industrial engineering and managerial perspective, the proposed framework acts as a critical bridge between historically siloed departments, namely maritime supply chain procurement and internal refinery operations. By functioning as an advanced DSS, it provides plant managers with a robust, explainable, and risk-aware tool capable of translating abstract logistical uncertainty into actionable, disruption-proof scheduling policies, generating optimized and highly robust schedules in less than three minutes of computational time.

As with any methodological framework, this research presents certain limitations. The algorithm has been validated on a single, albeit highly realistic, industrial case study; further benchmarking across diverse refinery topologies is necessary to fully guarantee its generalizability. Additionally, the stochastic validation assumes a known, stationary distribution for vessel delays, whereas real-world maritime logistics often exhibit nonstationary, data-driven patterns governed by macroeconomic and weather anomalies. The 30-day scheduling horizon, while standard for short-term tactical planning, truncates long-term inventory carryover effects. Furthermore, regarding process stability, while the current mathematical formulation enforces strict daily volumetric bounds, it currently lacks a continuous gradient constraint to safely govern the thermal ramp-up rate of the distillation unit immediately following a cold shutdown.

Despite these limitations, this research opens several promising avenues for future investigation. First, expanding the physical boundaries of the model downstream to include the continuous blending and scheduling of refined products would allow for a holistic, plant-wide optimization framework. This expansion must also incorporate precise thermodynamic and process stability gradients to mathematically govern unit startup and shutdown transients. Second, while the current approach effectively manages single-berth port dynamics, future research could integrate multi-berth maritime physics, incorporating complex tidal restrictions, weather windows, and heterogeneous vessel drafts. Finally, integrating Machine Learning techniques, such as utilizing Deep Reinforcement Learning to dynamically tune the simheuristic hyperparameters or training predictive models to forecast exogenous vessel arrival distributions based on real-time maritime tracking data, could further enhance the algorithm computational efficiency and predictive accuracy.

**Supplementary Materials:** The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/app16136733/s1: Table S1: Day-by-day operational scheduling dataset; Table S2: Scenario A dynamic rescheduling dataset; Table S3: Scenario B dynamic rescheduling dataset. All tables are provided as separate sheets within a single Excel file.

**Author Contributions:** Conceptualization, A.G. and G.M.; methodology, A.G. and G.M.; software, A.G. and G.M.; validation, A.G. and G.M.; formal analysis, A.G. and G.M.; investigation, A.G. and G.M.; data curation, A.G. and G.M.; writing—original draft preparation, A.G. and G.M.; writing—review and editing, A.G. and G.M.; visualization, A.G. and G.M.; supervision, A.G. and G.M. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The numerical data supporting the findings of this study are available from the corresponding author upon reasonable request.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Nomenclature

| | |
|---|---|
| $t \in \mathcal{T}$ | Discrete time periods (days) |
| $c \in \mathcal{C}$ | Set of crude oil types (Azeri, WTI, CPC) |
| $k \in \mathcal{K}$ | Set of storage tanks ($T_1, T_2, T_3$) |
| $v \in \mathcal{V}$ | Set of arriving vessels |
| $m \in \mathcal{M}_t$ | Set of feasible moves at day $t$ |
| $Cap_k$ | Maximum volumetric capacity of tank $k$ [kt] |
| $SS$ | Global Safety Stock limit [kt] |
| $SF$ | Global Safety Space limit [kt] |
| $R^{min}, R^{max}$ | Minimum and maximum topping feed rates [kt/day] |
| $\tau$ | Mandatory brine-settling time [days] |
| $\pi_c$ | Premium economic reward for crude $c$ |
| $\lambda$ | Unit penalty weights for physical violations |
| $\epsilon$ | Positive shift regularization constant |
| $M$ | Big-M constant for canonical MIP formulation |
| $M_{death}$ | Massive finite penalty for unrecoverable infeasibility |
| $Max_{FES}$ | Maximum computational budget (Function Evaluations) |
| $FES$ | Current function evaluation counter |

| | |
|---|---|
| $I_{k,c,t}$ | Inventory of crude $c$ in tank $k$ at end of day $t$ [kt] |
| $U_{v,k,c,t}$ | Volume unloaded from vessel $v$ to tank $k$ [kt] |
| $X_{k',k,c,t}$ | Internal transfer volume from tank $k'$ to tank $k$ [kt] |
| $F_{k,c,t}$ | Volume fed to the CDU from tank $k$ [kt] |
| $R_t$ | Total topping processing rate at day $t$ [kt/day] |
| $Y_{k,c,t}$ | Binary indicator: 1 if tank $k$ stores crude $c$ |
| $S_{k,t}$ | Consecutive days tank $k$ has been settling |
| $Q_{v,c,t}$ | Residual volume on waiting vessel $v$ [kt] |
| $B_t$ | Total port backlog (congestion) at day $t$ [kt] |
| $\alpha$ | Greediness parameter for the Reactive GRASP |
| $p_i, q_i$ | Selection probability and quality score for $\alpha_i$ |
| $SFR$ | Stochastic Feasibility Rate (%) |

## References


1. Assis, L.S.; Camponogara, E.; Menezes, B.C.; Grossmann, I.E. An MINLP formulation for integrating the operational management of crude oil supply. *Comput. Chem. Eng.* **2019**, *123*, 110–125. [CrossRef]
2. Zheng, W.; Gao, X.; Huang, F.; Zuo, X.; Chen, X. Integrated optimization of crude oil procurement planning and blending scheduling for property stabilization. *Comput. Chem. Eng.* **2024**, *186*, 108716. [CrossRef]
3. AlMajed, A.M.; Syed, M.N.; Malik, A.I. Multi-Objective Mathematical Model for Crude Oil Terminal Scheduling with Tank Farm Operations. *Mathematics* **2025**, *13*, 3817. [CrossRef]
4. Fuller, D.B.; Ferreira Filho, V.J.M.; de Arruda, E.F. Oil industry value chain simulation with learning agents. *Comput. Chem. Eng.* **2018**, *111*, 199–209. [CrossRef]
5. Dai, X.; Zhao, L.; He, R.; Du, W.; Zhong, W.; Li, Z.; Qian, F. Data-driven Wasserstein distributionally robust chance-constrained optimization for crude oil scheduling under uncertainty. *Chin. J. Chem. Eng.* **2024**, *69*, 152–166. [CrossRef]
6. Oliveira, F.; Nunes, P.M.; Blajberg, R.; Hamacher, S. A framework for crude oil scheduling in an integrated terminal-refinery system under supply uncertainty. *Eur. J. Oper. Res.* **2016**, *252*, 635–645. [CrossRef]
7. Kong, L.; Xiao, H.; Wang, Y.; Chen, K.; Liu, M. Coordinated Allocation of Channel-Tugboat-Berth Resources Under Tidal Constraints at Liquid Terminal. *Appl. Sci.* **2025**, *15*, 13263. [CrossRef]
8. Cerdá, J.; Pautasso, P.C.; Cafaro, D.C. Efficient Approach for Scheduling Crude Oil Operations in Marine-Access Refineries. *Ind. Eng. Chem. Res.* **2015**, *54*, 8219–8238. [CrossRef]
9. Stanzani, A.d.L.; Pureza, V.; Morabito, R.; da Silva, B.J.V.; Yamashita, D.; Ribas, P.C. Optimizing multiship routing and scheduling with constraints on inventory levels in a Brazilian oil company. *Int. Trans. Oper. Res.* **2018**, *25*, 1163–1198. [CrossRef]
10. Zhang, H.; Liang, Y.; Liao, Q.; Gao, J.; Yan, X.; Zhang, W. Mixed-time mixed-integer linear programming for optimal detailed scheduling of a crude oil port depot. *Chem. Eng. Res. Des.* **2018**, *137*, 434–451. [CrossRef]
11. Panda, D.; Ramteke, M. Reactive scheduling of crude oil using structure adapted genetic algorithm under multiple uncertainties. *Comput. Chem. Eng.* **2018**, *116*, 333–351. [CrossRef]
12. He, W.; Zhao, J.; Zhao, L.; Li, Z.; Yang, M.; Liu, T. Data-driven two-stage distributionally robust optimization for refinery planning under uncertainty. *Chem. Eng. Sci.* **2023**, *269*, 118466. [CrossRef]
13. Long, J.; Wang, N.; Zhai, J.; Liang, C.; Jiang, S.; Zhao, L. Data driven multi-objective economic-environmental robust optimization for refinery planning with multiple modes under uncertainty. *Comput. Ind. Eng.* **2024**, *198*, 110697. [CrossRef]
14. Floudas, C.A.; Lin, X. Continuous-time versus discrete-time approaches for scheduling of chemical processes: A review. *Comput. Chem. Eng.* **2004**, *28*, 2109–2129. [CrossRef]
15. Zhang, S.; Wu, N.; Li, Z.; Qu, T.; Li, C. Petri net-based approach to short-term scheduling of crude oil operations with less tank requirement. *Inf. Sci.* **2017**, *417*, 247–261. [CrossRef]
16. Reddy, P.C.P.; Karimi, I.A.; Srinivasan, R. A new continuous-time formulation for scheduling crude oil operations. *Chem. Eng. Sci.* **2004**, *59*, 1325–1341. [CrossRef]
17. Pan, M.; Li, X.; Qian, Y. New approach for scheduling crude oil operations. *Chem. Eng. Sci.* **2009**, *64*, 965–983. [CrossRef]
18. Yadav, S.; Shaik, M.A. Short-Term Scheduling of Refinery Crude Oil Operations. *Ind. Eng. Chem. Res.* **2012**, *51*, 9287–9299. [CrossRef]
19. Tao, Z.; Li, H.; Qiu, T. A hybrid framework of intrinsic constraint handling and safe reinforcement learning for crude oil scheduling. *Comput. Chem. Eng.* **2026**, *205*, 109480. [CrossRef]
20. Shah, N. Mathematical programming techniques for crude oil scheduling. *Comput. Chem. Eng.* **1996**, *20*, S1227–S1232. [CrossRef]
21. Moro, L.F.L.; Pinto, J.M. Mixed-Integer Programming Approach for Short-Term Crude Oil Scheduling. *Ind. Eng. Chem. Res.* **2004**, *43*, 85–94. [CrossRef]

22. Li, J.; Li, W.; Karimi, I.A.; Srinivasan, R. Improving the robustness and efficiency of crude scheduling algorithms. *AIChE J.* **2007**, *53*, 2659–2680. [CrossRef]
23. Karuppiah, R.; Furman, K.C.; Grossmann, I.E. Global optimization for scheduling refinery crude oil operations. *Comput. Chem. Eng.* **2008**, *32*, 2745–2766. [CrossRef]
24. Mouret, S.; Grossmann, I.E.; Pestiaux, P. A new Lagrangian decomposition approach applied to the integration of refinery planning and crude-oil scheduling. *Comput. Chem. Eng.* **2011**, *35*, 2750–2766. [CrossRef]
25. Oddsdottir, T.A.; Grunow, M.; Akkerman, R. Procurement planning in oil refining industries considering blending operations. *Comput. Chem. Eng.* **2013**, *58*, 1–13. [CrossRef]
26. Clapper, Y.; Berkhout, J.; Bekker, R. Adaptive budget allocation in simheuristics applied to stochastic home healthcare routing and scheduling. *Comput. Ind. Eng.* **2024**, *198*, 110651. [CrossRef]
27. Xu, J.; Zhang, S.; Zhang, J.; Wang, S.; Xu, Q. Simultaneous scheduling of front-end crude transfer and refinery processing. *Comput. Chem. Eng.* **2017**, *96*, 212–236. [CrossRef]
28. He, R.; Xie, Y.; Zhang, S.; Xu, F.; Long, J. Knowledge-assisted hybrid optimization strategy of large-scale crude oil scheduling integrated production planning. *Comput. Chem. Eng.* **2025**, *192*, 108904. [CrossRef]
29. Ramteke, M.; Srinivasan, R. Large-Scale Refinery Crude Oil Scheduling by Integrating Graph Representation and Genetic Algorithm. *Ind. Eng. Chem. Res.* **2012**, *51*, 5256–5272. [CrossRef]
30. Hou, Y.; Wu, N.; Li, Z.; Zhang, Y.; Qu, T.; Zhu, Q. Many-objective optimization for scheduling of crude oil operations based on NSGA-III with consideration of energy efficiency. *Swarm Evol. Comput.* **2020**, *57*, 100714. [CrossRef]
31. Li, J.; Misener, R.; Floudas, C.A. Scheduling of crude oil operations under demand uncertainty: A robust optimization framework coupled with global optimization. *AIChE J.* **2012**, *58*, 2373–2396. [CrossRef]
32. Dai, X.; Zhao, L.; Li, Z.; Du, W.; Zhong, W.; He, R.; Qian, F. A data-driven approach for crude oil scheduling optimization under product yield uncertainty. *Chem. Eng. Sci.* **2021**, *246*, 116971. [CrossRef]
33. Yalcin, D.; Sildir, H. Robust scheduling of crude oil farming and processing under uncertainty. *Chem. Eng. Res. Des.* **2026**, *227*, 354–373. [CrossRef]
34. Panda, D.; Ramteke, M. Preventive crude oil scheduling under demand uncertainty using structure adapted genetic algorithm. *Appl. Energy* **2019**, *235*, 68–82. [CrossRef]
35. Garcia-Verdier, T.G.; Gutierrez, G.; Méndez, C.A.; Palacín, C.G.; de Prada, C. Optimization of crude oil operations scheduling by applying a two-stage stochastic programming approach with risk management. *J. Process Control* **2024**, *133*, 103142. [CrossRef]
36. Liu, B.; Li, Z.-C.; Wang, Y. A branch-and-price heuristic algorithm for the bunkering operation problem of a liquefied natural gas bunkering station in the inland waterways. *Transp. Res. Part B Methodol.* **2023**, *167*, 145–170. [CrossRef]
37. Wu, N.; Chu, C.; Chu, F.; Zhou, M. Schedulability Analysis of Short-Term Scheduling for Crude Oil Operations in Refinery With Oil Residency Time and Charging-Tank-Switch-Overlap Constraints. *IEEE Trans. Autom. Sci. Eng.* **2011**, *8*, 190–204. [CrossRef]
38. Wu, N.; Bai, L.; Zhou, M. An Efficient Scheduling Method for Crude Oil Operations in Refinery With Crude Oil Type Mixing Requirements. *IEEE Trans. Syst. Man Cybern. Syst.* **2016**, *46*, 413–426. [CrossRef]
39. Feo, T.A.; Resende, M.G.C. Greedy Randomized Adaptive Search Procedures. *J. Glob. Optim.* **1995**, *6*, 109–133. [CrossRef]
40. Prais, M.; Ribeiro, C.C. Reactive GRASP: An application to a matrix decomposition problem in TDMA traffic assignment. *Inf. J. Comput.* **2000**, *12*, 164–176. [CrossRef]
41. Juan, A.A.; Faulin, J.; Grasman, S.E.; Rabe, M.; Figueira, G. A review of simheuristics: Extending metaheuristics to deal with stochastic combinatorial optimization problems. *Oper. Res. Perspect.* **2015**, *2*, 62–72. [CrossRef]
42. Rockafellar, R.T.; Uryasev, S. Optimization of conditional value-at-risk. *J. Risk* **2000**, *2*, 21–42. [CrossRef]
43. Rockafellar, R.T.; Uryasev, S. Conditional value-at-risk for general loss distributions. *J. Bank. Financ.* **2002**, *26*, 1443–1471. [CrossRef]