

Optimal Reactive Operation of General Topology Supply Chain and Manufacturing Networks under Disruptions

November 26, 2024

Daniel Ovalle^a, Joshua L. Pulsipher^b, Yixin Ye^{c,1}, Kyle Harshbarger^{d,2}, Scott Bury^e, Carl D. Laird^a, and Ignacio E. Grossmann^{a,3}

^a Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA

^b Department of Chemical Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada

^c The Dow Chemical Company, Core R&D, Lake Jackson, TX

^d The Dow Chemical Company, Supply Chain Innovation, Midland, MI

^e The Dow Chemical Company, Core R&D, Midland, MI

Abstract

Supply and manufacturing networks in the chemical industry involve diverse processing steps across different locations, rendering their operation vulnerable to disruptions from unplanned events. Optimal responses should consider factors such as product allocation, delayed shipments, and price renegotiation, among other factors. In such context, we propose a multiperiod mixed-integer linear programming model that integrates production, scheduling, shipping, and order management to minimize the financial impact of such disruptions. The model accommodates arbitrary supply chain topologies and incorporates various disruption scenarios, offering adaptability to real-world complexities. A case study from the chemical industry demonstrates the scalability of the model under finer time discretization and explores the influence of disruption types and order management costs on optimal schedules. This approach provides a tractable, adaptable framework for developing responsive operational plans in supply chain and manufacturing networks under uncertainty.

1 Introduction

Real-world events highlight the importance of effective supply chain disruption management, as exemplified by the COVID-19 pandemic, which introduced unprecedented challenges including facility shutdowns, labor shortages, and transportation constraints [1, 2]. More specifically, supply chain networks of high value-added chemicals are particularly vulnerable to disruption. These products often require several raw materials and multiple manufacturing steps in series and/or parallel at various geographical locations. Hence, supply chain networks from the chemical industry often deviate from traditional multi-echelon network topologies where material flows in a feed-forward manner [3]. More specifically, these networks have arbitrary network topology, meaning that no simplifying assumptions are made about the structure or connectivity of its components [4]. As shown in Figure 1, in a traditional multi-echelon supply chain, material is shipped and transformed sequentially, and the permissible connections among node types are defined. In contrast, in arbitrary topologies the notion of an echelon becomes obscured as material flow is no longer unidirectional

¹Currently at DoorDash Inc.

²Currently at Huntsman Corporation.

³Corresponding author. Email: grossmann@cmu.edu

and can take various paths, including looping through multiple facilities. While material generally flows “downstream,” it may be reused at multiple stages, particularly in scenarios like sustaining a reaction. Moreover, plant and warehouse nodes are allowed to have both internal and external demand introducing a significant level of complexity that is not commonly addressed in the supply chain literature [3]. This is given by the fact that in these complex supply and manufacturing networks the demand cannot be easily aggregated or centralized. Furthermore, the demand for such products is significantly smaller and more volatile than commodity products. Accordingly, the profit margins are higher, which gives rise to an even higher loss rate per disruption when compared to supply chains from other industries. All of these aspects increase both the challenges and the importance of managing disruptions in high-value chemical supply chains.




Figure 1: Comparison of a traditional feed forward multi-echelon topology and a potential arbitrary topology with similar number of components.

From the perspective of the customer, supply chain disruptions and interruptions can lead to delayed deliveries, increases in prices, or impacts on the product quality. From the perspective of the company, they can lead to significant economic losses, including increased costs to rush shipment to meet customer orders, penalties from delayed deliveries, or even loss of revenue due to failure in fulfilling existing orders. Furthermore, customer satisfaction levels can be significantly impacted when an order is delivered late or canceled. Therefore, effective optimization of decisions along the supply chain is needed to mitigate the economic and operational impacts of disruptions and interruptions. Some examples of these decisions include rerouting current inventories, increasing production rate at non-disrupted locations, buying intermediate or finished products from competitors or third parties, and managing existing order delivery dates. In practice, response decisions are often made locally and empirically leading to sub-optimal strategies. Furthermore, cascading impacts of decentralized decisions can lead to poor system economics, and even situations heuristic policies are not possible to implement.

To address the problem of supply chains under disruptions, a significant amount of literature focuses on supply chain resilience which refers to the ability to adapt to unexpected events and quickly recover the desired level of operations [5, 6]. Researchers have introduced metrics to measure resilience, considering aspects like network information [7, 8], the handled products [9], and overall performance [10, 11]. For a more detailed exploration of supply chain resilience, interested readers can refer to the comprehensive reviews by Katsaliaki et al. [12] and Ribeiro and Barbosa-Povoa [13].

In the area of supply chain resilience, significant attention is given to the facility location problem, which deals with strategic planning for the network design while considering uncertainty associated with demand and disruptions [14, 15, 16]. This problem balances the trade-off between centralizing the demand on specific nodes to overcome individual site variability [17, 18] with the implicit vulnerability that arises at the aggregated critical nodes [15]. To address this problem, authors have suggested various strategies, including considering site-dependent distribution probabilities [19, 20], using approximation algorithms [21, 22], implementing tailored cuts and bounds [23], and employing stochastic programming approaches using sample average approximation [24, 25]. For a more comprehensive review of robust supply chain design under uncertainty, readers can refer to Klibi et al. [26]. While the aforementioned approaches effectively address resilient supply chain network design, gaps remain in the literature regarding short-term operational planning

functioning within the restrictions of the physical supply chain design. There is a need to understand (and optimize) how the operation of the supply chain system can effectively react to given disruptions within a shorter time frame (e.g., daily or hourly) instead of solving a long-term network design problem.

In the operational supply chain literature from the Operations Research (OR) community, several authors derive inventory policies aimed at mitigating the impacts of uncertainty and disruptions, as highlighted by Glasserman and Tayur [27]. Utilizing knowledge about network topology proves beneficial in designing and operating supply chains under disruptions. When the network exhibits a specific multi-echelon structure, assumptions can be made about its behavior, allowing for development of tailored algorithms or heuristics based on graph and queuing theory [3, 28]. Various heuristics and exact algorithms have been developed to determine optimal replenishment and allocation policies under uncertainty for spanning serial networks [29, 30, 31], single supplier three-component networks [32], assembling networks [33], multi-location networks [34], and distribution networks [35, 36]. Although these approaches can be optimal and computationally inexpensive, they are not applicable to the case of interest due to two main reasons. First, supply chains in the chemical industry frequently have arbitrary topologies, and the assumptions made when developing these methods are not applicable. Second, these approaches often yield fixed policies rather than the optimal dynamic response to recover from a specific disruption, leaving a gap in the literature that still needs to be addressed. Furthermore, traditional OR rely on distributional assumptions about the demand during lead time which might not always hold in practice [37, 38].

Chemical industry supply chains have been widely addressed by the Process Systems Engineering community [39, 40] as they play a key role in efforts towards enterprise-wide optimization [41]. While the majority of contributions from this community focus on network design, several authors have addressed the operational dynamic response to disruptions as described in the review by Garcia and You [42]. The operational response was first addressed in Perea-Lopez et al. [43], studying multiperiod supply chain reactions and exploring how closed-loop control techniques can mitigate the bullwhip effect in a multi-echelon supply chain. More recently, Badejo and Ierapetritou [44] proposed a multiperiod mathematical model coupled with a rolling horizon approach to tackle supply chain disruptions. However, there are several key features that this work did not address. First, the coarse discretization is not able to capture more detailed dynamics of the supply chain system. Moreover, this work does not consider important order management decisions, such as late delivery and order cancellation in their response portfolio.

In summary, there is a need for models capable of generating an optimal dynamic response for operating a supply chain network with an arbitrary topology following a disruption. This response should include optimal schedules in a actionable time-discretization for production, material acquisition, shipment, and inventory levels. It should also address order management decisions, covering both delivery dates and cancellations of existing orders. To bridge the aforementioned gaps, in this paper we develop a multiperiod mixed-integer linear programming (MILP) model that determines an optimal disruption-reactive schedule, throughout a specified operation horizon, for a multi-material supply chain and manufacturing network of arbitrary topology. This model considers actionable time discretizations that can capture shipments with different transportation durations. Furthermore, the model is able to account for late deliveries that may accumulate after a disruption, and it is able to provide a schedule on how to optimally deliver those backorders. The model also includes binary variables for order cancellation, used when demands cannot be met post-disruption and late delivery is not cost-effective. This strategic cancellation approach allows scarce resources to be reallocated to fulfill critical or high-value orders and enables supply chain managers to notify customers in advance, facilitating opportunities for renegotiation or rescheduling. To achieve this objective, we use Generalized Disjunctive Programming (GDP) to express discrete decisions at a higher level. The resulting disjunctions are then reformulated directly using MILP constraints. For further background on GDP, readers are referred to Raman and Grossmann [45], and Grossmann and Trespalcios [46].

The remainder of the paper is as follows: Section 2 defines the problem, outlining assumptions, decisions, and primary objectives. In Section 3, we first introduce the base MILP model and later relax some assumptions to capture more complex supply chain dynamics. Section 4 defines various disruption types and demonstrates their natural integration into the model. Section 5 presents a case study inspired by The Dow Chemical Company’s silicone rubber business. Finally, Section 6 discusses computational results, and Section 7 provides concluding remarks and suggests potential future research directions.

2 Problem Statement

In this work we consider a multi-material supply chain and manufacturing network where a disruption has occurred. The network has an arbitrary topology consisting of supplier, plant, warehouse, and customer nodes (see Figure 1). We allow for multiple transportation modes, which are represented as different connections between the same nodes. An operational time horizon is defined during which scheduled customer orders are known, and the reactive operation occurs. The shipment times and costs are known for every arc. Production coefficients, indicating the relationship between plant-produced recipes and material quantities, are given. Capacities for production, shipment, and inventory are specified. The material selling price and costs related to material acquisition, production, and inventory holding are also known.

There are several operational variables that need to be determined in the response to disruption. Optimal schedules for shipment amounts, routes, plant production, and acquisition need to be constructed for each material at every node. By determining these variables alongside existing stock, inventory levels are specified throughout the operation. Furthermore, given that once a disruption happens it might not always be feasible to deliver orders according to the original schedule, order management decisions need to be made. This means determining which orders are fulfilled on time, which will be delayed (and by how much), and which are to be canceled.

The primary goal of this problem is to derive an optimal operational response to the disruption, that aims to maximize the overall profit of the operation throughout the entire horizon while still accounting for customer satisfaction. At the end of the time horizon, the obtained response should induce recovery by bringing the system as close as possible to its pre-disruption state. This operational strategy involves balancing the existing trade-off between maintaining high customer satisfaction, characterized by delivering orders on time and avoiding cancellations, and obtaining a cost-effective operation during and after the disruption.

Certain assumptions are made in the problem formulation. The production is assumed to be continuous, thus rigorous batch scheduling models such as state-task [47] and resource-task network models [48] are not considered. Moreover, orders from the same customer and involving the same material within a given time period, are aggregated and treated as a singular entity. Additionally, transportation is considered continuous, nevertheless, any discrete unavailability of transportation resources can be effectively modeled as a disruption within the framework.

3 Proposed Base Model

In this section, we present the multiperiod mixed-integer linear programming (MILP) formulation for the optimal integrated response given a disruption. This model considers a set of different materials \mathcal{M} , as well as suppliers \mathcal{S} , plants \mathcal{P} , warehouses \mathcal{W} , and customers \mathcal{C} that are connected through a set of arcs \mathcal{A} that operate throughout a discrete time horizon \mathcal{T} .

The base model operates under several further assumptions. The model assumes transportation costs are linear with respect to the amount moved and have no minimum shipping amount. Furthermore, fixed costs associated with transportation are ignored. The base model does not consider service level agreements with suppliers, meaning there is no mandated minimum quantity for procurement from suppliers. To facilitate recovery post-disruption, the base model aligns end-of-operation inventories with their pre-disruption levels, i.e., the beginning of the operation. Finally, the model imposes a strict lower bound on the buffer stock value, preventing inventories from falling below this threshold. All of these assumptions are relaxed in Section 3.7 where different extensions to this base model are discussed. A comprehensive summary of the nomenclature for all the sets, parameters and variables can be found in the Appendix.

3.1 Material-Node and Material-Arc Set Encoding

Flexibility in managing various materials is crucial in supply chain optimization with arbitrary network topologies. Conventional distinctions between raw materials, intermediate products, and final products can become ambiguous, as the same material may serve different functions based on its location within the supply chain. For example, consider a supply chain network with a manufacturing facility, Plant 1, as depicted in Figure 2.

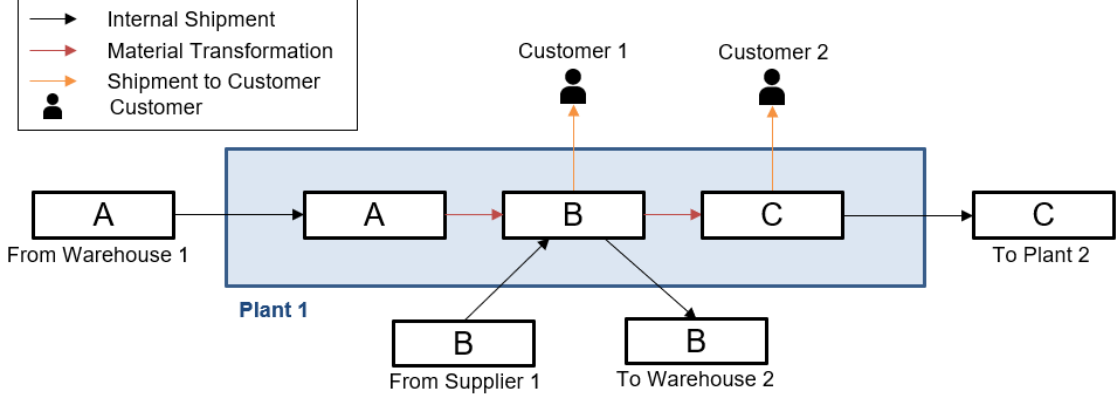


Figure 2: Illustrative example of material transformation at Plant 1.

In this plant, material B can be viewed either as a raw material purchased from Supplier 1, or as an intermediate product from the internal reaction. Furthermore, it can be considered as a product given external demand from Customer 1. Figure 2 emphasizes the need of adopting a flexible material handling approach, avoiding rigid classifications.

In response to this, the proposed model employs a single aggregated set \mathcal{M} containing all materials. This allows for material tracking irrespective of its specific location or role within the network. Moreover, to avoid creation of unnecessary variables (e.g., materials that are not sold by a supplier), we define subsets of materials existing at a given node n as $\mathcal{M}_n \forall n \in \mathcal{N}$. Similarly, not every material has to traverse every arc, leading to the definition of the analogous subset $\mathcal{M}_a \forall a \in \mathcal{A}$, as all the species capable of traveling through arc a . Also, we incorporate the transportation method directly into the arc definition, where each arc is characterized by origin, destination, material, and transportation method. This approach, together with the set of materials \mathcal{M}_a , provides the modeler with flexibility, allowing the designation of materials that can exclusively traverse a route when shipped via a particular transportation mode.

3.2 Balance Constraints

The parameter τ_{mat} represents the shipping time (duration) for material m entering arc a at time period t . Also, F_{mat}^{In} represents the amount of material m entering arc a during a given time period t , while F_{mat}^{Out} represents the same flow emerging from the arc after τ_{mat} time periods as:

$$F_{mat}^{In} = F_{mat\{t+\tau_{mat}\}}^{Out} \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} : t + \tau_{mat} \leq |\mathcal{T}|. \quad (1)$$

This also allows us to model transportation delays as disruptions within the framework. Note also that extra constraints are included to disable flow variables that will lead to times outside the time horizon:

$$F_{mat}^{In} = 0 \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} : t + \tau_{mat} > |\mathcal{T}| \quad (2a)$$

$$F_{mat}^{Out} = 0 \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} : t - \tau_{mat} < 0 \quad (2b)$$

The warehouse inventory (I) balance is shown below where F^{Out} variables include flows going out of an upstream arc into the inventory, and F^{In} refer to flows out of the inventory into a downstream arc.

$$I_{mw} = I_{mw\{t-1\}} + \sum_{a \in \mathcal{A}_w^{In}} F_{mat}^{Out} - \sum_{a \in \mathcal{A}_w^{Out}} F_{mat}^{In} \quad \forall w \in \mathcal{W}, m \in \mathcal{M}_w, t \in \mathcal{T} \setminus \{0\}. \quad (3)$$

Plants are allowed to store inventory and utilize existing stock in their reactive response. This flexibility allows for policies, such as producing with available inventory to alleviate upstream shortages, or accumulating production in the early stages of a reaction for later depletion. Consequently, an inventory balance similar to Equation (3) is used along with an additional consumption/generation term to account for production as:

$$I_{mpt} = I_{mp\{t-1\}} + \sum_{a \in \mathcal{A}_p^{In}} F_{mat}^{Out} - \sum_{a \in \mathcal{A}_p^{Out}} F_{mat}^{In} + \sum_{r \in \mathcal{R}_p} \phi_{rm} P_{prt} \quad \forall p \in \mathcal{P}, m \in \mathcal{M}_p, t \in \mathcal{T} \setminus \{0\}, \quad (4)$$

where production balance (P) is modeled using a recipe-based approach, and the set \mathcal{R}_p denotes the possible recipes at plant p . In this method, P represents the amount of a recipe produced. To map from recipe to material, the mass-based bill of materials ϕ is employed, indicating whether a material is consumed (-) or produced (+). This method allows for the modeling of different recipes by tracking the net production or consumption of each material. In the base model production is assumed to occur instantaneously per time period, however, an extension to time-dependent production is discussed in Section 3.7.1.

Inputs and outputs are calculated in the network in terms the flows leaving supplier nodes and the flows entering demand nodes, respectively. This allows every plant and warehouse to either receive external supply, or fulfill external demand, as long as the corresponding arc to the supplier or customer exists. The satisfied demand (D) is then computed considering all arcs entering a customer c (A_c^{In}) through the equation:

$$D_{mct} = \sum_{a \in \mathcal{A}_c^{In}} F_{mat}^{Out} \quad \forall c \in \mathcal{C}, m \in \mathcal{M}_c, t \in \mathcal{T}. \quad (5)$$

Here, D_{mct} determines the customer demand amount that is satisfied at time t , and may or may not be the same as the desired amount ordered. This is further discussed in Section 3.4. Analogously, the amount to buy (B) is computed by accounting for all arcs leaving supplier s (A_s^{Out}) as indicated below:

$$B_{mst} = \sum_{a \in \mathcal{A}_s^{Out}} F_{mat}^{In} \quad \forall s \in \mathcal{S}, m \in \mathcal{M}_s, t \in \mathcal{T}. \quad (6)$$

3.3 Determining Initial and Final States

The proposed formulation is a reactive model that plans supply chain operations once disruption information becomes available. Hence, once there is awareness of a disruption, the initial state of the system, characterized by the inventory levels in both plants and warehouses at the time the unplanned event takes place (I^0), is fixed as:

$$I_{mn0} = I_{mn}^0 \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n. \quad (7)$$

Similarly, the initial state of all remaining decision variables is determined by their pre-planned values in the absence of disruption as:

$$F_{ma0}^{In} = F_{ma}^{In,0} \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a \quad (8a)$$

$$F_{ma0}^{Out} = F_{ma}^{Out,0} \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a \quad (8b)$$

$$P_{pr0} = P_{pr}^0 \quad \forall p \in \mathcal{P}, r \in \mathcal{R}_p \quad (8c)$$

$$D_{mc0} = D_{mc}^0 \quad \forall c \in \mathcal{C}, m \in \mathcal{M}_c \quad (8d)$$

$$B_{ms0} = B_{ms}^0 \quad \forall s \in \mathcal{S}, m \in \mathcal{M}_s \quad (8e)$$

In practice, this approach can be deployed via a rolling horizon, or a shrinking horizon technique [4, 49]. These strategies not only allow decision-makers to update model parameters with current information but also facilitate the incorporation of new disturbances, providing the robust advantages of feedback [50].

Ensuring a notion of recovery post-disruption is important to facilitate a smoother transition to nominal inventory policies after the disruption has been addressed. Defining a sense of normality in a supply chain system prone to constant perturbations can be subject to multiple interpretations. Thus, in this work our objective is to bring the system back to the state it was before the disruption occurred. In the base model we do this by imposing hard terminal constraints as:

$$I_{mn|\mathcal{T}} = I_{mn}^0 \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, \quad (9)$$

which have been used successfully to recover disrupted systems in the literature [51]. Nevertheless, since this approach may appear overly restrictive to certain modelers, we propose an alternative method for handling the final state in Section 3.7.4.

If the final state is not handled properly, the optimizer may drive inventories to zero once the reactive operation concludes due to the existence of holding costs in the objective function [52]. Furthermore, the practical deployment of this model can be perceived as an instantiation of model predictive control, wherein solution profiles (i.e., schedules) are derived at each time step, and these solutions are implemented using a rolling horizon strategy [43, 53]. In such a setup, ensuring stability during the rolling horizon requires proper handling of the end time point, and the use of a hard terminal constraint in Equation (9) constitutes a viable option [50]. Furthermore, in the following subsection, we introduce additional order management options that ensure feasibility, even with the imposition of a hard terminal constraint.

3.4 Order Management Dynamics

An important innovation of this work is accounting for late deliveries and order cancellations as potential reactive actions, which is key given that after a disruption occurs, adhering to the original order schedule may become infeasible or cost prohibitive.

In this model, orders may be satisfied later than their initially planned delivery date by paying a late delivery penalty. To achieve this, unmet demand is accumulated, and it can be progressively fulfilled over time as resources become available in the system. With this model, fractions of the original order can be delivered to the customer incrementally until the initially accepted order is entirely fulfilled. To incentivize the prompt satisfaction of delayed orders, unmet demand is penalized for each period that it remains undelivered.

The reactive response can be customized for individual customers and products by tailoring the late delivery cost profile for each customer and material, as illustrated in Figure 3a. Correspondingly, each profile results in a distinct cost accumulation for units that remain unmet, as depicted in Figure 3b. For this, we assume that orders for the same material and customer are sufficiently spaced in time, which is reasonable given in the chemical product supply chain orders for the same product are typically large in size and widely dispersed in time. This assumption is made to minimize order interference in the profiles. Nevertheless, if this condition does not hold, a constant penalty per customer and material will still provide a sense of prioritization as needed.

The example given in Figure 3 examines an order scheduled for period 2 for a specific customer and material to illustrate the flexibility in prioritization that can be achieved. For simplicity, we show the first order in the schedule post-disruption, hence, no cost profile is required for previous time periods. Alternative 1 features a constant cost that increases in the second half of the time horizon representing an average customer where no particular priority is required, maintaining a constant cost with an increased penalty for older unmet demand. In contrast, Alternative 2 outlines a profile for a higher-priority customer featuring a larger initial cost and an exponentially increasing profile, incentivizing a faster delivery than the previous alternative. Alternative 3 reflects a scenario where a customer tolerates late deliveries for five days, keeping the cost low during this grace period, but significantly escalating afterward, showcasing how more complex customer dynamics can be integrated into the model by adjusting the cost profile. Analogously, Figure 3b illustrates how the late delivery cost accumulates over time for a single unit of unmet demand, providing the total penalty cost associated with delivering at each time period for the previous alternatives.

In scenarios where delivering an order significantly past its original due date is unavoidable, it might be worth considering the option of canceling the order rather than delivering late. Canceling an order has the potential to alleviate the network, enabling the allocation of resources to other orders more efficiently and leading to a more cost-effective response. Consequently, our model incorporates a feature for order cancellation, allowing the optimizer to make informed decisions about whether to accept an order, and if accepted, how late it will be progressively (or entirely) fulfilled. This flexibility enhances the response of the model, allowing it to optimally manage the fulfillment of orders and their potential cancellation in the face of disruptions. If the model decides to cancel an order, it incurs a fixed penalty reflecting the impact of this cancellation on the goodwill of the company with its customers.

By leveraging the GDP framework, we can formulate a disjunction for the cancellation and its impact on




Figure 3: Potential late delivery cost profiles and respective cost accumulations to account for different customer priorities.

the unmet demand (U) balance as follows:

$$\begin{bmatrix} Y_{mct} \\ U_{mct} = U_{mc\{t-1\}} - D_{mct} \end{bmatrix} \vee \begin{bmatrix} -Y_{mct} \\ U_{mct} = U_{mc\{t-1\}} - D_{mct} + \delta_{mct} \end{bmatrix} \quad \forall c \in \mathcal{C}, m \in \mathcal{M}_c, t \in \mathcal{T}. \quad (10)$$

This formulation implies that when an order is canceled ($Y=\text{True}$), that order amount δ should not be accounted for in the balance allowing the operation to focus on the unmet demand accumulated from previously accepted orders. Conversely, when the order is accepted (or not canceled, $Y=\text{False}$), the amount ordered δ is included in the unmet balance. This allows the model to determine if the order will be delivered on time or late, and if so, how it will be progressively fulfilled.

As mentioned in Section 1, GDP is introduced as a modeling framework to pose the constraints more clearly. Since, the problem is not solved in disjunctive form, we introduce an equivalent MILP reformulation where y is the binary variable associated with Y (meaning $y = \mathbb{1}\{Y\}$, where $\mathbb{1}\{\cdot\}$ is the indicator function) as:

$$U_{mct} = U_{mc\{t-1\}} - D_{mct} + \delta_{mct}(1 - y_{mct}) \quad \forall c \in \mathcal{C}, m \in \mathcal{M}_c, t \in \mathcal{T} \quad (11)$$

The resulting expression for the unmet demand is calculated as an inventory balance to account for accumulated unmet demand across the time periods. Since unmet demand is nonnegative, the model can choose to cancel a specific order ($y = 1$), relieving the system but incurring a high fixed price (λ^δ). On the other hand, if a given order is not canceled ($y = 0$), the reactive operation tries to fulfill both the given order and any accumulated unmet demand. The above expression applies only when an order is present, meaning that δ is non-zero. If there is no order (i.e., $\delta = 0$), the cancellation variable y can be set to one. The impact of canceled orders on the objective function is discussed in Section 3.6.

3.5 Inventory bounds

Upper bounds for the inventory levels in plants and warehouses are typically given by the physical capacity allocated to each material within the facility (I_{mnt}^U). However, when materials are not segregated, inventory may be restricted relative to the total facility volume (I_{nt}^U) as:

$$\sum_{m \in \mathcal{M}_n} I_{mnt} \leq I_{nt}^U \quad \forall n \in \mathcal{P} \cup \mathcal{W}, t \in \mathcal{T} \quad (12)$$

Determining the lower bound for inventory levels is a decision influenced by the risk aversion of the modeler. The riskier approach permits inventories to reach zero, allowing all available buffer stock to be utilized during disruptions if necessary. Conversely, the more conservative approach retains a lower bound on the buffer stock (SS), based on the argument that buffer stock is primarily intended for hedging against demand uncertainty rather than disruptions. Both perspectives can be captured and combined in the following expression:

$$I_{mnt}^L = \alpha_{mnt}SS_{mnt} \leq I_{mnt} \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (13)$$

where $\alpha \in [0, 1]$ represents the fraction of the buffer stock SS that the model preserves regardless of the disruption. The inventory bounds used in the base model are explicitly shown in the next section (see Equation (15c)) where a capacity is assumed for each material individually and no lower bound is placed to yield the least conservative schedules ($\alpha = 0$). Nevertheless, an alternative where this assumption is relaxed is proposed in Section 3.7.5.

3.6 Objective Function

The objective of the model is to maximize profit from the sales (λ^D), while accounting for several costs such as buying material from suppliers, manufacturing cost at different facilities (λ^P), shipping cost using different transportation types (road, marine or air-freight) (λ^F), and holding inventory cost at both plants and warehouses (λ^I). Furthermore, this model accounts for cost of delivering late (λ^U), and for a high fixed cost in the case where an order needs to be canceled (λ^δ) as shown below:

$$\begin{aligned} \max \sum_{t \in \mathcal{T}} [& \sum_{c \in \mathcal{C}} \sum_{m \in \mathcal{M}_c} (\lambda_{mct}^D D_{mct} - \lambda_{mct}^U U_{mct} - \lambda_{mct}^\delta y_{mct}) - \sum_{a \in \mathcal{A}} \sum_{m \in \mathcal{M}_a} \lambda_{mat}^F F_{mat}^{Out} - \sum_{s \in \mathcal{S}} \sum_{m \in \mathcal{M}_s} \lambda_{mst}^B B_{mst} \\ & - \sum_{p \in \mathcal{P}} \sum_{r \in \mathcal{R}_p} \lambda_{prt}^P P_{prt} - \sum_{n \in \mathcal{P} \cup \mathcal{W}} \sum_{m \in \mathcal{M}_n} \lambda_{mnt}^I I_{mnt}] \end{aligned} \quad (14)$$

The values of λ^U and λ^δ capture the impact that delays in delivery or order cancellations can have on the reputation of the company and its customer relations. Determining the exact value for these parameters is a modeling decision that must obtain a balance between financial considerations and customer relationships. These penalties have bounds that can be inferred from the system. The upper bound is the net present value of all forecasted sales from the affected customer over 15 to 20 years, representing the worst-case scenario of customer loss due to a canceled order or late delivery. The lower bound is the cost of sourcing the product from a competitor and selling it at cost, covering logistical expenses without affecting goodwill.

Finally, all continuous variables are nonnegative and bounded as shown below:

$$0 \leq F_{mat}^{In} \leq F_{mat}^U \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} \quad (15a)$$

$$0 \leq F_{mat}^{Out} \leq F_{mat}^U \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} \quad (15b)$$

$$0 \leq I_{mnt} \leq I_{mnt}^U \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (15c)$$

$$0 \leq P_{prt} \leq P_{prt}^U \quad \forall p \in \mathcal{P}, r \in \mathcal{R}_p, t \in \mathcal{T} \quad (15d)$$

$$0 \leq D_{mct} \leq D_{mct}^U \quad \forall c \in \mathcal{C}, m \in \mathcal{M}_c, t \in \mathcal{T} \quad (15e)$$

$$0 \leq B_{mst} \leq B_{mst}^U \quad \forall s \in \mathcal{S}, m \in \mathcal{M}_s, t \in \mathcal{T} \quad (15f)$$

where the bounds are time-indexed. Additional considerations on bounding variables to model disruptions will be presented in Section 4.

3.7 Modeling Extensions

The base model is composed of Equations (1)-(9), (11)-(15) and it provides a reactive supply chain operation solution in response to the disruptions under mild assumptions. The following subsections illustrate how the base model can be extended to accommodate more flexible responses and dynamics, when some of the initial assumptions are relaxed. Later in this paper (see Section 6.2), we will discuss the impact that each of the extensions, individually and in combination, has on the computational performance.

3.7.1 Production Across Time Periods (PATP)

The base model assumes time-independent production, meaning that the recipe starts and concludes within the same period as shown in the Equation (4). However, there are instances where this assumption may not hold true, either due to processes inside the plant taking longer than the time discretization (e.g., brewing or fermenting), or because the modeler opts for a finer time scale to capture more detailed schedules. To model production across time periods (PATP), we follow a similar approach to that used for modeling shipping time delays, as demonstrated in Equations (1)-(2). We define P^{In} and P^{Out} as the quantity of produced recipe that initiates and concludes (respectively) within the specified period. Now, the time-dependent production can be expressed as:

$$P_{prt}^{In} = P_{pr\{t+\tau_{prt}^P\}}^{Out} \quad \forall p \in \mathcal{P}, r \in \mathcal{R}_p, t \in \mathcal{T} \quad (16a)$$

$$P_{prt}^{In} = 0 \quad \forall p \in \mathcal{P}, r \in \mathcal{R}_p, t \in \mathcal{T} : t + \tau_{prt}^P > |\mathcal{T}| \quad (16b)$$

$$P_{prt}^{Out} = 0 \quad \forall p \in \mathcal{P}, r \in \mathcal{R}_p, t \in \mathcal{T} : t - \tau_{prt}^P < 0 \quad (16c)$$

where τ_{prt}^P is the time it takes to execute the recipe when production starts at period t . Consequently, the plant inventory balance originally expressed in Equation (4) can be written in terms of the time-dependent production as:

$$I_{mpt} = I_{mp\{t-1\}} + \sum_{a \in \mathcal{A}_p^{In}} F_{mat}^{Out} - \sum_{a \in \mathcal{A}_p^{Out}} F_{mat}^{In} + \sum_{r \in \mathcal{R}_p | \phi_{rm} > 0} \phi_{rm} P_{prt}^{Out} + \sum_{r \in \mathcal{R}_p | \phi_{rm} < 0} \phi_{rm} P_{prt}^{In} \quad \forall p \in \mathcal{P}, m \in \mathcal{M}_p, t \in \mathcal{T} \setminus \{0\} \quad (17)$$

Therefore, the model that supports manufacturing that occurs across time periods is given by Equations (1)-(3), (5)-(9), and (11)-(17), where the production variable is redefined.

3.7.2 Fixed Transportation Cost (FTC) and Minimum Shipping Amount

The base model assumes that shipping costs are proportional to the quantity of material shipped. However, in many situations, there may be fixed transportation costs (FTC) associated with the use of a particular route. Furthermore, there may also be minimum shipment amount requirements.

To extend the model to account fixed transportation costs and minimum quantities, we define X as a Boolean variable that indicates flow existence according to the following disjunction:

$$\left[\begin{array}{c} X_{mat} \\ F_{mat}^L \leq F_{mat}^{In} \leq F_{mat}^U \end{array} \right] \vee \left[\begin{array}{c} \neg X_{mat} \\ F_{mat}^{In} = 0 \end{array} \right] \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} \quad (18)$$

where flow is activated only when $X = \text{True}$, and it is bounded by both the minimum quantity and the capacity of that arc $[F^L, F^U]$. Similar to the disjunction shown in Equation (10), we propose an MILP reformulation using the associated binary variable x (where $x = \mathbb{1}\{X\}$) as:

$$F_{mat}^L x_{mat} \leq F_{mat}^{In} \leq F_{mat}^U x_{mat} \quad \forall a \in \mathcal{A}, m \in \mathcal{M}_a, t \in \mathcal{T} \quad (19)$$

allowing the flow to be non-zero only when there is material traversing the arc. Also, x can be used to account for fixed transportation costs by subtracting:

$$FixedCost = \sum_{t \in \mathcal{T}} \sum_{a \in \mathcal{A}} \sum_{m \in \mathcal{M}_a} \lambda_{mat}^{Fix} x_{mat} \quad (20)$$

from the objective function of the base model depicted in Equation (14).

The extension that considers fixed transportation cost and minimum shipping amount is given by Equations (1)-(9), (11)-(15), and (19)-(20). This extension introduces a significant number of binary variables and big-M constraints (like Equation (19)) which can impact the solution time and the tightness of the linear programming relaxation of the MILP [54]. The impact of these extensions on the solution time and scalability of the model are shown in Section 6.2.

3.7.3 Service Level Agreements (SLA)

Service level agreements (SLA) are contracts that may exist with the potential suppliers. If the supply chain company desires to buy material from the supplier, then a minimum amount may need to be purchased throughout a period of time.

Service level agreements can be modeled in two different ways. The first option is to buy more than a minimum permitted amount and can be used when the time discretization (i.e. daily or weekly) of the model is equal or larger than the one in the contract. In that case, a simple minimum quantity bound is imposed as:

$$\left[\begin{array}{c} W_{mst} \\ B_{mst}^{SLA} \leq B_{mst} \leq B_{mst}^U \end{array} \right] \vee \left[\begin{array}{c} -W_{mst} \\ B_{mst} = 0 \end{array} \right] \quad \forall s \in \mathcal{S}^{SLA}, m \in \mathcal{M}_s, t \in \mathcal{T} \quad (21)$$

where W is the Boolean variable that determines if the SLA is executed at time t , and the model commits to purchase the required quantities. Furthermore, we define B^{SLA} as the minimum purchase amount.

In the second option, where the time partitioning of the model is smaller than the one in the agreement, a time window has to be considered. Here, the model has to guarantee that the quantity purchased over the SLA time window is within the contract specification. This can be modeled as follows:

$$\left[\begin{array}{c} W_{mst} \\ B_{mst}^{SLA} \leq \sum_{t'=t}^{t+\tau_{mst}^{SLA}} B_{mst'} \leq B_{mst}^U \end{array} \right] \vee \left[\begin{array}{c} -W_{mst} \\ B_{mst} = 0 \end{array} \right] \quad \forall s \in \mathcal{S}^{SLA}, m \in \mathcal{M}_s, t \in \mathcal{T} : t + \tau_{mst}^{SLA} \leq |\mathcal{T}| \quad (22)$$

where τ^{SLA} is the duration of the time window. The equation above assumes that a contract cannot extend beyond the time horizon; thus, the disjunction is only imposed when the SLA can be fulfilled.

The two disjunctions shown in Equations (21) and (22) can be reformulated to equivalent MILP constraints using the associated binary variable w as expressed in Equations (23a) and (23b) respectively.

$$B_{mst}^{SLA} w_{mst} \leq B_{mst} \leq B_{mst}^U w_{mst} \quad \forall s \in \mathcal{S}^{SLA}, m \in \mathcal{M}_s, t \in \mathcal{T} \quad (23a)$$

$$B_{mst}^{SLA} w_{mst} \leq \sum_{t'=t}^{t+\tau_{mst}^{SLA}} B_{mst'} \leq B_{mst}^U w_{mst} \quad \forall s \in \mathcal{S}^{SLA}, m \in \mathcal{M}_s, t \in \mathcal{T} : t + \tau_{mst}^{SLA} \leq |\mathcal{T}| \quad (23b)$$

This discussion only applies for the specific set of suppliers \mathcal{S}^{SLA} that require service level agreements where $\mathcal{S}^{SLA} \subseteq \mathcal{S}$.

The model that integrates SLA includes Equations (1)-(9), (11)-(15), and either Equation (23a) or (23b). As above, this extension requires additional binary variables and big-M constraints, which may negatively affect the solution time. These effects will be studied in Section 6.2.

3.7.4 Final Inventory Deviation (FID)

As mentioned earlier, providing a notion of recovery from the disruption is key to resume “normal” operation. In the base model, this notion was given by Equation (9) as a hard terminal constraint that matched pre and post-disruption states exactly. These constraints can sometimes be too stringent and may potentially yield an unnecessary amount of late deliveries and cancellations. The hard constraints can be relaxed by replacing them with an L1 penalty in the objective to the final inventory deviation (FID) as:

$$Deviation_{mn} = |I_{mn}^0 - I_{mn|\mathcal{T}}| \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n \quad (24)$$

which can be reformulated as:

$$Deviation_{mn} \geq I_{mn}^0 - I_{mn|\mathcal{T}} \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n \quad (25a)$$

$$Deviation_{mn} \geq I_{mn|\mathcal{T}} - I_{mn}^0 \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n \quad (25b)$$

and which in turn, can replace Equation (9) in the formulation. Furthermore, the total deviation is calculated as:

$$TotalDeviation = \sum_{n \in \mathcal{P} \cup \mathcal{W}} \sum_{m \in \mathcal{M}_n} \lambda_{mn}^{Dev} Deviation_{mn} \quad (26)$$

where a penalty (λ^{Dev}) incentivizes the difference to be as small as possible, and it can be subtracted from the objective function. Intuitively, a large penalty can be applied when prioritizing rapid recovery from the disruption. Therefore, the formulation extended to account for a final inventory deviation penalty includes Equations (1)-(8), (11)-(15), and (25)-(26).

The penalty cost is indexed by location and material enabling us to prioritize which products and places more urgently require a base inventory to resume their normal operation. It is worth noting that while this approach can help achieve solutions with fewer backorders and cancellations, certain assumptions (related to a gradual shrinkage of the deviations) must hold to ensure stability in a rolling horizon framework. However, addressing these assumptions is outside the scope of this work, therefore we refer the readers to Rawlings [50] for more information on this subject.

3.7.5 Negative Inventory Deviation (NID)

The base model considers hard lower bounds on inventory variables. In this section, we extend the model with soft bounds to allow for more flexible responses. Contrary to Equation (24) where a symmetric L1 penalty was enforced as constraint, here we only impose a penalty when the inventory level is below a threshold given by the lower bound I^L . We refer to this one-sided deviation as negative inventory deviation (NID). This alternative allows the model to go below the determined threshold at the cost of a penalty (λ^K). Naturally, larger values for λ^K model a more conservative approach where inventory values above the threshold value are preferred. The logic is captured by the disjunction:

$$\left[\begin{array}{l} Z_{mnt} \\ I_{mnt} - I_{mnt}^L \leq 0 \\ K_{mnt} = I_{mnt} - I_{mnt}^L \end{array} \right] \vee \left[\begin{array}{l} -Z_{mnt} \\ I_{mnt} - I_{mnt}^L \geq 0 \\ K_{mnt} = 0 \end{array} \right] \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (27)$$

where Z is the Boolean variable that indicates if there is a violation below the threshold, and K represents the total negative deviation. We also use here an MILP reformulation for the disjunction. First, we note that the original upper bounds on inventory posed in Equation (15c) are still required. Therefore, we can derive the following bound for the deviation:

$$-I_{mnt}^L \leq I_{mnt} - I_{mnt}^L \leq I_{mnt}^U - I_{mnt}^L \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (28)$$

The disjunction indicates that the negative deviation (K) behaves like a rectified linear unit (ReLU) reflected over the y -axis and displaced by I^L . Figure 4 illustrates this penalty as a function of inventory. We reformulate the disjunction using the ReLU formulation from Ceccon et al. [55], reflected over the y -axis and displaced by the threshold, as follows

$$K_{mnt} \leq I_{mnt} - I_{mnt}^L \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (29a)$$

$$K_{mnt} \geq I_{mnt} - I_{mnt}^L + (I_{mnt}^L - I_{mnt}^U)(1 - z_{mnt}) \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (29b)$$

$$K_{mnt} \geq -I_{mnt}^L z_{mnt} \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (29c)$$

$$K_{mnt} \leq 0 \quad \forall n \in \mathcal{P} \cup \mathcal{W}, m \in \mathcal{M}_n, t \in \mathcal{T} \quad (29d)$$

The deviation along the time horizon is aggregated and penalized to obtain the total cost of inventory bounds violation:

$$TotalNegativeDeviation = \sum_{n \in \mathcal{P} \cup \mathcal{W}} \sum_{m \in \mathcal{M}_n} \sum_{t \in \mathcal{T}} \lambda_{mnt}^K K_{mnt} \quad (30)$$

which has to be added to the objective of the base model to penalize lower bound violations of the threshold. Therefore, the model incorporates a lower bound penalization for the inventory is given by Equations (1)-(9), (11)-(15), and (28)-(30).




Figure 4: Negative deviation from the threshold I^L as a function of the inventory I of material A in Plant 1. The inventory is bounded as $I \in [0, 100]$ and a constant threshold $I^L = 50$ is imposed.

In the section above, we established the base model as a powerful formulation for reactive responses in supply chain and manufacturing networks. Additionally, we expanded the model to accommodate scenarios with relaxed initial assumptions, yielding more adaptable responses. In Section 6.2 we will discuss the implications of these extensions on the computational performance. In the following section, we will demonstrate how to integrate various disruptions that occur in supply chains into the model.

4 Disruption Modeling

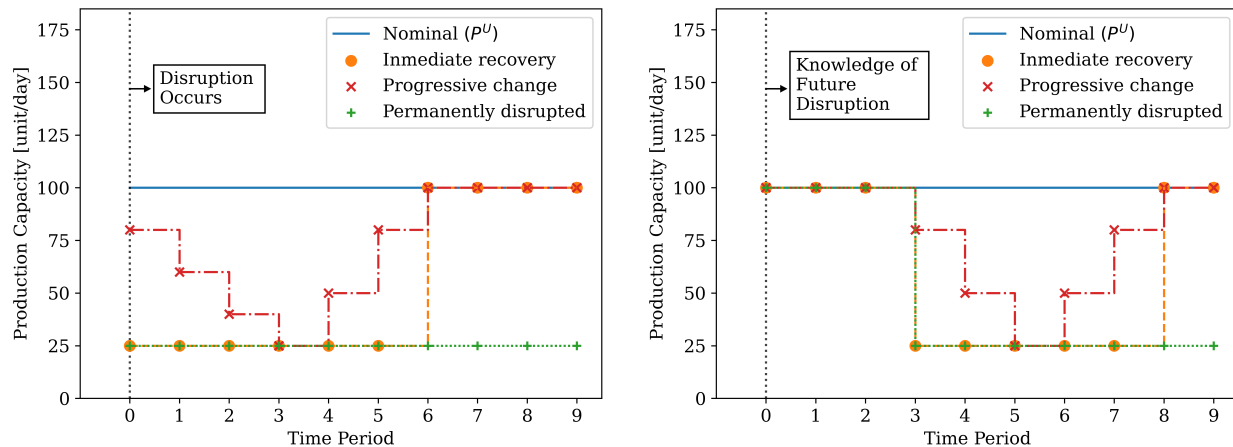
The proposed framework enables the modeler to incorporate various types of disruptions that may arise in a supply chain system. Generally, a disruption denotes a deviation of the system from its nominal operation. This study classifies disruptions into two primary categories: immediate disruptions and scheduled disruptions. Immediate disruptions occur unexpectedly or in an unplanned manner, including events like equipment failures, roadblocks, or supply shortages. The model formulation is designed to generate a reactive schedule in response to these disruptions, allowing the system to adapt to the unexpected changes. Conversely, scheduled disruptions are events for which abnormal circumstances are known in advance. This category includes instances such as regular maintenance or planned shutdowns. In these scenarios, the formulation can offer proactive policies, enabling the network to prepare for the disruption ahead of time and operate effectively throughout the disruption period.

Characterizing a disruption involves two parts. First, identifying the disrupted horizon meaning when the disruption starts and ends. The start can be at $t = 0$ for an unplanned disruption. Similarly, the end can be at period $t = |\mathcal{T}|$ if the end of the disruption is unclear or if the system is known to be permanently disrupted. Second, the set of disrupted parameters must be identified, which captures the negative impacts on the system, causing deviations from normal operation throughout the disrupted horizon.

A significant portion of disruptions occurring in real-life systems can be mapped, in some manner, to a reduction of available capacity. Examples of this concept include a roadblock impacting edge capacity, a supply shortage reducing purchasing capacity, or an equipment failure decreasing production capacity. By indexing both upper and lower bounds of continuous variables over time, bound profiles can be generated to model disruptions by tightening the bounds.

Consider the example where there is a reduction in plant production, hence the disrupted parameter corresponds to production bounds. Figure 5 illustrates how both immediate and scheduled disruptions can

be modeled by designing a bound time profile. These profiles not only encode information about the time of disruption and recovery, but also about how the capacity precisely decays and increases over time.



(a) Bound profiles to model unexpected (immediate) disruptions.

(b) Bound profiles to model known (scheduled) disruptions.

Figure 5: Potential production upper bound profiles for the production of a recipe to model a disruption that reduces the capacity by 75% in both immediate and scheduled cases.

The nominal upper bound, depicted in blue, represents the capacity of the current operation, which serves as a baseline for the undisrupted system. In certain disruptions, the system recovers immediately after the event has passed, as shown in orange, where no additional time to return to its usual capacity. In other cases, disruptions may progressively tighten and then relax the bound over time, illustrated in red, offering the modeler flexibility to characterize the disruption dynamically. In green are modeled scenarios where the system is assumed to be permanently disrupted, either due to a severe disruption, or a lack of knowledge about when the disrupted horizon will end. Furthermore, both bounds can collapse to a single value at a specific time if necessary. For example, having the upper bound go to zero can model a complete shutdown in plant production.

Other disruptions can be modeled by changing parameters other than variable bounds. For instance, if shipping through a route takes longer due to external factors, the reactive response might involve finding alternative routes, and adjusting downstream operations to handle the lead time disruption. In our approach, transportation lead times can vary over time. To model an increase in transportation delay, simply adjust the nominal lead times. This discussion can also be extended to production delays as proposed in Section 3.7.1.

In supply chain operations, there are instances where economic values, such as costs and prices change often due to external factors that affect the various stakeholders involved in the operation. Examples include market fluctuations impacting product prices, or supply shortages affecting not only supplier capacity but also the cost of the product. To model these changes in economic parameters, it suffices to impose the disrupted prices and costs. By doing so and obtaining an optimal reactive schedule, decision-makers can assess how significantly the optimal reaction would differ from the usual operation.

There is one case of disruption we consider in this work that aligns with what most of the literature has focused on when it comes to supply chain networks under uncertainty: changes in the demand. Unlike most literature that provides fixed policies to address sudden changes or spikes in demand, our reactive approach aims to optimize the operation of the system in response to unexpected orders. If a new order is placed, our approach evaluates how to readjust the system to integrate the incoming order. Here the model may decide to deliver the order late or cancel it to avoid inducing stress on the system. In this case, the set of disrupted parameters must not only include the size of the order but also the associated late delivery cost profile and cancellation cost of the order place at period t' for the specific material m' and customer c' . Furthermore, if

there are specific reasons (such as customer importance or contractual obligations) for which the new order cannot be delivered late or canceled, this policy can be enforced by:

$$U_{m'c't} = 0 \quad \forall t \in \mathcal{T} \quad (31a)$$

$$y_{m'c't'} = 0 \quad (31b)$$

in which case the final state of model has to be extended, as proposed in Section 3.7.4 to guarantee feasibility given that an order is being forced to be delivered.

This proposed approach is capable of integrating various types of disruptions that occur in real life in an equation-oriented manner to provide an optimal response. Impacts in operations are modeled independently by design, such that the modeler determines if multiple disruptions are present in a single event. Furthermore, the proposed method gives a way to bridge most of the existing literature to account for consequences of unexpected demands. In the following sections of the paper, we present a case study and evaluate the computational performance of the model.

5 Motivating Case Study

The supply chain network we consider is motivated by a part of the silicone rubber business of The Dow Chemical Company. It manufactures ten different products, and consists of two raw material suppliers, two plants, two warehouses, and nine customers that are satisfied from both manufacturers and warehouses. We consider the operation over a four month period ($|T| = 120$ days) and the structure of the motivating example is presented in Figure 6.




Figure 6: Supply chain and manufacturing network topology for the motivating case study

The first plant (Plant 1) uses the two different raw materials (Raw A and Raw B) to make six products (DP, Fluid F, Blend A, Blend B, Blend C and Blend C1). In this plant, two products (labeled in the figure as DP and Blend C1) are sold to customers directly, while the rest of the products are packed in steel drums and shipped to Warehouse 1 through road transportation. Note that Blend C is both an intermediate material in the production of Blend C1 as well as a final product that is delivered to Warehouse 1. Once in Warehouse 1, these four products can be sold directly to local customers or shipped to Warehouse 2 through marine or air transportation. There, three products (labeled in the figure as Blend A, Blend B, and Blend C) are sold directly to local customers. Fluid F is shipped to Plant 2 to be further processed into four different final products (SR 1 - SR 4) and sold to local customers. All nine customers have independent demand frequencies and order quantities.

Due to its size and complexity, the network can be exposed to several unplanned events that cause disruptions in its operation. This particular network can face, for example, four different types of disruptions that could negatively impact the overall operation.

- Raw material shortage: This can occur when an upstream supplier is unable to deliver the amount of raw material originally requested by the company for scheduled operation.
- Plant equipment failure: This causes a reduction or complete loss of production at a plant facility, impacting downstream consumers and inventory levels for both upstream and downstream elements.
- Route unavailability: A given arc in the supply chain may not be capable of transporting any material for a specific time window. In this case, the network becomes partially disconnected and alternative routing options may be needed.
- Logistics resource limitation: The transport of material between supply chain elements can be impacted by availability of appropriate logistics resources (e.g., packing materials like steel drums or availability of drivers for truck transport). Furthermore, a limitation in the resources used for storage can affect the capacity in warehouses.

These examples can cause a reduction in available material for downstream processes and delays in deliveries of customer orders. However, these can also impact upstream operations, and cause inventory management challenges. For example, a plant failure causes an increase in raw material inventory that can impact upstream suppliers and producers. Hence, an integrated approach that considers the network holistically is desired.

6 Results

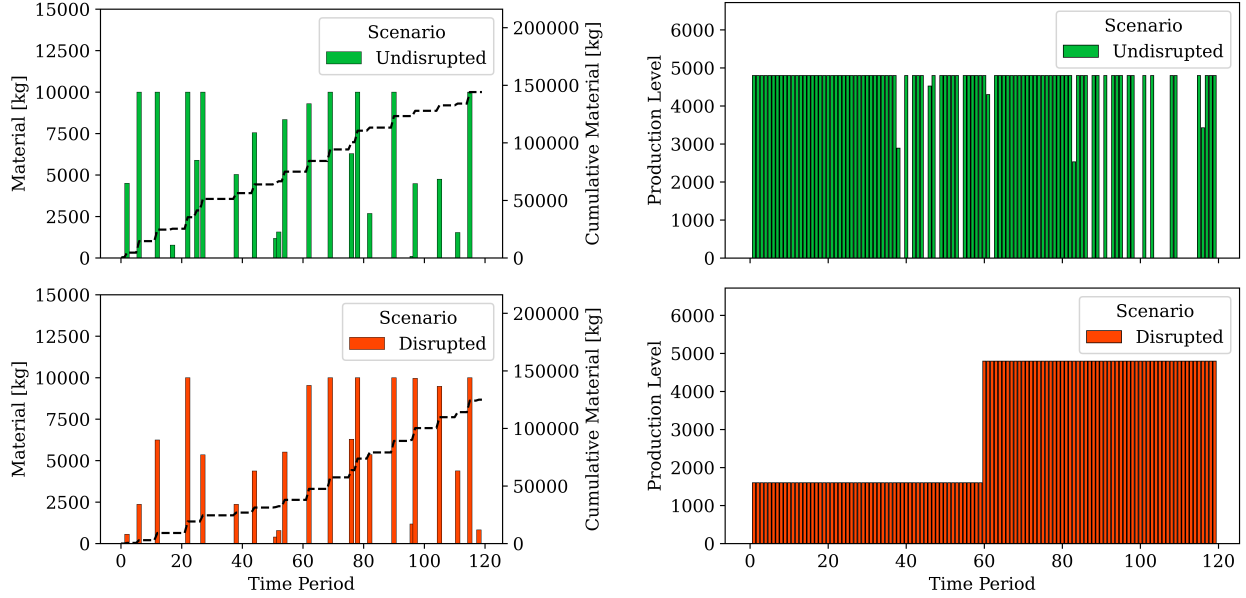
In the following subsections, we present the computational results of the proposed model. We begin by presenting the optimal reactive schedules generated by the model for two scenarios. This provides insight into how the model determines an integrated schedule for supply chain operations and demonstrates how disruptions can propagate beyond the directly affected nodes. Then, we investigate how the computational solution time is influenced by increasing the size of the model, particularly by extending the time horizon and examining the impact of different model extensions. This section provides details on the size of the models studied, including the number of binary variables, continuous variables, constraints, and non-zero elements. Subsequently, we evaluate the different optimal operations obtained with the model across various characterizations of the same disruption. Finally, we conduct a sensitivity analysis on the cancellation and late delivery penalties, which are key parameters that impact optimal decisions.

The proposed solution approach, which includes all the extensions, was implemented in Pyomo [56] and the resulting MILP was solved with Gurobi v9.5.1 as the solver on a Linux machine with 8 Intel® Xeon® Gold 6234 CPUs running at 3.30 GHz with 8 total hardware threads and 1 TB of RAM running Ubuntu. All the data presented from the Dow Chemical Company has been anonymized, therefore, the values and magnitudes showcased in this section do not reflect the actual operational figures of the company. The source code and sample data for illustration are freely available in the following open-source [Github repository](https://github.com/dovallev/supply_chain_disruptions)⁴.

6.1 Optimal Reactive Operation

In this section, we demonstrate the proposed formulation on a single case study and analyze the optimal solution. We consider the supply chain and manufacturing network outlined in Section 5, and we compare the solution from a single disruption scenario with the solution in the undisrupted case. The resulting models have 17,345 continuous variables, 497 binary variables, and 15,480 constraints for the base model with a solution time of 0.23 seconds. The disruption involves the unplanned failure (meaning that it occurs at time zero) of two out of three reactors at Plant 1, which produce DP, Fluid F, and Blend A. This reduces their production capacity to one-third of the nominal value for two months. We examine the optimal response of the base model from Section 3 over the next four months (120-day discretization), focusing on the transformation of Raw A into Fluid F and its subsequent use in producing SR 3. While our discussion centers on these specific materials, it is important to note that the model provides a comprehensive schedule for all materials in the network, with the disruption impacting the entire operation due to the critical nature of Raw A.

⁴https://github.com/dovallev/supply_chain_disruptions



(a) Total procurement of Raw A from supplier.

(b) Production of Fluid F from Raw A in Plant 1.

Figure 7: Procurement and production levels in Plant 1 for Raw A and Fluid F.

Figure 7 illustrates the transformation of Raw A into Fluid F that takes place in Plant 1 for both undisrupted and disrupted scenarios. More specifically, Figure 7a presents the procurement schedule of Raw A from an external supplier, while Figure 7b depicts the production schedule for Fluid F from Raw A in Plant 1. The disruption shifts larger procurement orders to the second half of the operation in contrast to the undisrupted case where most orders are placed in the first half. This can be seen best by comparing the slope of cumulative procurement line in each scenario: the undisrupted scenario shows faster procurement in the first half and the disrupted scenario shows faster procurement in second half. This occurs because, with limited production capacity in the first half of the horizon, the model minimizes raw material procurement to avoid excess inventory. As production ramps up in the second half, increased raw material is then procured to meet the increased demand. This shows how disruptions can impact upstream operations: a disruption at Plant 1 affects the production rate and subsequently affects the need for raw materials and the flow from the upstream supplier.

In the disrupted scenario, the production capacity for Fluid F is reduced to one-third, requiring the plant to operate at full capacity for the entire period to mitigate the disruption. In contrast, the undisrupted case allows for periods where Fluid F production is not at full capacity. Naturally, the plant operates at full capacity once production is restored, aiming to recover quickly and minimize the impact of the disruption.

Figure 8 illustrates how the disruption extends beyond the affected node. Specifically, Figures 8a and 8b depict Fluid F shipments from Warehouse 1 to Warehouse 2 via truck and air freight, respectively. The disruption significantly impacts the optimal shipment of Fluid F. In the undisrupted case, large, sporadic shipments are sent via truck, while the disrupted case sees more frequent, smaller shipments. Notably, airfreight, unused in the undisrupted scenario, is employed from Warehouse 1 to Warehouse 2 in the second half of the disrupted case. This shift occurs as the operation tries to “catch up” using faster, but more expensive transportation modes.

The reactor disruption not only directly impacts Fluid F production, but also affects other materials downstream. This phenomenon is illustrated by the production and inventory profiles of SR 3 at Plant 2, as shown in Figures 9a and 9b, respectively. The SR 3 profiles reveal how the upstream disruptions propagate downstream. In the undisrupted scenario, production orders are larger and more sporadic, whereas, in the disrupted case, orders are smaller and more frequent. These behaviors directly reflect the downstream propagation of the shipping patterns discussed in Figure 8a. Since Fluid F is the raw material used in production, the production schedule in each scenario depends on the quantity and frequency of Fluid F




Figure 8: Shipment of Fluid F from Warehouse 1 to Warehouse 2.

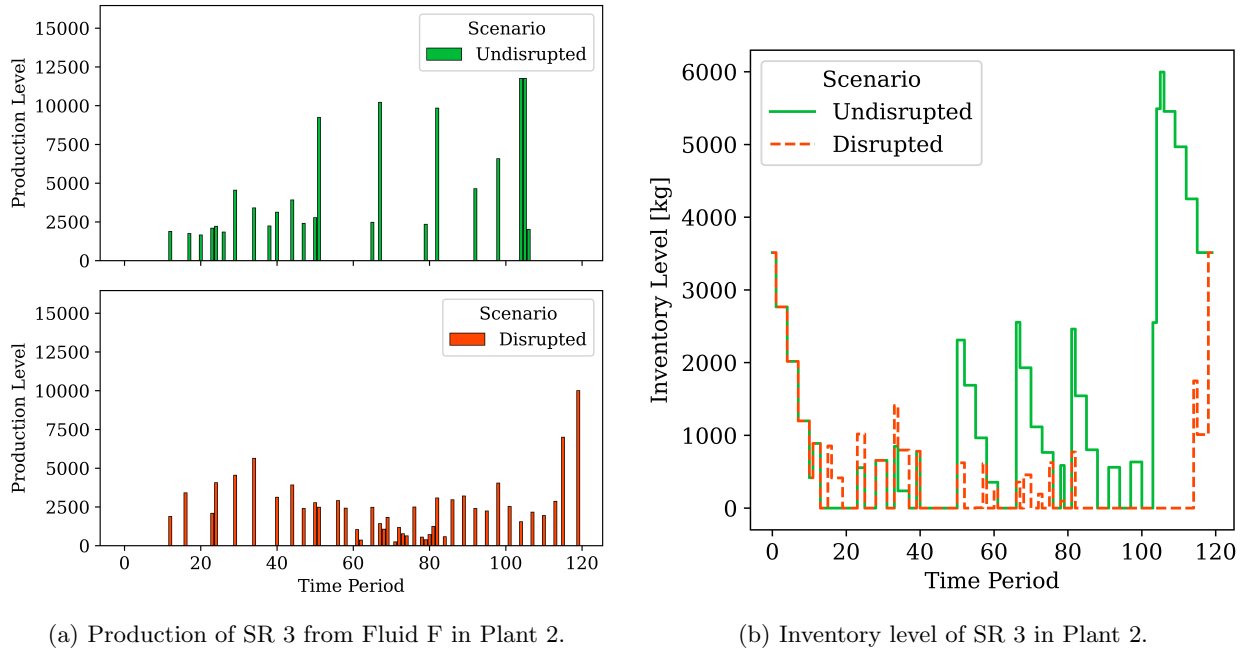



Figure 9: Production and inventory levels of SR 3 in Plant 2.

shipments from upstream. Notably, the disrupted scenario demonstrates the ability of the model to generate an integrated and centralized response, as the strategy to recover quickly (using small, frequent shipments and production runs as they become available) is coordinated across the entire network.

The inventory profiles initially align in both scenarios, but they diverge over time as the disruption propagates through the network. This behavior occurs because the disruption takes time to propagate to this node in the network. This makes sense given that operations remain unaffected as they rely on

existing inventory rather than upstream supplies. Notably, this approach in the initial stages is cost-effective, as utilizing existing stock is cheaper than shipping material from other nodes. In the undisrupted case, inventory increases correspond to production periods with gradual depletion as external demand arises. In the disrupted scenario, we observe lower inventory levels and a higher turnover frequency. In the latter case, the high turnover frequency results from an integrated operation that is trying to recover quickly and is unable to hold inventory for extended periods.



(a) Total amount of material delayed.

(b) Total amount of material canceled.

Figure 10: Total amount of material delayed and canceled throughout the operation.

One of the most significant consequences of disruptions is their impact on demand coverage, affecting the timely delivery of orders to customers. Figure 10 illustrates the order management outcomes of the model. Specifically, Figure 10a shows the total amount of delayed material, while Figure 10b displays the total amount of material from canceled orders. This analysis aggregates data across all 10 materials sold by the network. Analyzing both figures, we observe that the undisrupted case involves no delays or cancellations, as expected for an operation designed for steady-state nominal conditions (hence, no green profiles are present in the figures). In the disrupted scenario, the model responds by canceling large orders in the first half of the time horizon due to insufficient resources, avoiding the high costs of long delays. This early cancellation strategy is complemented by delaying orders in the middle of the time horizon, where there is less time to wait for resources, allowing the fulfillment of these orders with some delay. Notably, the overall order management response is integrated. These decisions are known immediately upon the disruption, enabling decision-makers to communicate with customers and negotiate alternative solutions in advance.

6.2 Time Scalability

Determining the number of time periods in the horizon is a modeling choice that can significantly impact the computational time. In practice, increasing the number of time periods can be done to increase the horizon or to increase the frequency of decisions. In this section we study the impact that the number of time periods has on the solution time of the different extensions to an undisrupted model. We evaluated all possible combinations of models that can be created by including (or not) the extensions discussed in Section 3.7 yielding a total of 48 models. For each model type, a set of 50 replicas are constructed by varying order schedules. These were all solved for increasing number of time periods resulting in models with up to hundreds of thousands of variables and constraints as shown in Table 1. Figure 11 shows the average solution time of all the replicas where each individual line corresponds to a different model, summarizing over 290 hours of computing time.

	Min	Max	Mean	Median
Binary Vars.	25	98,231	26,540.81	19,370
Continuous Vars.	1,199	266,666	122,111.84	121,294
Constraints	898	377,889	132,613.65	122,516
Non-zeros	2,421	881,549	332,432.11	319,199

Table 1: Model size summary statistics.

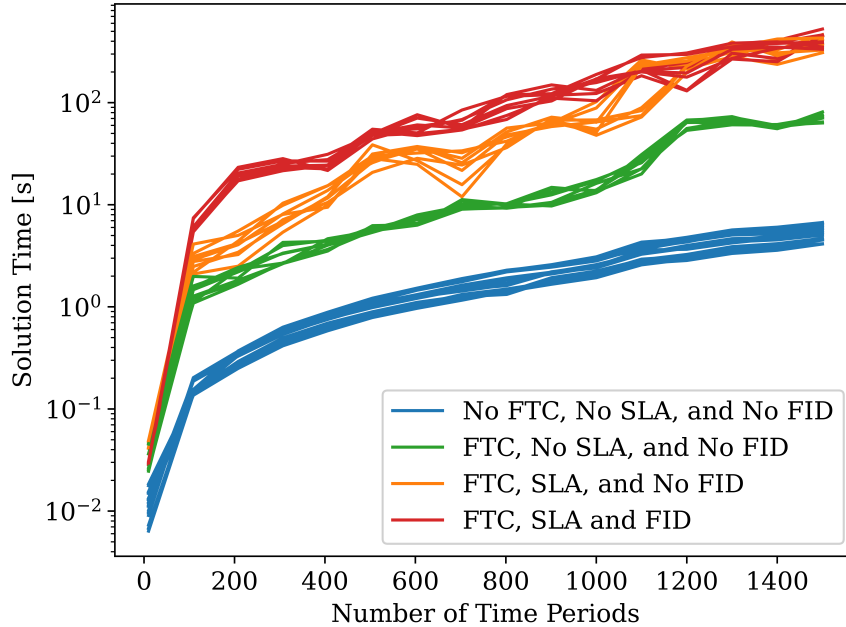


Figure 11: Time scalability of the different model extensions.

While 48 different models were considered, we have identified several key extensions that have the strongest impact on solution time. Figure 11 includes the results for all models; however, we group them according to these high-impact extensions. The y-axis of the figure represents computational time on a logarithmic scale, thus unit increments on this axis reflect orders of magnitude increments in solution time. The extension with the most significant impact on solution time is the fixed transportation cost (FTC) extension, as indicated by all 24 models without this extension solving the fastest (shown in blue). Next, we examine how the FTC extension interacts with other extensions. The extension that most affects computational time when combined with FTC is the service level agreements (SLA), distinguishing the performance between the green and orange groups. Lastly, the addition of final inventory deviation FID notably impacts solution time as evident in the difference between the orange and red groups.

As expected, any models that include fixed transportation cost (FTC) and service level agreements (SLA either in its simple or window version) require more time (as depicted in red and orange), likely due to the additional binary variables and big-M constraints [54]. Including the final inventory deviation (FID) extension seems to interact with FTC and SLA, resulting in an increase of computational time. Nevertheless, this difference is reduced as we increase the number of time periods. In models where FTC is considered and SLAs are not, the computational time decreases independently from all other extensions as shown in green. This implies that, on average, FID only represents an increase in solution time if paired with SLA. Therefore, if the service level agreement extension is not considered, FID (as well as PATP and NID) can be added without significantly impacting the solution time. We conclude that only three different extensions, and their respective interactions, significantly impact the solution time of the model. In this sense, the production across time periods (PATP) and negative inventory deviation (NID, which requires additional binary variables) extensions do not have a major impact on solution time. The model proves to be scalable

with time given that even the largest formulations, which yield hundreds of thousands of binary variables, take no more than 9 minutes to solve on average.

6.3 Disruption Characterization

Effective response optimization and recovery relies on appropriate characterization of the disruption. For example, consider a scenario where a main reactor within Plant 1 is damaged and needs replacement. First, we must evaluate the impact on the plant production capacity, which involves understanding how production will be affected. Relevant information includes: the number of recipes reliant on the damaged reactor, whether there are similar reactors operating in parallel, the existence of alternative routes within the plant can perform the same tasks, and if this disruption implies a plant shut-down. Second, we need to determine the appropriate time-frame considering factors such as the purchasing lead time for a new reactor, the time required to replace the reactor, and how long it will take to restore the plant to nominal production capacity. With this information, we can adjust the production bounds to capture the impact of the disruption.

To illustrate the impact of different disruption characterizations, we obtain the optimal disruption response by solving the optimization model for the various disruption scenarios. The system was solved over 120 periods with a 12-hour time discretization in a fully reactive context with immediate recovery (as discussed in Section 4). Here, we vary plant production bounds and the disruption duration in terms of a fraction. In both scenarios, a fraction equal to zero yields the most adverse situation, which corresponds to either a complete production shutdown in the plant or a perpetually disrupted state. Conversely, a fraction of one reflects the nominal state of the system. Figure 12 offers a view of how these different disruption levels impact both the profit and the number of canceled orders. On the other hand, Figure 13 provides insights into how these disruption levels affect the optimal operation itself. In this context, we analyze the number of shipments needed and the total inventory in the warehouses of the network. It is worth mentioning that these figures represent an overview of the aggregated data for 10 different materials, giving us a big-picture perspective on the solution. However, the model is still calculating optimal inventory and routing for each material individually, ensuring a tailored product-specific approach to managing supply chain disruptions. Note that the color in these figures is proportional to the vertical axis, and it is meant to help visualize the results.

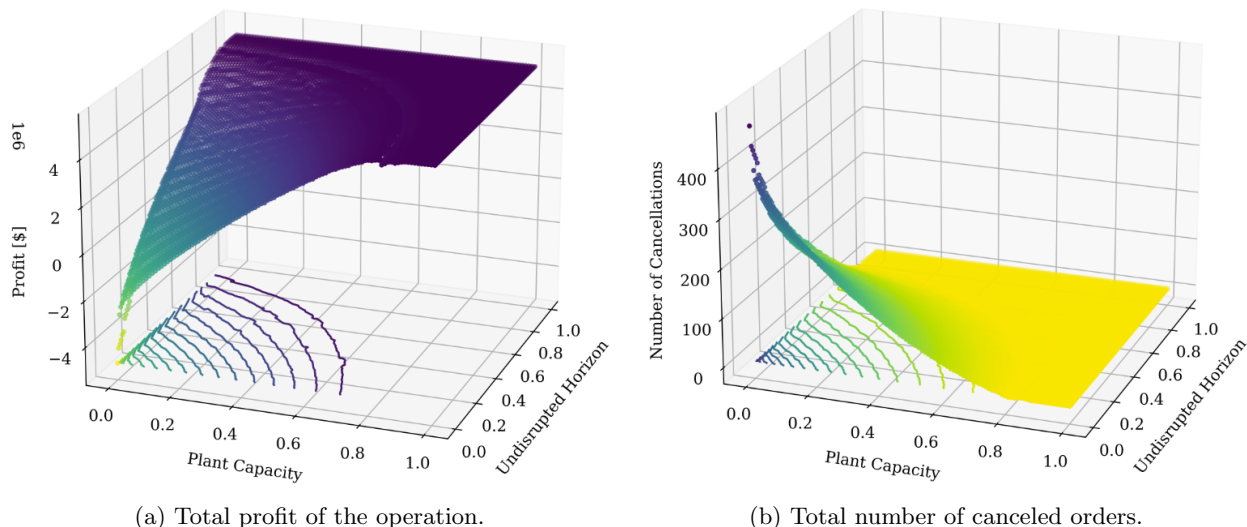


Figure 12: Impacts on the profit of the operation and number of orders canceled for different characterizations of a damaged reactor in Plant 1.

We observe that both subplots depicted in Figure 12, representing the variation in the number of order cancellations and profit, exhibit a monotonic trend, which aligns with our expectation. Here, reduced production capacity together with longer disruption periods results in fewer orders being delivered. Notably,




Figure 13: Impacts on the amount of material flow and warehouse inventory levels for different characterizations of a damaged reactor in Plant 1.

these behaviors are directly and inversely correlated, as any decision to cancel an order incurs a cost (represented by λ^δ) and lowers overall profit. In addition to illustrating the results of an approach, figures like these are useful for decision-makers to analyze the reliability of their supply chain in the face of disruptions (with optimal decision-making).

The subfigures in Figure 13 show the number of shipments and total warehouse inventory levels revealing non-monotonic behavior. This indicates that different disruption characterizations impact the entire system, not just the plant. The dynamic responses are due to the interconnectivity of the network and complex node relationships, facilitating disruption propagation in time and space. Additionally, the subfigures exhibit complementarity; within the mid-range of characterization values, there is a region with a high number of shipments and low warehouse inventory levels. This suggests a rapid inventory turnover policy in warehouses as a reactive response to mitigate disruption effects. While the figures provide aggregated high-level information, further investigation of per-scenario models confirmed that high inventory turnover is an optimal response policy within this range.

Both Figures 12 and 13 exhibit a plateau as we approach the undisrupted conditions (e.g., fraction values close to one in both axes). Here, the model demonstrates the capability to achieve performance levels closely resembling those of an undisrupted operation when either the impact of the disruption on plant production capacity is limited or it does not affect the system for an extended duration. The behavior makes intuitive sense in Figure 12, where near-undisrupted scenarios consistently yield high profits and eliminate the need for order cancellations and delays. However, in the operation plots in Figure 13, a plateau near the values of near-undisrupted characterization is observed, which is unexpected since these models need not behave similarly in aggregation. Additionally, Figure 13a shows the plateau positioned lower than some middle values, indicating this is not due to variables being close to their bounds but rather reflects genuine similarity in optimal schedules under these scenarios. This demonstrates the capability of the model to deliver consistent performance under near-undisrupted conditions despite variations in the operational context.

When analyzing Figures 12 and 13, we observe that except for the origin representing an indefinite plant shutdown, the studied variables change continuously and gradually without abrupt shifts. This consistency holds even in the non-smooth plots in Figure 13, indicating that within a narrow range of disruption characterizations the output variable values remain close. Thus, with a good understanding of expected plant capacity and disruption duration, similar aggregated metrics can be expected. While individual schedules may differ slightly within this range, overall patterns exhibit similarity and continuity.

Given the inherent complexity of supply chain and manufacturing networks, accurate characterization of disruptions is a challenging task. Disruptions can interact, impacting multiple network nodes simultaneously,

making it difficult to precisely measure their effects on the entire system. However, it is crucial to realize that precise characterization is not a requisite for effectively utilizing the proposed model and obtaining value from it. As demonstrated in Section 6.2, the model is computational efficient, enabling the simultaneous evaluation of a multitude of scenarios. For instance, in generating the results presented in this section, which encompassed schedules for a two-month operation, more than 10,200 different scenarios were assessed and run in under 30 minutes of computation time. The scalability of the model allows decision-makers to utilize it as a tool for optimal “simulation”, affording the opportunity to experiment with diverse characterizations and evaluate the resulting schedules. This feature, coupled with the rolling horizon approach, makes the proposed model as a valuable instrument for conducting “what-if” scenarios, facilitating responsive decision-making as information is revealed with time and disruptions become better characterized.

6.4 Late Delivery and Cancellation Cost Sensitivity

As we discussed in Section 3.6, quantifying the impact of order cancellations (λ^δ) and late deliveries (λ^U) is not necessarily straightforward. Beyond the costs alone, these can have an impact on the reputation of the company and customer relations. In this section we conducted a sensitivity analysis that provides insights into how cancellation and late deliveries affect the model, not only in terms of the primary objective function, but also in terms of the variables that affect the goodwill of the company.

Consider a case where an unplanned disruption occurs due to road maintenance, entirely blocking the route from Plant 1 to Warehouse 1 for a full month. In this scenario, the decision-maker is tasked with determining the optimal reactive response for the operations for the next two months in a 12-hour time discretization. Achieving this schedule requires the use of alternative routing, like airfreight, and a smart allocation of inventory downstream to overcome the trucking limitation.

Figures 14a and 14b provide an overview of the effects that varying late delivery and cancellation costs have on the overall profit for this system. Figure 14a shows the total quantity of materials delivered late, and Figure 14b shows the number of orders that had to be canceled. We select the same cost for all materials, costumers, and periods to simplify the analysis, and all materials are aggregated to give a holistic overview of the system.

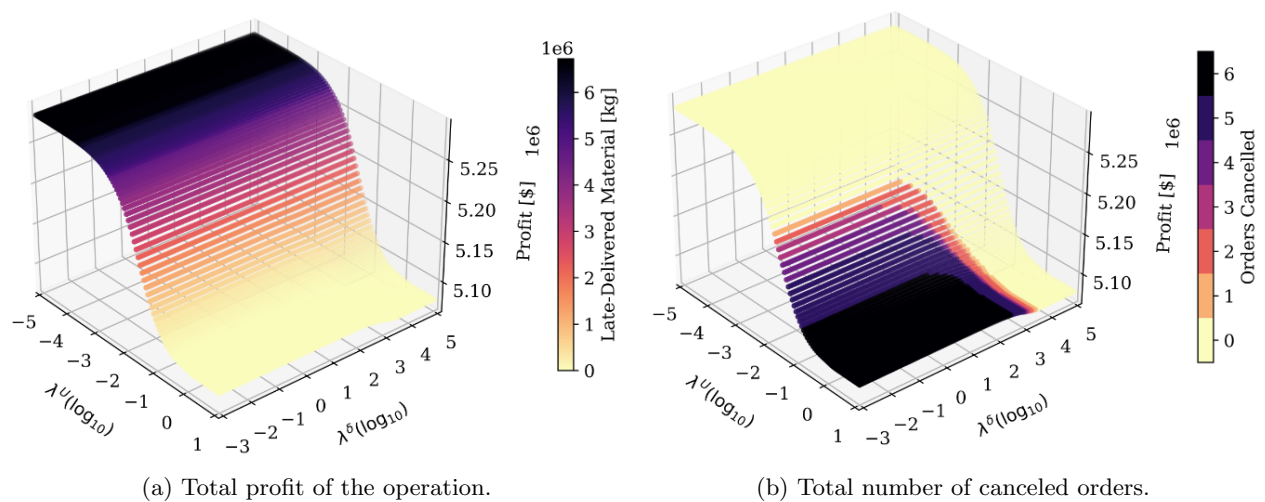


Figure 14: Impacts on the profit of the operation and number of orders canceled for different characterizations of a damaged reactor in Plant 1.

The figures show that the profit of the operation is primarily influenced by the late delivery cost, with the cancellation cost having limited impact. More specifically, while profit values remain relatively constant across various cancellation penalties, there is a distinct reduction in profit as the cost associated with late deliveries increases. Notably, the values become stable on both ends of the λ^U spectrum implying that the range of costs has been sufficiently explored. For low values, the model initially refrained from delivering any of the orders affected by the disruption because delivering late was inexpensive. On the opposite end,

the model decided not to make any late deliveries, even when facing the highest cost penalties, as it had preferred to cancel them.

Figure 14a confirms our intuition about covering the full spectrum of late delivery costs based on the color gradient. The top section is uniformly black, indicating that, for a range of late delivery costs, all affected orders were consistently delivered late. Here, operational adjustments like re-routing, air freight, and production schedule changes outweigh the income from on-time deliveries. The low late delivery penalty leads the model to opt for paying the accumulated late delivery cost rather than delivering on time. In contrast, the lower portion of the plot is flat and light yellow, indicating no late deliveries for these cost values. The model prioritizes cost-effective late deliveries over on-time deliveries based on the trade-off between operational adjustments and associated income.

In the upper part of Figure 14b, the model refrains from canceling orders since delays are not as costly. As delays and cancellations become closer in cost, the model identifies a cost optimal decision that is impacted by both parameters. Overall, these behaviors align with our expectations considering the nature of the cost penalties. With late deliveries being penalized on a per-unit basis, the model retains better control over their allocation. Conversely, the cost of cancellation remains fixed, regardless of the order size, which restricts the flexibility in managing these penalties. This key difference in cost structures influences the strategic choices when dealing with late deliveries and order cancellations in response to supply chain disruptions.

Once the model is solved, it is crucial to evaluate the resulting fill rate within the supply chain, as demonstrated using the color representations. The reason behind this lies in the fact that customer satisfaction is not always precisely encapsulated within the purely economic objective. Instead, we seek a balanced consideration of the trade-off between customer satisfaction and the overall cost of the supply chain.

Conducting a sensitivity analysis can be a useful tool for fine-tuning the cost penalties in order to find the right balance between stressing the system and maintaining customer satisfaction at the desired level. Fortunately, these analyses can be performed rapidly thanks to the demonstrated solution time, which allows for the quick exploration of multiple scenarios in a manageable and practical manner. This adaptive approach ensures that the supply chain is well-prepared to respond to a spectrum of disruptions with a tailored cost strategy.

7 Conclusions and Future Work

In this study, we present a multiperiod mixed-integer linear programming model designed to optimize the operational response of a supply chain and manufacturing network under disruption. The model generates integrated production, shipping, and procurement schedules to mitigate the economic impacts of disruptions. It also incorporates customer management decisions such as on-time or delayed deliveries, the extent of delays, and order cancellations. We extend the model to account for practical supply chain dynamics including time-dependent production, service level agreements, fixed transportation costs, and inventory deviations. Our analysis shows that these extensions have minimal impact on solution time, while identifying which factors and their interactions most affect computational performance.

The model is flexible, allowing for independent or combined disruption scenarios. It bridges gaps in the literature by addressing unexpected demands as disruptions and offering effective solutions. Using a real-life case study from the chemical industry, we demonstrate the ability of the model to provide tractable solutions as the discretization scheme is refined. We also explore the effects of various disruption scenarios and cost penalties for late deliveries and order cancellations showing that, despite significant differences in these factors, the model remains computationally efficient enabling robust “what-if” analyses to guide decision-making.

As a future research direction, we aim to investigate the scalability of the model concerning network size and assess whether different configurations of arbitrary networks yield varied computational performance. Currently, the problem is addressed deterministically, with uncertainty indirectly addressed through multiple deterministic instances in a simulation scheme, capitalizing on the computational efficiency of the model. However, future research will focus on directly incorporating uncertainty by modeling it in the disruption characterization. Moreover, chance constraints will be explored to ensure the enforcement of a minimum service level within this framework.

Acknowledgments and Disclaimers

We gratefully acknowledge the Dow Chemical Company for motivating this work and providing a real-life case study. Parts of this work were presented in the proceedings of the FOCAPO/CPC 2023 Conference.

References

- [1] J. Moosavi, A. M. Fathollahi-Fard, and M. A. Dulebenets, "Supply chain disruption during the COVID-19 pandemic: Recognizing potential disruption management strategies," *International Journal of Disaster Risk Reduction*, vol. 75, p. 102983, 2022.
- [2] D. Ivanov, "Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case," *Transportation Research Part E: Logistics and Transportation Review*, vol. 136, p. 101922, 2020.
- [3] L. V. Snyder and Z.-J. M. Shen, *Fundamentals of Supply Chain Theory*. John Wiley & Sons, 2019.
- [4] F. You, J. M. Wassick, and I. E. Grossmann, "Risk management for a global supply chain planning under uncertainty: models and algorithms," *AIChE Journal*, vol. 55, no. 4, pp. 931–946, 2009.
- [5] H. Elleuch, E. Dafaoui, A. Elmhamedi, and H. Chabchoub, "Resilience and vulnerability in supply chain: Literature review," *IFAC-PapersOnLine*, vol. 49, no. 12, pp. 1448–1453, 2016.
- [6] M. Kamalahmadi and M. M. Parast, "A review of the literature on the principles of enterprise and supply chain resilience: Major findings and directions for future research," *International Journal of Production Economics*, vol. 171, pp. 116–133, 2016.
- [7] S. R. Cardoso, A. P. Barbosa-Póvoa, S. Relvas, and A. Q. Novais, "Resilience metrics in the assessment of complex supply-chains performance operating under demand uncertainty," *Omega*, vol. 56, pp. 53–73, 2015.
- [8] B. Adenso-Diaz, C. Mena, S. García-Carbajal, and M. Liechty, "The impact of supply network characteristics on reliability," *Supply Chain Management: An International Journal*, vol. 17, no. 3, pp. 263–276, 2012.
- [9] S. G. Azevedo, H. Carvalho, and V. Cruz-Machado, "LARG index: A benchmarking tool for improving the leanness, agility, resilience and greenness of the automotive supply chain," *Benchmarking: An International Journal*, vol. 23, no. 6, pp. 1472–1499, 2016.
- [10] A. Munoz and M. Dunbar, "On the quantification of operational supply chain resilience," *International Journal of Production Research*, vol. 53, no. 22, pp. 6736–6751, 2015.
- [11] V. Dixit, N. Seshadrinath, and M. Tiwari, "Performance measures based optimization of supply chain network resilience: A NSGA-II+ Co-Kriging approach," *Computers & Industrial Engineering*, vol. 93, pp. 205–214, 2016.
- [12] K. Katsaliaki, P. Galetsi, and S. Kumar, "Supply chain disruptions and resilience: A major review and future research agenda," *Annals of Operations Research*, pp. 1–38, 2021.
- [13] J. P. Ribeiro and A. Barbosa-Povoa, "Supply Chain Resilience: Definitions and quantitative modelling approaches—A literature review," *Computers & Industrial Engineering*, vol. 115, pp. 109–122, 2018.
- [14] S. H. Owen and M. S. Daskin, "Strategic facility location: A review," *European Journal of Operational Research*, vol. 111, no. 3, pp. 423–447, 1998.
- [15] L. V. Snyder, "Facility location under uncertainty: A review," *IIE Transactions*, vol. 38, no. 7, pp. 547–564, 2006.
- [16] L. V. Snyder and M. S. Daskin, "Reliability models for facility location: the expected failure cost case," *Transportation Science*, vol. 39, no. 3, pp. 400–416, 2005.
- [17] P. Tsiakis, N. Shah, and C. C. Pantelides, "Design of multi-echelon supply chain networks under demand uncertainty," *Industrial & Engineering Chemistry Research*, vol. 40, no. 16, pp. 3585–3604, 2001.

- [18] F. You and I. E. Grossmann, “Mixed-integer nonlinear programming models and algorithms for large-scale supply chain design with stochastic inventory management,” *Industrial & Engineering Chemistry Research*, vol. 47, no. 20, pp. 7802–7817, 2008.
- [19] M. Lim, M. S. Daskin, A. Bassamboo, and S. Chopra, “A facility reliability problem: Formulation, properties, and algorithm,” *Naval Research Logistics (NRL)*, vol. 57, no. 1, pp. 58–70, 2010.
- [20] T. Cui, Y. Ouyang, and Z.-J. M. Shen, “Reliable facility location design under the risk of disruptions,” *Operations Research*, vol. 58, no. 4-part-1, pp. 998–1011, 2010.
- [21] Z.-J. M. Shen, R. L. Zhan, and J. Zhang, “The reliable facility location problem: Formulations, heuristics, and approximation algorithms,” *INFORMS Journal on Computing*, vol. 23, no. 3, pp. 470–482, 2011.
- [22] X. Li and Y. Ouyang, “A continuum approximation approach to reliable facility location design under correlated probabilistic disruptions,” *Transportation Research Part B: Methodological*, vol. 44, no. 4, pp. 535–548, 2010.
- [23] P. Garcia-Herreros, J. M. Wassick, and I. E. Grossmann, “Design of resilient supply chains with risk of facility disruptions,” *Industrial & Engineering Chemistry Research*, vol. 53, no. 44, pp. 17 240–17 251, 2014.
- [24] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro, “A stochastic programming approach for supply chain network design under uncertainty,” *European Journal of Operational Research*, vol. 167, no. 1, pp. 96–115, 2005.
- [25] W. Klibi and A. Martel, “Scenario-based supply chain network risk modeling,” *European Journal of Operational Research*, vol. 223, no. 3, pp. 644–658, 2012.
- [26] W. Klibi, A. Martel, and A. Guitouni, “The design of robust value-creating supply chain networks: A critical review,” *European Journal of Operational Research*, vol. 203, no. 2, pp. 283–293, 2010.
- [27] P. Glasserman and S. Tayur, “Sensitivity analysis for base-stock levels in multiechelon production-inventory systems,” *Management Science*, vol. 41, no. 2, pp. 263–281, 1995.
- [28] L. V. Snyder, Z. Atan, P. Peng, Y. Rong, A. J. Schmitt, and B. Sinsoysal, “OR/MS models for supply chain disruptions: A review,” *IIE Transactions*, vol. 48, no. 2, pp. 89–109, 2016.
- [29] F. Chen and Y.-S. Zheng, “Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand,” *Management Science*, vol. 40, no. 10, pp. 1262–1275, 1994.
- [30] A. J. Clark and H. Scarf, “Optimal policies for a multi-echelon inventory problem,” *Management science*, vol. 6, no. 4, pp. 475–490, 1960.
- [31] K. H. Shang and J.-S. Song, “Newsvendor bounds and heuristic for optimal policies in serial supply chains,” *Management Science*, vol. 49, no. 5, pp. 618–638, 2003.
- [32] C. Kim, J. Jun, J. Baek, R. Smith, and Y.-D. Kim, “Adaptive inventory control models for supply chain management,” *The International Journal of Advanced Manufacturing Technology*, vol. 26, no. 9, pp. 1184–1192, 2005.
- [33] K. Rosling, “Optimal inventory policies for assembly systems under random demands,” *Operations Research*, vol. 37, no. 4, pp. 565–579, 1989.
- [34] A. J. Schmitt, S. A. Sun, L. V. Snyder, and Z.-J. M. Shen, “Centralization versus decentralization: Risk pooling, risk diversification, and supply chain disruptions,” *Omega*, vol. 52, pp. 201–212, 2015.
- [35] C. C. Sherbrooke, “METRIC: A multi-echelon technique for recoverable item control,” *Operations Research*, vol. 16, no. 1, pp. 122–141, 1968.

- [36] Y. Rong, Z. Atan, and L. V. Snyder, “Heuristics for base-stock levels in multi-echelon distribution networks,” *Production and Operations Management*, vol. 26, no. 9, pp. 1760–1777, 2017.
- [37] A. S. Eruguz, E. Sahin, Z. Jemai, and Y. Dallery, “A comprehensive survey of guaranteed-service models for multi-echelon inventory optimization,” *International Journal of Production Economics*, vol. 172, pp. 110–125, 2016.
- [38] V. Achkar, B. Brunaud, R. Musa, and I. Grossmann, “Efficient computational strategies for a mathematical programming model for multi-echelon inventory optimization based on the guaranteed-service approach,” *Computers & Chemical Engineering*, vol. 182, p. 108567, 2024.
- [39] N. Shah, “Process industry supply chains: Advances and challenges,” *Computers & Chemical Engineering*, vol. 29, no. 6, pp. 1225–1235, 2005.
- [40] E. N. Pistikopoulos, A. Barbosa-Povoa, J. H. Lee, R. Misener, A. Mitsos, G. V. Reklaitis, V. Venkatasubramanian, F. You, and R. Gani, “Process systems engineering—the generation next?” *Computers & Chemical Engineering*, vol. 147, p. 107252, 2021.
- [41] I. E. Grossmann, “Advances in mathematical programming models for enterprise-wide optimization,” *Computers & Chemical Engineering*, vol. 47, pp. 2–18, 2012.
- [42] D. J. Garcia and F. You, “Supply chain design and optimization: Challenges and opportunities,” *Computers & Chemical Engineering*, vol. 81, pp. 153–170, 2015.
- [43] E. Perea-Lopez, B. E. Ydstie, and I. E. Grossmann, “A model predictive control strategy for supply chain optimization,” *Computers & Chemical Engineering*, vol. 27, no. 8-9, pp. 1201–1218, 2003.
- [44] O. Badejo and M. Ierapetritou, “Mathematical programming approach to optimize tactical and operational supply chain decisions under disruptions,” *Industrial & Engineering Chemistry Research*, vol. 61, no. 45, pp. 16 747–16 763, 2022.
- [45] R. Raman and I. E. Grossmann, “Modelling and computational techniques for logic based integer programming,” *Computers & Chemical Engineering*, vol. 18, no. 7, pp. 563–578, 1994.
- [46] I. E. Grossmann and F. Trespalacios, “Systematic modeling of discrete-continuous optimization models through generalized disjunctive programming,” *AIChE Journal*, vol. 59, no. 9, pp. 3276–3295, 2013.
- [47] C. T. Maravelias and I. E. Grossmann, “New general continuous-time State- Task network formulation for short-term scheduling of multipurpose batch plants,” *Industrial & Engineering Chemistry Research*, vol. 42, no. 13, pp. 3056–3074, 2003.
- [48] E. Kondili, C. C. Pantelides, and R. W. Sargent, “A general algorithm for short-term scheduling of batch operations—I. MILP formulation,” *Computers & Chemical Engineering*, vol. 17, no. 2, pp. 211–227, 1993.
- [49] J. Balasubramanian and I. Grossmann, “Approximation to multistage stochastic optimization in multi-period batch plant scheduling under demand uncertainty,” *Industrial & engineering chemistry research*, vol. 43, no. 14, pp. 3695–3713, 2004.
- [50] J. B. Rawlings, D. Q. Mayne, M. Diehl *et al.*, *Model predictive control: Theory, computation, and design*. Nob Hill Publishing Madison, WI, 2017, vol. 2.
- [51] Z. Chong and C. L. Swartz, “Model-based control of multi-unit systems under partial shutdown conditions,” in *2009 American Control Conference*. IEEE, 2009, pp. 160–165.
- [52] R. M. Lima, I. E. Grossmann, and Y. Jiao, “Long-term scheduling of a single-unit multi-product continuous process to manufacture high performance glass,” *Computers & Chemical Engineering*, vol. 35, no. 3, pp. 554–574, 2011.

- [53] K. Subramanian, J. B. Rawlings, C. T. Maravelias, J. Flores-Cerrillo, and L. Megan, “Integration of control theory and scheduling methods for supply chain management,” *Computers & Chemical Engineering*, vol. 51, pp. 4–20, 2013.
- [54] F. Trespalacios and I. E. Grossmann, “Improved Big-M reformulation for generalized disjunctive programs,” *Computers & Chemical Engineering*, vol. 76, pp. 98–103, 2015.
- [55] F. Ceccon, J. Jalving, J. Haddad, A. Thebelt, C. Tsay, C. D. Laird, and R. Misener, “OMLT: Optimization & machine learning toolkit,” *The Journal of Machine Learning Research*, vol. 23, no. 1, pp. 15 829–15 836, 2022.
- [56] M. L. Bynum, G. A. Hackebeil, W. E. Hart, C. D. Laird, B. L. Nicholson, J. D. Siirola, J.-P. Watson, and D. L. Woodruff, *Pyomo–optimization modeling in python*, 3rd ed. Springer Science & Business Media, 2021, vol. 67.

A Appendix: Nomenclature

This section provides a comprehensive nomenclature of all sets, parameters, and variables used in the base model and its proposed extensions.

A.1 Sets

- Nodes ($\mathcal{N} := \mathcal{S} \cup \mathcal{P} \cup \mathcal{W} \cup \mathcal{C}$)
 - Suppliers (\mathcal{S})
 - Plants (\mathcal{P})
 - Warehouses (\mathcal{W})
 - Customers (\mathcal{C})
 - Suppliers that require service level agreements ($\mathcal{S}^{SLA} \subseteq \mathcal{S}$)
- Arcs ($\mathcal{A} := (\mathcal{N} \times \mathcal{N})$)
 - Arcs coming into node $n \in \mathcal{N}$ ($\mathcal{A}_n^{In} := \{(i, j) \in \mathcal{A} : j = n\}$)
 - Arcs coming out of node $n \in \mathcal{N}$ ($\mathcal{A}_n^{Out} := \{(i, j) \in \mathcal{A} : i = n\}$)
- Materials (\mathcal{M})
 - Materials that can exist at node $n \in \mathcal{N}$ (\mathcal{M}_n)
 - Materials that can move through at arc $a \in \mathcal{A}$ (\mathcal{M}_a)
- Time periods (\mathcal{T})
- Recipes that occur in plant $p \in \mathcal{P}$ (\mathcal{R}_p)

A.2 Continuous Variables ($\in \mathbb{R}_*^+$)

- F_{mat}^{In} : Flow of material $m \in \mathcal{M}_a$ that is going into arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- F_{mat}^{Out} : Flow of material $m \in \mathcal{M}_a$ that is going out of arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- P_{prt} : Production of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ during time period $t \in \mathcal{T}$.
- I_{mnt} : Inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$.
- B_{mst} : Amount to buy of material $m \in \mathcal{M}_s$ from supplier $s \in \mathcal{S}$ during time period $t \in \mathcal{T}$.
- D_{mct} : Demand met of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- U_{mct} : Unmet demand met of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- P_{prt}^{In} : Production of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ starting at time period $t \in \mathcal{T}$.
- P_{prt}^{Out} : Production of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ concluding at time period $t \in \mathcal{T}$.
- $Deviation_{mn}$: Absolute value of difference between the initial and final inventories of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$.
- K_{mnt} : Distance between the inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ and the threshold value I_{mnt}^L at time period $t \in \mathcal{T}$. This distance is also referred in the paper as negative inventory deviation.

A.3 Binary Variables ($\in \{0, 1\}$)

- y_{mct} : Indicates if the order of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$ is canceled ($y_{mct} = 1$) or not ($y_{mct} = 0$).
- x_{mat} : Indicates the existence of flow of material $m \in \mathcal{M}_a$ going into arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- w_{mst} : Determines if the service level agreement of material $m \in \mathcal{M}_s$ and supplier $s \in \mathcal{S}^{SLA}$ is executed at time period $t \in \mathcal{T}$ or not.
- z_{mnt} : Indicates that the inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ is below I_{mnt}^L at time period $t \in \mathcal{T}$.

A.4 Parameters

- ϕ_{rm} : Mass production (+) or consumption (-) of material $m \in \mathcal{M}_p$ in for recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ during time period $t \in \mathcal{T}$.
- δ_{mct} : Order of material $m \in \mathcal{M}_c$ placed by customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- SS_{mnt} : Buffer stock of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$.
- α_{mnt} : Fraction of the buffer stock of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$ that should be enforced regardless of the disruption.
- B_{mst}^{SLA} : Minimum purchase amount stated by the service level agreement of material $m \in \mathcal{M}_s$ and supplier $s \in \mathcal{S}^{SLA}$ executed at time period $t \in \mathcal{T}$.

A.4.1 Variable bounds

- $[F_{mat}^L, F_{mat}^U]$: Lower and upper bounds for the flow of material $m \in \mathcal{M}_a$ that traversing arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- $[P_{prt}^L, P_{prt}^U]$: Lower and upper bounds for the production of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ during time period $t \in \mathcal{T}$.
- $[I_{mnt}^L, I_{mnt}^U]$: Lower and upper bounds for the inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$.
- $[B_{mst}^L, B_{mst}^U]$: Lower and upper bounds for the amount to buy of material $m \in \mathcal{M}_s$ from supplier $s \in \mathcal{S}$ during time period $t \in \mathcal{T}$.
- $[D_{mct}^L, D_{mct}^U]$: Lower and upper bounds for the demand met of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- $[U_{mct}^L, U_{mct}^U]$: Lower and upper bounds for the unmet demand met of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.

A.4.2 Initial conditions

- $F_{ma}^{In,0}$: Pre-planned flow of material $m \in \mathcal{M}_a$ that is going into arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$ in absence of disruption.
- $F_{ma}^{Out,0}$: Pre-planned flow of material $m \in \mathcal{M}_a$ that is going out of arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$ in absence of disruption.
- P_{pr}^0 : Pre-planned production of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ during time period $t \in \mathcal{T}$ in absence of disruption.

- I_{mn}^0 : Existent inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$ at the time disruption occurs.
- B_{ms}^0 : Pre-planned amount to buy of material $m \in \mathcal{M}_s$ from supplier $s \in \mathcal{S}$ during time period $t \in \mathcal{T}$ in absence of disruption.
- D_{mc}^0 : Pre-planned demand met of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$ in absence of disruption.

A.4.3 Time durations (τ)

- τ_{mat} : Time it takes material $m \in \mathcal{M}_a$ to traverse arc $a \in \mathcal{A}$ when shipped at time period $t \in \mathcal{T}$.
- τ_{prt}^P : Time it takes to execute recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ starting at time period $t \in \mathcal{T}$.
- τ_{mst}^{SLA} : Time of the window to fulfill a service level agreement of material $m \in \mathcal{M}_s$ to supplier $s \in \mathcal{S}^{SLA}$ for a contract that started at time period $t \in \mathcal{T}$.

A.4.4 Economic parameters (λ)

- λ_{mct}^D : Price per unit of material $m \in \mathcal{M}_c$ sold to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- λ_{mct}^U : Cost penalty per unit of unmet demand of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- λ_{mct}^δ : Cost penalty of canceling an order of material $m \in \mathcal{M}_c$ to customer $c \in \mathcal{C}$ during time period $t \in \mathcal{T}$.
- λ_{mat}^F : Cost of transporting a unit material $m \in \mathcal{M}_a$ through arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- λ_{mst}^B : Cost of buying a unit material $m \in \mathcal{M}_s$ from supplier $s \in \mathcal{S}$ during time period $t \in \mathcal{T}$.
- λ_{prt}^P : Cost of producing a unit of recipe $r \in \mathcal{R}_p$ in plant $p \in \mathcal{P}$ during time period $t \in \mathcal{T}$.
- λ_{mnt}^I : Cost of holding a unit material $m \in \mathcal{M}_n$ as inventory in node $n \in \mathcal{P} \cup \mathcal{W}$ during time period $t \in \mathcal{T}$.
- λ_{mat}^{Fix} : Fixed cost of transporting material $m \in \mathcal{M}_a$ through arc $a \in \mathcal{A}$ during time period $t \in \mathcal{T}$.
- λ_{mn}^{Dev} : Penalty for the deviation between the initial and final inventories of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$.
- λ_{mnt}^K : Penalty for the deviation between the inventory of material $m \in \mathcal{M}_n$ in node $n \in \mathcal{P} \cup \mathcal{W}$ and the threshold value I_{mnt}^L during time period $t \in \mathcal{T}$.