Optimizing the Recovery of Disrupted Multi-Echelon Assembly Supply Chain Networks

Huy Nguyen* Thomas C. Sharkey*† John E. Mitchell* Willam A. Wallace*

February 21, 2019

Abstract

We consider optimization problems related to the scheduling of multi-echelon assembly supply chain (MEASC) networks that have applications in the recovery from large-scale disruptive events. Each manufacturer within this network assembles a component from a series of sub-components received from other manufacturers We develop scheduling decision rules that are applied locally at each manufacturer and are proven to optimize two industry-relevant *global* recovery metrics: (i) minimizing the maximum tardiness of any order of the final product of the MEASC network (ii) and minimizing the time to recover from the disruptive event. Our approaches are applied to a data set based upon an industrial partner's supply chain to show their applicability as well as their advantages over integer programming models. The developed decision rules were proven to be optimal, faster, and more robust than the equivalent IP formulations. In addition, they provide conditions under which local manufacturer decisions will lead to globally optimal recovery efforts. These decision rules can help managers to make better production and shipping decisions to optimize the recovery after disruptions and quantitatively test the impact of different pre-event mitigation strategies against potential disruptions. They can be further useful in MEASCs with or expecting a large amount of backorders.

*Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute
†Corresponding Author, sharkt@rpi.edu

1 Introduction

In recent years, supply chain managers have become increasingly concerned with disruptions to their systems. This is because over the past 20 years, there have been many large-scale disruptive events, both man-made and natural, such as the September 11th terrorist attacks in 2001, Hurricane Katrina in 2005, the 2011 Japan earthquake and tsunami, and Hurricane Sandy in 2012. In 2011, the Japan earthquake and tsunami generated 300 billion dollars in damage, forced all Japan's ports to be closed, and caused the decline of Japanese auto manufacturers' production and market share for several months (CNN Library 2011).

Many manufacturing supply chains are complex with many tiers and various parts manufacturers, some of which are the sole-source manufacturers for certain parts. We are particularly interested in supply chains that can be classified as multi-echelon assembly supply chain (MEASC) networks. In a MEASC network, each supply node produces a component part from a series of components (including, in some cases, raw material). That part is then delivered and used as a component for an upper tier manufacturer to assemble into a product. At the final assembly, all major components are assembled to make the final product of the MEASC network. It should be noted that while there are studies in the literature (Snyder et al. 2016) on assembly supply chain networks, they often restrict the networks to have a fixed number of echelons (usually two), or to have tree network structures (i.e., a component manufacturer only sends its part to one upper tier assembler). These studies will be discussed in Section 2. Our classification of MEASC networks is not subject to those restrictions and, therefore, has more general industrial applications.

The challenge for this type of supply chain is that the end product is assembled from numerous components with widely varying lead times, geographical locations, bill of materials and costs. MEASC networks are often used in the manufacturing of electronics, appliances, automobiles, aircrafts, and consumer packaged goods. Single-sourced components within MEASC networks are especially important in industries with high manufacturer qualifications, such as defense aircraft manufacturing, or with strict regulations, such as biopharmaceutical manufacturing. Therefore, after a disruptive event, these MEASC networks are more vulnerable (due to the lack of alternatives for components) to cascading disruptions throughout the system when a subset of manufacturers are impacted by a large-scale event (see Figure 1 for illustration).

A disruption does not necessarily have to occur at the location of the final assembly node to induce delay to the production. For instance, the Japan disaster in 2011 not only disrupted Toyota's supplies of made-in-

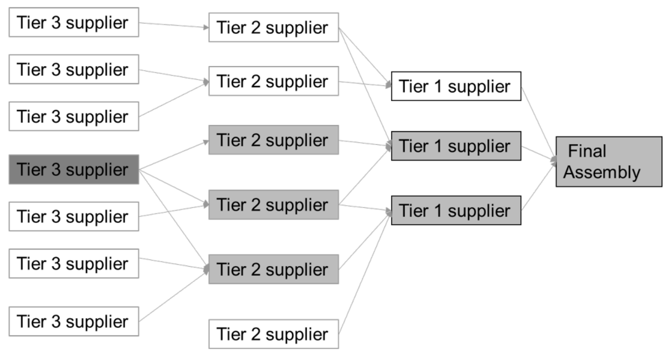

Figure 1: The impact of a disrupted manufacturer on an example MEASC network. An arc from manufacturer i to manufacturer j represents that manufacturer j requires the component produced by manufacturer i. The dark gray Tier 3 manufacturer is disrupted and this disruption will 'cascade' to all light grey manufacturers.

Japan models like the Toyota Prius but also diminished the flow of parts for some cars assembled in the U.S. After re-examining their supply chain, Toyota discovered that while "they may have attempted to mitigate disruption risk by multi-sourcing with Tier 1 manufacturers, many Tier 1 manufacturers shared some or all of their Tier 2 manufacturers, leading to a significant degree of correlation in risk for supply coming from Tier 1" (Ang, Iancu, and Swinney 2016). Like many other manufacturing firms, Toyota has very limited visibility beyond their Tier 1 manufacturers. Toyota's Tier 2 manufacturers can range from large firms with specific expertise like automotive semiconductor manufacturer Renesas to smaller firms with generic capabilities such as the rubber manufacturer Fujikura (Ang, Iancu, and Swinney 2016). Production shortages and delays can be costly for the manufacturer. In 2012, Boeing lost an estimated $100 million for each 787 it sold due to compensation to buyers for delays and cost overruns (Ostrower 2012). Lastly, the previous examples focused on companies whose products allow for multi-sourcing the same component which is not necessarily possible (or cost-effective if it is possible) for MEASC networks with high manufacturer qualification standards. For example, the manufacturing of defense aircraft may require that parts are single-sourced due to the confidential nature of the production process. This is also true for biopharmaceutical manufacturing where regulatory agencies require the production process to be the exact same for each run of the product (and, therefore, its components).

Hence, there is the need to improve the resiliency of a supply chain and, in particular, reducing the delays, or tardiness, of orders and the time to recover of a MEASC network after a disruption. This motivates the analysis of new classes of scheduling problems in MEASC networks. These two important objectives were discussed with our industrial partner and both have merits analyzing the post-event recovery of their networks. In the context of our new MEASC scheduling problems, the corresponding objectives for these goals are minimizing the maximum tardiness of the MEASC network in meeting demand for its final product and minimizing the time it takes the MEASC network to recover from a disruption.

The results of this paper are the first to optimize the two global objectives of the MEASC network by solving the *local* scheduling problems faced by the manufacturers. In addition to the application in disruption recovery planning, the problem can be configured for non-disrupted supply chains to optimize those objectives in normal operation. This can be suitable in a setting where a company is currently dealing with or expecting a large amount of backorders. This approach is applicable in practice since the individual manufacturers within the MEASC network retain their autonomy and it is not assumed that they follow a single centralized decision-maker in the recovery process. This local scheduling problem is unique in the traditional scheduling literature (see (Pinedo 2012)) since a set of sub-tasks (corresponding to the set of paths from the manufacturer to the final assembler) all need to be completed in order for a 'main' task (i.e., all sub-components from that manufacturer required for a product) to be considered complete at the manufacturer. This paper provides decision rules that solve these local scheduling problems for both the maximum tardiness objective and the time to recover objective. Our results then prove that the performance of the *global* recovery problem faced by the MEASC network will be optimized should all the manufacturers solve their local scheduling problems with these decision rules. Furthermore, very little previous work has been done on disruptions and scheduling rules in MEASC networks. Therefore, this paper fills a critical gap in the research by providing optimization algorithms/decision rules to dynamically recover the MEASC network from a disruption, which is a key aspect of supply chain resiliency.

The rest of the paper is organized as follows. Section 2 provides an overview of related literature. Section 3 provides the problem statement. Section 4 describes the decision rule for the maximum tardiness metric and proves that applying it locally leads to the optimal global recovery. Section 5 presents the decision rule for the time to recover objective and provides an overview of its proof of global optimality. Detailed proofs are provided in the supplemental material. Section 6 shows our testing and potential uses of the recovery problems, including the robustness of our rules to errors in the estimates of the disruption.

2 Literature Review

Resiliency is often defined as the *ability to withstand and rapidly recover from unexpected, disruptive events* (see The White House (2013)). The overwhelming majority of the research done on supply chain disruptions (see Snyder et al. (2016) for overviews) focuses on the *ability to withstand* and neglects the *rapidly recover* part of this definition.

2.1 Supply Chain Resiliency

The research on supply chain resiliency (Tukamuhabwa et al. (2015)) tends to focus on how pre-event strategies such as flexibility, redundancy, agility, and collaboration impact operations after the disruptive event without necessarily modeling the direct recovery of the supply chain. The focus of supply chain risk management (e.g., Tang (2006), Tomlin (2006), Ho et al. (2015)) is on identifying, monitoring, assessing, and mitigating the risk of disruptive events (broadly defined) on supply chains which is clearly related to supply chain resiliency, yet does not address the recovery of the network from the disruptive event.

There has been related work on infrastructure resiliency (e.g., Bruneau et al. (2003)). However, there is a fundamental difference between these areas in that unmet demand within the MEASC network must be met in the future by the network (e.g., if power is out to a demand point for a week, we do not need to meet twice the power the next week but if we miss the delivery of a final product in a supply chain network, we will need to produce it in the future). Therefore, the MEASC network cannot be considered recovered until it fulfills all unmet demand, i.e., back-orders, that resulted from the disruption whereas an infrastructure network would be considered recovered once it returns to its normal level of functionality.

2.2 Supply Chain and Network Recovery

There has been recent research on network restoration/recovery that is related to supply chain recovery. Çelik (2016) provides an overview of this literature that includes discussions on *integrated network design and scheduling* (INDS) problems (e.g., Nurre et al. (2012)), incremental network design problems (e.g., Baxter et al. (2014)), and problems focused on connectivity between supply and demand nodes (e.g., Averbakh and Pereira (2015), Celik, Ergun, and Keskinocak (2015)). However, this previous research on network restoration cannot be readily applied to MEASC networks since the models assume that flows (or goods) do not change as they are processed by nodes within the network. In particular, these papers would model the supply chain as a network where goods flow through *any* path in making their way from a supply node to a demand node. This is quite different than MEASC networks where components at manufacturer node i must flow through *every* path from i to the final manufacturer B in order to meet the demand within the network. Therefore, it is necessary to examine new recovery models for MEASC networks.

On the disruption side of the supply chain resilience literature, the vast majority of the research examines either (i) the immediate impact of the disruption, i.e., how well will the supply chain operate during the time the disrupted manufacturers are non-operational, or (ii) small-scale disruptions that occur somewhat frequently and impact an individual manufacturer. Snyder et al. (2016) provides a review of modeling supply chain disruptions which includes strategies on how to mitigate their impact (which overlap with supply chain risk management strategies). While there has been research on disruptions to multi-echelon supply chains, it typically focuses on distribution systems and thus is not applicable to the assembly aspects of MEASC networks. Rong, Bulut, and Snyder (2009), Atan and Snyder (2012), Schmitt and Snyder (2012) and Schmitt et al. (2015) work on heuristics and approximations to solve for the base-stock levels which minimize the total system cost for one warehouse, multiple retailers systems under disruptions. Bimpikus, Fearing, and Tahbaz-Salehi (2013) use a simple example with two tiers and four manufacturers to show that optimizing procurement (sourcing) decisions in each tier of the supply chain in isolation may not be sufficient for mitigating risks at an aggregate level. Rong, Atan, and Snyder (2014) develop two heuristics to approximate the optimal base-stock levels for local distributors in an infinite-horizon problem for multi-echelon, continuous-review distribution systems. While these works are relevant for the analysis of multi-echelon systems, they are primarily applicable only to distribution networks rather than MEASC networks.

There is very little literature on disruptions to MEASC networks. Some relevant works to MEASC networks is on the impact of inventory decisions on mitigating disruptions (DeCroix (2013), He (2014)) where the disrupted state of manufacturers follows a stochastic process (i.e., the state of the manufacturer is a random variable that transitions between operational and disrupted states). Therefore, these results are not applicable to the large-scale disruptions that are the motivation of this research. Our research addresses the recovery from a large-scale event impacting a number of manufacturers (i.e., those in a particular region vulnerable to the same hazard). DeCroix (2013) focuses on minimizing the expected discounted cost of operating the system over an infinite horizon. This approach requires the MEASC networks to be trees, which is only applicable to an assembly system where every manufacturer can only supply at most one upper tier assembler (i.e., a manufacturer has only one path from it to the final demand node). This is also a common assumption of the assembly and assemble-to-order literature (e.g., Schmidt and Nahmias (1985), Yao (1988), Rosling (1989), Gerchak, Wang, and Yano (1994), Gurnani, Akella, and Lehoczky (2000), Song and Yao (2002), Lu, Song, and Yao (2003), Song and Zipkin (2003), Bollapragada, Rao, and Zhang (2004), Benjaafar and ElHafsi (2006), Benjaafar et al. (2011), Ceryan, Duenyas, and Koren (2012) and Ji, Wang, and Hu (2016)).

Our work is applicable to a broader range of MEASC networks in that we are considering situations where a manufacturer node may have multiple paths from it to the final demand node. Our research only requires that the MEASC network is acyclic, which is a more realistic assumption for an assembly network, especially those of our industrial partner (see Table 2 for the average out-degrees of manufacturers where previous work assumes the out-degree of all manufacturers is 1). In addition, DeCroix (2013) only considers disruptions occurring at a manufacturer level where a disrupted manufacturer has a probability of recovering before the next planning period and an operational manufacturer has a probability of becoming disrupted. After a disruption, meeting deadlines for products is critical in the recovery of the supply chain system. Therefore, minimizing the tardiness of demand becomes important. A recent model has considered a tardiness-based objective for emergency relief supply chains. Lei, Lee, and Dong (2016) considers a problem of assembling emergency supply kits for distribution during a large-scale emergency where the components of the kit can come from one of multiple manufacturers. Their work only considers a network with a single tier of manufacturers, a single tier of distribution centers (that assembles kits from components), and a single tier of demand points. This restriction does not allow their model or heuristics to be extendable to multiple echelons. Therefore, there is a significant opportunity for research on the resilience of MEASC networks, particularly with the objectives related to tardiness in meeting the demand of final products.

The most significant related works to our paper are Simchi-Levi, Schmidt, and Wei (2014), Simchi-Levi et al. (2015) and Gao et al. (2016) where they create mathematical programming models to optimize financial loss and time to survival for MEASC networks. Simchi-Levi, Schmidt, and Wei (2014), Simchi-

Levi et al. (2015) assume that at most one manufacturer in the supply chain can be disrupted. While these models are useful in identifying critical nodes in the network, the impact of large scale disruptions should be considered when multiple manufacturers are disrupted in the network. Gao et al. (2016) extends the models to incorporate simultaneous disruptions. There are some differences in term of methods and applications between this research and the models proposed in Gao et al. (2016).

While Simchi-Levi, Schmidt, and Wei (2014), Simchi-Levi et al. (2015) and Gao et al. (2016) incorporate the notion of *time to recover* (TTR) at the individual manufacturer's level, there are a few key differences in how we define TTR. At a node level, Simchi-Levi, Schmidt, and Wei (2014), Simchi-Levi et al. (2015) define TTR as 'the time it takes for a node to recover to full functionality after a disruption'. This is similar to the parameter u_i, the time manufacturer i is restored, in our model and *is not impacted* by the decisions of the MEASC network. Our TTR not only accounts for the time the supply chain returns to its normal performance level but also until it fulfills all unmet demand, i.e. back-orders, that resulted from the disruption. Our TTR definition is important in the situation where all the back-orders must be met, such as the case for aircraft manufacturing where orders are agreed upon years before the delivery. The objectives of our decision rules focus on tardiness and the TTR of the MEASC network. Their models' objectives are financial metrics rather than the TTR or the tardiness of the MEASC networks as a whole. Our approach offers different insights about the network resilience, especially to managers of more time-sensitive supply chains that are under critical contracts to deliver the final products (e.g., defense aircraft manufacturing). In addition, supply chain managers can utilize both the models in Gao et al. (2016) and the decision rules in this paper to evaluate the tradeoffs between these objectives.

In terms of our methodological contributions compared to Simchi-Levi, Schmidt, and Wei (2014), Simchi-Levi et al. (2015) and Gao et al. (2016), this research proposes decision rules that are polynomial time algorithms with theoretical proofs of optimality. This gives a better computational time and performance guarantee, which is especially important in the context of disruption management because of the urgency of the matter. In addition to the application in disruption recovery planning, the problem described in this paper can be applied to non-disrupted MEASC networks to optimize our objectives under normal operations. This can be suitable in a situation where the MEASC network is experiencing a large amount of demand that could lead to backorders. Therefore, these decision rules are a valuable contribution to the acyclic MEASC scheduling literature.

2.3 Acyclic Supply Chain Scheduling

There has been recent research on planning and scheduling that may be applicable to acyclic MEASC networks. These models often utilize meta-heuristics such as genetic algorithms (e.g., Jiang, Yang, and Li (2016)) or mathematical programming (Yao and Liu (2009), Sarrafha et al. (2015), Sawik (2016)). While these models are useful in a variety of applications, our decision rules can provide significantly better performance in terms of optimality guarantees, runtime, and memory usage. Furthermore, since manufacturers can carry out these decision rules locally, limited knowledge needs to be shared between manufacturers which can be beneficial in practice.

The scheduling of acyclic supply chains is similar to hybrid flow shop (HFS) scheduling with precedence constraints, which consists of flow shops with multiple parallel machines in multiple stages. Ruiz and Vázquez-Rodríguez (2010) and Ribas, Leisten, and Framiñan (2010) present literature reviews of different approaches to HFS scheduling problems with tardiness and completion time objectives. The employed techniques are mostly heuristic approximations, meta-heuristics, and mathematical programming formulations. Some studies such as Lee and Yu (2008), Diedrich et al. (2009), and Yin et al. (2017) provide approximation algorithms for scheduling with machine unavailability/disruptions. It should be noted that Potts et al. (1995) proves that minimizing the makespan for two-stage (or two-echelon, in our context) assembly scheduling problems without disruptions is NP-hard. Potts et al. (1995)'s problem has that each 'final' job can have different processing requirements at each of the manufacturers in the first echelon, which is different than in our context since each final job is the same product (or very similar products) thus implying that the processing requirements at each manufacturer are identical. Therefore, our results can be applied to a special case of the two-stage problem considered by Potts et al. (1995) but also can be applied to a more general n-echelon problem. In particular, we are able to develop a polynomial time algorithm for the problem for n echelons given that the number of paths in the network is fixed. There are polynomial time exact algorithms to a few variants of the HFS problems such as those found in Dessouky and Verma (1998), Djellab and Djellab (2002), Guirchoun, Martineau, and Billaut (2005) and Yang (2013), but they are not applicable to the scheduling of acyclic MEASC networks due to limitations on number of stages, machines and their inability to extend to manufacturers that produce different components.

This paper provides strategies to optimize the performance of the *global* recovery problem faced by the MEASC network. The manufacturers within the MEASC network are independent decision-makers but their recovery decisions can be influenced by the final assembler through contracts and incentives. Hence, the objective of this paper is to develop decision rules that the manufacturers can use to solve their local scheduling problems and the optimal global recovery for the key performance metrics of minimizing the maximum tardiness of customer demand and minimizing the time to recover of the MEASC network.

3 Problem Statement

In our MEASC recovery problems, we are given a network of manufacturers, $G = (N, A)$ where the manufacturers are represented by the nodes and an arc $(i, j) \in A$ implies that the component produced by manufacturer i is necessary in manufacturer j's assembly process. We call a node $i \in N$ with zero in-degree a raw materials manufacturer since no components are necessary to assemble their component. We assume that the network is acyclic, which captures many realistic MEASC networks including those of our industrial partner. There is exactly one *final manufacturer* within the network, node B, where the final product is assembled (the out-degree of B is zero). This can be assumed without loss of generality since we can always add in a 'super-final manufacturer' node if there is more than one final manufacturer. There is a set of demand, K, for the final product (which we will refer as *final demand*) of the MEASC network, where each unit of demand $k \in K$ has a deadline D^k. For ease of presentation, we assume that these demands are arranged in non-decreasing order of their deadlines, i.e., if $k_1 < k_2$ then $D^{k_1} \leq D^{k_2}$. For all nodes $i \in N$, we say that they belong to *Tier* E if the path with the largest number of nodes from i to B contains exactly E arcs (B belongs to Tier 0). Without loss of generality, there are M tiers in this network, where Tier M is the tier of raw material manufacturers. It should be noted that although a raw material manufacturer is not always classified as Tier M in reality, we can represent it as Tier M in the network by adding a path of $M - 1$ dummy nodes to B. We define $\mathcal{E}$ to be the set of tiers in the network and $N(E)$ to be the set of manufacturers in tier $E \in \mathcal{E}$.

Each manufacturer $i \in N$ has a production lead time p_i. The capacity of a raw material manufacturer i, cap^i, is the maximum number of products that can be produced in a time period. We define $R(i)$ to be the set of paths r from a manufacturer i to B, which will be used in constructing local scheduling problems. Each arc (i, j) has a shipping lead time, $p_{(i,j)}$. We let $Q_{(i,j)}$ represent the number of components of type i necessary to produce *one* component of type j and define $pred(j)$ to be the set of manufacturers which j requires components (i.e., $i \in pred(j)$ implies that $(i, j) \in A$). Assembly manufacturers (not raw material manufacturers) require a sufficient quantity of all necessary parts before production begins. We also define $Q_r^{i,k}$ as the total components i needs to send along path $r \in R(i)$ in a final product k. For example, if a final product requires 2 parts X and 1 part Y (i.e. $Q_{(X,B)} = 2, Q_{(Y,B)} = 1$), X requires 4 components Z ($Q_{(Z,X)} = 4$), and Y requires 1 components Z ($Q_{(Z,Y)} = 1$), then $Q^{i,k}_{Z-X-B} = Q_{(Z,X)} \times Q_{(X,B)} = 8$ and $Q^{i,k}_{Z-Y-B} = Q_{(Z,Y)} \times Q_{(Y,B)} = 1$. Note that disruptions to manufacturer i can be modeled by reducing the assembly capacity of i to zero (or some value smaller than the capacity) during the disruption period from time 0 to u^i, where u^i is the time that manufacturer i is restored. Alternatively, disruptions can be modeled by adding an uncapacitated (or sufficiently large capacity) dummy node $\tilde{i}$ with a shipping leadtime, $p_{(di,i)} = u^i$. The disruption at node i in the original network is equivalent to i having to wait for the raw material to arrive from manufacturer $\tilde{i}$, and since $\tilde{i}$ is uncapacitated, production at i is no longer restricted by $\tilde{i}$ after time u^i.

We assume that assembly manufacturers have uncapacitated production once they are restored, i.e., their only limitation is the availability of said components while raw material manufacturers' production is still limited by the source of the raw materials, e.g. mining. However, our decision rules can also incorporate a more flexible form of capacity for assembly manufacturers, which we refer to as *cumulative capacity*. If an assembly manufacturer i is cumulatively capacitated, at any time t, its cumulative production from 1 to t has to be no more than it cumulative capacity from 1 to t, $cap^{i,1\to t}$. An example in practice would be an assembly manufacturer that can 'shift' around the capacity it has allocated to the final assembler to different time periods but never violate the overall capacity they have allocated. Cumulative capacity at manufacturer i can be modeled by adding a raw material dummy node $\tilde{i}$ with a shipping leadtime, $p_{(\tilde{i},i)} = 0$, and a production leadtime, $p_{\tilde{i}} = 0$. The capacity of the raw material dummy node $\tilde{i}$ at time t, $cap^{\tilde{i},t}$ is equal to the changes in cumulative capacity at i, $cap^{\tilde{i},t} = cap^{i,1\to(t+1)} - cap^{i,1\to t}$. In other words, the cumulative production at manufacturer i from time 1 to t is capacitated as $cap^{i,1\to t} = \sum_{\tilde{t}=1}^{t} cap^{\tilde{i},\tilde{t}}$.

Additionally, we assume that each manufacturer only produces/assembles one type of component and is the single sole source for that component, which is realistic for MEASC networks with high manufacturer qualifications, like that of our industrial partner.

3.1 Decision Variables and Constraints

We are interested in deciding the timing of the assembly and shipping decisions within this network, $C_{r,q}^{i,k}$. $C_{r,q}^{i,k}$ is the decision variable deciding the time to start shipping a product q ($q = 1, 2 \ldots Q_r^{i,k}$) from manufacturer i to meet final product demand k through path r. Since it is best to start assembling as soon as all the required subcomponents are available, we are able to combine both assembly and shipping decisions into a single variable $C_{r,q}^{i,k}$. It should be noted that this is a pseudo-polynomial representation of the decisions, since we have a decision variable for each q and a deadline for each individual demand k. However, we believe that the proofs of our results are more intuitive with this representation. In terms of implementation, instead of storing completion time $C_{r,q}^{i,k}$ and D^k, we can calculate and store the quantity of production/shipping and demand quantity per time period as shown in the supplemental material. Therefore, there are equivalent polynomial-time representations of our results. There are the following constraints on $C_{r,q}^{i,k}$:

1. every manufacturer is constrained by the assembly capacity schedule, which depends on production capacity and the disruption/availability of that particular manufacturer,

2. we can only begin processing a component at manufacturer j once we have the set of required subcomponents for j's component at the manufacturer, i.e., we must have $Q_{(i,j)}$ of component i for all $(i, j) \in A$,

3. we only ship components out of i once their assembly is complete.

The assembly and shipping decisions then provide completion times of the final product which can be assigned to the demand $k \in K$, generating the completion time out of the final assembly node B, $C^{B,k}$, of each unit of final demand. This paper focuses on two objectives:

- minimizing the maximum tardiness of any final demand k, i.e., let $t^{B,k} = \max\{0, C^{B,k} - D^k\}$ be the tardiness of the delivery individual final product k, then our objective is $\min T_{out}^B = \min \max_{k \in K} t^{B,k}$

- minimizing the time to recover of the MEASC network, i.e., $\min TTR^B = \min \max_{k:t^{B,k}>0} C^{B,k}$, which is the time when the last tardy demand is met, i.e., all the demands after that are on time.

Note that, for a non - raw material manufacturer, we have

$$C^{i,k}_{r,q} \geq \max\left(u^i, \max_{s\in pred(i)}\left(C^{s,k}_{r,q} + p_{s,i}\right)\right) + p_i \tag{1}$$

because the manufacturer must be restored from the disruption and sufficient quantities of all necessary parts are required before production begins.

Additional tardiness and time to recover variables and constraints will be discussed when the decision rules are presented and analyzed.

3.2 Summary of Assumptions

In this subsection, we provide the summary of all the assumptions in our problem:

- Each manufacturer only produces/assembles one product and is the single source for that product. This is appropriate for MEASC networks with strict manufacturer qualification processes.

- Raw material manufacturers do not require parts, and have a fixed production capacity schedule (which may include periods with zero capacity due to the disruption).

- Unless the manufacturer is one of the raw material manufacturers, production is uncapacitated and is only limited by the quantity of available parts, duration of assembly, and disruption status. Recall that this assumption can capture, without loss of generality, assembly manufacturers with cumulative capacities.

- Without loss of generality, the MEASC network can be transformed to an equivalent one where every path from a manufacturer in a given tier goes to a manufacturer in the next tier up. This can be accomplished by adding dummy nodes to the network so that the length of any path from a manufacturer to the final assembly node is of the same length. In the case of multiple final assembly nodes, a super sink can be added.

- The impact of transportation disruptions is modeled by adding a dummy node and treating that node as a manufacturer. For example, if an arc from A to B with shipping leadtime of 7 is disrupted for

5 periods, the we can replace that arc with a dummy manufacturer D with processing time of 7 and disrupt D for 5 periods.

3.3 Constructing the Local Scheduling Problem

Our approach to solve the *global* recovery problem of the MEASC network will be to solve *local* scheduling problems at each manufacturer, based on the recovery metric. Therefore, it is necessary to add parameters associated with these local scheduling problems. For a path $r \in R(i)$ from manufacturer i to the final manufacturer B, we define p_r^i as the minimum lead time required for a product from manufacturer i to be processed and shipped to the final manufacturer through the path, assuming their is no delay as it travels along the path. In particular, we have that $p_r^i = p_{(i,j_1)} + p_{j_1} + p_{(j_1,j_2)} + \cdots + p_{j_n,B} + p_B$ where $r = i - j_1 - j_2 - \ldots - j_n - B$. In other words, we must send the required number of components, $Q_r^{i,k}$, along path r by time $D^k - p_r^i$ in order for final product k to have a chance not to be late. Therefore, the local due date at manufacturer i for demand k is defined to be $D_r^{i,k} = D^k - p_r^i$. Note that we have dropped the q subscript as all components along path r for k must the same due date. Based on the local due dates, we define the following tardiness-related variables:

- $t_{r,q}^{i,k} = \max(C_{r,q}^{i,k} - D_r^{i,k}, 0)$: the tardiness of the product q from manufacturer i for demand k through path r based on the decision $C_{r,q}^{i,k}$;
- $t^{i,k} = \max_{r,q}(C_{r,q}^{i,k} - D_r^{i,k}, 0)$: the tardiness of the manufacturer i for demand k based on decisions at i;
- $t^k = \max_{i,r,q}(C_{r,q}^{i,k} - D_r^{i,k}, 0)$: the tardiness of demand k at the final assembly node.

It is important to note that the final tardiness of task k at B will be greater than or equal to the tardiness of any of the path tasks (i.e. the tasks associated with satisfying the demand k along the path from a manufacturer i to B) associated with k at manufacturer i. Therefore, we have a 'global task' k whose tardiness at manufacturer i is measured as the maximum tardiness of any of the path tasks associated with k. Note that, although the number of paths from manufacturer i in tier M can be $O(n^M)$, it is, in practice, much lower. For our industrial data set, the maximum number of paths from a manufacturer to B is 5.

3.4 Overview of Decision Rules

Table 1 summarizes the definitions of the decision rules that are discussed in the next two sections.

Table 1: Definitions of decision rules

Decision rule	Definition
Earliest due date (EDD) first	Products are shipped in the non-decreasing order of local due date, i.e. if $D_{r1}^{i,k_1} < D_{r2}^{i,k_2}$, then $C_{r1,q1}^{i,k_1} \leq C_{r2,q2}^{i,k_2}$.
Reverse EDD	Works backward from the task with the largest deadline, $D_r^{i,k}$, and assigns the latest available production that is less than or equal to $D_r^{i,k}$.
Longest leadtime (LLT) first	Products are shipped in non-increasing order of p_r^i.

4 Minimizing the Maximum Tardiness of the MEASC: Applying the Earliest Due Date (EDD) Rule

The focus of this section is on examining the recovery metric of minimizing the maximum tardiness of any demand for the final product. In particular, the approach to solve this problem will be to have manufacturers apply the earliest due date (EDD) first rule in meeting their required tasks and, further, to begin production as soon as all components are available. We formally define the global **EDD Rule** as:

- Starting from Tier M to 1, each manufacturer i ships its products to the earliest delivery deadline, i.e. if $D_{r1}^{i,k_1} < D_{r2}^{i,k_2}$, then $C_{r1,q1}^{i,k_1} \leq C_{r2,q2}^{i,k_2}$.

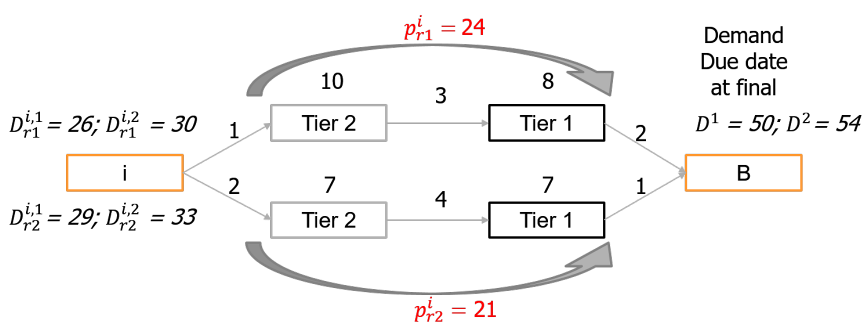

Figure 2: Decision rules example with the production at i ($u^i = 29, cap^i = 1, p_i = 2$) is completed at time $t = 29, 31, 33, 35$

As an illustration, in Figure 2, the EDD shipping decisions($C_{r_1}^{i,1} = 29, C_{r_2}^{i,1} = 31, C_{r_1}^{i,2} = 33, C_{r_2}^{i,2} = 35$ based on the order of local deadlines $26, 29, 30, 33$) gives the optimal maximum tardiness of 3. The application of the local EDD rule at each manufacturer leads the optimal global recovery when the MEASC is focused on minimizing the maximum tardiness of demand.

Theorem 4.1 *Applying the EDD rule for each individual manufacturer minimizes the maximum tardiness of the whole MEASC network, i.e.,* T_{out}^{B}.

Sketch of Proof The detailed proof is presented in the supplemental material. We first show that, using the interchange argument, the EDD rule minimizes the maximum tardiness at the manufacturer's local scheduling level. After that, the key idea behind this proof is to show that the maximum tardiness in each manufacturer's local scheduling problem is a lower bound on the maximum tardiness at B. We can determine these lower bounds at the raw materials manufacturers by applying the EDD rule and show that we only increase the maximum tardiness at the non-raw materials manufacturers if they were disrupted; the length of disruption at those manufacturers are also each a lower bound on the overall maximum tardiness. Therefore, as we move through the tiers of the MEASC network, the current tardiness is equal to a lower bound on the overall maximum tardiness.

5 Minimizing the Time to Recover of the MEASC: Applying the DR-TTR Rule

The time to recover metric focuses on the last time in which there is late demand within the MEASC network. It is noted that while this definition implies that the normal operations of the supply chain result in its ability to meet demand, there may be tardiness built into the supply chain without being caused by the disruption. If we knew that there typically are jobs with ΔT tardiness in its normal operations, we can add ΔT to each of the final demand deadlines D^k. Therefore, the TTR objective can be understood as the last time where an order goes beyond this 'normal' level of tardiness. For the ease of interpretation, the tardiness discussed in this paper is disruption-related.

In terms of the scheduling problem, we would like to determine the best makespan, i.e. completion time, for the demands that are late while ensuring the remaining demands are on time. In fact, if there is subset of demands such that the minimum makespan of the corresponding tasks is strictly larger than the highest final demand deadline of the subset, then one of the tasks in the subset *must be late*. Our goal is to determine a global cutoff $\bar{k}$ where at least one final demand $k \leq \bar{k}$ (recall that $k \in K$ are in non-decreasing order of deadlines) must be late. This leads us to the following definition.

Definition 5.1 *Let $\bar{K} \subseteq K$ be a subset of final demands. If the minimal makespan for the subset, $MS^{\bar{K}*}$, is larger than the largest deadline for that subset, $MS^{\bar{K}*} > \max_{k \in \bar{K}} D^k$, we say that the subset $\bar{K}$* ***must be late****.*

In order to determine the global cutoff $\bar{k}$, we determine local cutoffs $k(i)$ for each raw materials manufacturer $i \in N(M)$. The idea is to identify the local time to recover and then determine how this propagates to the global recovery objective. By definition, there must be a tardy task completed at the manufacturer at this local time to recover and we are then interested in minimizing the amount of time that it takes for this tardy task to reach B. In order to calculate the completion time at local manufacturer i, we introduce the function, $\alpha_r^i(\phi)$. It is the earliest arrival time at the final assembler through path $r \in R(i)$, considering disruptions on the path $r = i - j_1 - j_2..j_n - B$, should the shipment leave i at time ϕ:

$$\alpha_r^i(\phi) = \max\left(\phi + p_r^i, u^{j_1} + p_r^{j_2}, \ldots, u^B + p_B\right) = \max\left(\phi + p_r^i, \max_{m=r \setminus i}\left(u^m + p_r^m\right)\right) \tag{2}$$

The left term of the maximum accounts, $\phi + p_r^i$, for the completion time of the tasks at i and their *best possible* travel times from i to B (i.e., they do not encounter a disrupted manufacturer or have to wait for other components at the nodes). Note that even tasks that were met on-time at manufacturer i may be tardy at B since they could encounter disruptions upstream either due to other components being unavailable (which will be accounted for in analyzing other manufacturers) or disruptions to manufacturers in the upper tiers. Thus, the right term accounts for the potential disruptions *at manufacturers* they may encounter in their path from i to B. This is distinct from our approach for the maximum tardiness objective since we account for upstream disruptions in analyzing manufacturer i's bound on TTR^B. For a manufacturer i, the makespan is defined based on α as followed:

$$MS^{i,K} = \max_{r,q,k \in K} \alpha_r^i(C_{r,q}^{i,k}) \tag{3}$$

We can then define TTR^i as the 'global time to recover' based on manufacturer i, i.e., we have that

$$TTR^i = \max_{r,k,q:t^k>0} \alpha^i_r(C^{i,k}_{r,q}) \quad (4)$$

To determine the best $C^{i,k}_{r,q}$ (as we will show in Theorem 2), we want to select the path task that has the *shortest time* path from i to B and complete it as this last tardy task. Thus, over the set of tasks that need to be tardy, we apply the *longest lead time* (**LLT**) rule (i.e., ship products in non-increasing order of p^i_r) in terms of processing them out of the manufacturer. Using the example in Figure 2, the LLT shipping decisions give the optimal time to recover of 56 ($C^{i,1}_{r_1} = 29, C^{i,2}_{r_1} = 31$ for the top path with leadtime of 24, $C^{i,1}_{r_2} = 33, C^{i,2}_{r_2} = 35$ for the bottom path leadtime of 21). The tasks after the time to recover will be scheduled according to the reverse EDD rule which ensures that they are completed on-time at the manufacturer since we know there is enough capacity to complete the remaining tasks on-time. At this point, we have determined $C^{i,k}_{r,q}$ at the manufacturer i for all final tasks k, products q, and paths r from i to B and can determine a lower bound on the global time to recover, TTR^B, based on how the products flow along the paths to B. In particular, we have

$$TTR^B \geq \max_{i,r\in R(i),k,q:t^{i,k}_{r,q}>0} \left(C^{i,k}_{r,q} + p^i_r, \max_{m\in r, r'\in R(m)} (u^m + p^m_{r'}) \right). \quad (5)$$

or more simply,

$$TTR^B \geq \max_i TTR^i = \max_{i,r\in R(i),k,q,t^k>0} MS^{i,K}. \quad (6)$$

Note that each TTR^i is a lower bound on TTR^B. We can focus just on the raw material manufacturers (i.e., $i \in N(M)$) since we account for upstream disruptions in the calculation of TTR^i. We define $k(i)$ as the local cutoff for must-be-late demands. $k(i)$ is generated by iterating through each final task k from 1 to $|K|$ for each raw material manufacturer i and attempting to see if it is possible to decompose the problem into the set of must be late tasks and on-time tasks based on that final task at B. The global cutoff $\bar{k}$ is then set as $\bar{k} = \max_{i\in N(M)} k(i)$. We then set $\bar{K} = \{1, 2, ..\bar{k}\}$ as the set of tasks that must be late at B. Each manufacturer then receives the set of their tasks associated with $\bar{K}$ and then schedules them according to the longest lead-time (LLT) rule. The tasks that are not associated with $\bar{K}$, i.e., the on-time demands, are scheduled using the reverse EDD rule. Note that the reverse EDD rule will schedule the set of on-time demands in a backwards fashion and in such a manner that we use the *latest* possible availability for each task that still results in it being completed on-time. We determine $\bar{k}$ by solving the local time to recover problems at each manufacturer $i \in N(m)$ which then provides the set of tasks $k \in \bar{K}$ at B that must be late. We then provide each manufacturer $i \in N(m)$ with the set of their tasks that they should consider 'must be late.' We now have a more detailed overview of the decision rule to minimize the time to recover and a sketch of the proof that it provides the optimal global time to recover for the network.

5.1 Decision Rule for Minimizing the Time to Recover (DR-TTR)

We first focus on determining the global cutoff $\bar{k}$ by solving the local manufacturers' problems of minimizing their time to recover. We determine TTR^i for each $i \in N(m)$ by examining the time to recover of the system should it only rely on i and the paths from i to B. We do so by iterating through each final task k from 1 to $|K|$ and find the decomposition that results in the biggest subset of on-time demands at i. For each iteration k, we want to check if it is possible to schedule the reverse EDD rule for the subset of demands $K2 = \{k, k{+}1, \dots |K|\}$. If so, we calculate the makespan for the subset of demands $K1 = \{1, 2, \dots k{-}1\}$, $MS^{i,K1*}$. One additional condition is that this makespan should not exceed any deadlines of the on-time subset of demands, i.e. $MS^{i,K1*} \leq D^k$, or else it would require us to include job k in the must-be-late jobs. Thus, the decomposition is possible and we can break the loop. We prove that $K1$ is either a empty set or a must-be-late subset of demands in Theorem 4. We then set $k(i) = k - 1$ as the local cutoff, and $\mathcal{C}(i) = MS^{i,K1*}$ be the corresponding makespan. The global cutoff is then the maximum of all local cutoffs, $\bar{k} = \max_{i \in N(M)} k(i)$ and $\bar{\mathcal{C}} = \max_{i \in N(M)} \mathcal{C}(i)$ is the lower bound of time to recover at B. The following algorithm provides the details:

Algorithm to determine $\bar{k}$:

- Calculate the production completion time at each of the raw material manufacturers i, based on capacity cap^i and time of availability u^i. Let X be the production completion time vector.

$$X = [\underbrace{\underbrace{u^i + p_i, u^i + p_i, \dots, u^i + p_i,}_{cap^s \text{ terms}} \dots, u^i + 2 \times p_i, u^i + 2 \times p_i, \dots, u^i + \left\lceil \frac{|K|}{cap^i} \right\rceil \times p_i}_{|K| \text{ terms}}]$$

- For manufacturer $i \in N(m)$,
  - For each final demand k from 1 to $|K|$,
    - Let $K2 = \{k_2 | k \leq k_2 \leq |K|\}$ and apply the reverse EDD to $K2$: Start with the largest D_r^{i,k_2}, assign the largest production completion time from X that is less than or equal to D_r^{i,k_2}. If the reverse EDD schedule is possible, then continue; else, skip to the next iteration.
    - Let $X2$ be the subset of X used for $K2$ and $X1 = X \setminus X2$
    - Let $K1 = \{k_1 | k_1 \leq k - 1\}$
    - Apply the LLT rule to $K1$ with $X1$ to calculate $MS^{i,\{1,2,\ldots,k-1\}*}$
    - If $MS^{i,\{1,2,\ldots,k-1\}*} \leq D^k$, then the decomposition is possible. Do the following:
      - Set $k(i) = k - 1$
      - $\mathcal{C}(i) = MS^{i,K1*}$
      - BREAK
- $\bar{k} = \max_{i \in N(M)} k(i)$
- $\bar{\mathcal{C}} = \max_{i \in N(M)} \mathcal{C}(i)$

Our decision rule for the time to recover problem is summarized below.

Decision rules for time to recover (DR-TTR):

- Stage 1: Determine $\bar{k}$ with the algorithm described above. This means that the set $\bar{K}$ defined by the final demands before or at $\bar{k}$ must be tardy given the disruptions and the production capacities, no matter the schedule. Decompose the final demands into two sub-problems:
  - Subproblem 1 ($SP1$) with the subset $\bar{K}$ of all the final demands $k \leq \bar{k}$. These final demands must be late regardless of schedule (equivalently $MS^{\bar{K}*} > \max_{k \in \bar{K}} D^k$), due to lack of production capacity and/or disrupted upper tier manufacturers.
  - Subproblem 2 ($SP2$) with all the final demands after $\bar{k}$. There exists a schedule that achieves the maximum tardiness of 0. This means that $TTR^{SP2*} = 0$ because there is no tardy demand, and $TTR^{SP1 \cup SP2*} = TTR^{SP1}$

- Stage 2: Starting from Tier M to 1, at each manufacturer i:
  - Start production as soon as all components are available and the manufacturer is restored.
  - For the demands in $SP2$, we apply the reverse EDD rule to schedule them.
  - For $SP1$ with all the final demands up to $\bar{k}$, we apply the longest leadtime (**LLT**) first rule: Each manufacturer ships its products to the path with the longest leadtime first, i.e. if $p^i_{r_1} > p^i_{r_2}$, then $C^{i,k_1}_{r_1,q_1} \leq C^{i,k_2}_{r_2,q_2}$.

5.2 Sketch of Proof

This section provides an overview of the proof that the decision rule provides the optimal time to recover for the system and the theorems in the supplemental material containing the key steps.

1. For a single raw material manufacturer i, if we assume that the TTR of the system only depends on the scheduling of i (we call this TTR^i), then decomposing the subproblems at $k(i)$ from the algorithm in Subsection 5.1 and applying the LLT rule to the tasks up to $k(i)$ (Supplemental Material (SM)-Lemma 4.2) and the reverse EDD rule for tasks after $k(i)$ (SM-Theorem 4.3) gives the optimal local time to recover $TTR^{i*} = \mathcal{C}(i)$. TTR^{i*} is the optimal time to recover of manufacturer i and $\mathcal{C}(i)$ is the time to recover given by the algorithm (SM-Theorem 4.4). Note that $TTR^{i*} = \mathcal{C}(i)$, is a lower bound on TTR^B.

2. The algorithm determines the subset of final demands that must be late by setting the global cutoff to be $\bar{k} = \max_{i \in N(M)} k(i)$. Let $\bar{\mathcal{C}} = \max_{i \in N(M)} \mathcal{C}(i)$ be the maximum optimal individual TTR across raw material manufacturers. We can show that decomposing the main problem into 2 subproblems based on $\bar{k}$ and applying the decision rules at a raw material manufacturer i gives a local time to recover, $TTR^i \leq \bar{\mathcal{C}}$ (SM-Theorem 4.5).

3. If we assume that the TTR of the system only depends on the scheduling of all raw material manufacturers (i.e. Tier M manufacturers), applying the decision rule at all raw material manufacturers i will give the optimal $TTR^{M*} = \bar{\mathcal{C}}$ (SM-Theorem 4.6).

4. Let TTR^E be the TTR of the system if it only depends on the scheduling of all manufacturers from Tier M to E. We can prove that if the decision rule is applied in all tiers E from Tier M to 1, the time to recover does not increase: $TTR^E = TTR^{E+1}$ (SM-Theorem 4.7).

5. Hence, the decision rules minimize the time to recover for the whole system $TTR^B = TTR^{M*} = \bar{\mathcal{C}}$ (SM-Theorem 4.8).

6. We then show that applying the reverse EDD rule to minimize the maximum tardiness for the demands after $\bar{k}$ results in a maximum tardiness of 0. Thus, $TTR^B = \bar{\mathcal{C}}$ and is optimal.

6 Computational testing

6.1 Case Study Description

We test our algorithms on a MEASC network of our industrial partner in defense aircraft manufacturing. The production of this particular aircraft involves 512 manufacturers across 5 tiers (see Table 2) distributed over 44 states and 10 countries (including the United States).

	Tier 5	Tier 4	Tier 3	Tier 2	Tier 1	Final Assembly
Number of manufacturers	3	26	116	322	44	1
Average out-degree	1	1.31	1.22	1.18	1	0
Average in-degree	0	0.12	0.23	0.39	9.05	51
Max out-degree	1	2	4	4	1	0
Max in-degree	0	3	10	20	46	51
Number of raw material manufacturers	3	25	111	302	0	0

Table 2: Initial characteristics of the data set.

The global distribution of their MEASC networks makes them more vulnerable to a diverse set of disruptive events, especially ones that impact a ‘popular’ area for their manufacturers. The production lead time at each manufacturer is drawn from a uniform distribution $[3, 5]$ and the shipping lead time is drawn from the uniform distribution $[1, 2]$. There is a demand of 5 final products every 7 periods of time. The time horizon H for this problem is 80, i.e., we consider production over 80 time periods after the initial disruption. At a raw material manufacturer, we calculate the output of the manufacturer required to meet the demand during normal operating conditions and set production capacity to be 1.5 times this amount. For assembly manufacturers, the production is uncapacitated and only limited by the arrivals of all necessary sub-components from the lower tier manufacturers and their disruption status. For each pair of $(i, j) \in A$, the number of components of type i necessary to produce one component of type j, $Q_{(i,j)}$, is drawn randomly from uniform distribution $[2, 4]$.

For the purpose of demonstrating the application of our decision rules, we test two case studies based on two large-scale disruptions: (i) all manufacturers in New Jersey are disrupted for an extended period of time (modeling the impact of a Hurricane Sandy-like scenario) and (ii) all manufacturers around Northridge, CA are disrupted for an extended period of time (modeling the impact of a large earthquake). For Case (i), the disruption length, u^i, of the 11 manufacturers are generated depending on its location relative to the impact of the 2012 Hurricane Sandy. For example, disruption times generated from a uniform distribution over $[10, 30]$ for manufacturers in high impact areas, and over $[2, 8]$ for mildly impactedareas. In Case (ii), the disrupted lengths of the 15 manufacturers are generated randomly depending on their individual locations relative to the impact of the 1994 earthquake in the area (e.g. disruption lengths are generated from uniform distribution$[15, 30]$ for the highly affected area, and $[5, 15]$ for less affected areas). For each case study, we have examined 5 randomly generated instances, along with the worst case scenario (all lengths at their maximum) and the best case scenario (all lengths at their minimum). We run our decision rules against the equivalent integer programming (IP) formulations (see the supplemental material) on a quad-core i7-4980HQ CPU with 16 GB of RAM. For minimizing maximum tardiness, the decision rule provides the optimal solution to these problems in under half a second while the IP formulations require, on average, 1.5 minutes. For minimizing the time to recover, the runtime for the decision rule is also under half a second while the IP models take around 4 minutes. Preliminary testing has examined situations where the planning horizon ranges from 2000 time periods (e.g., if each period was a day, this would cover around five and a half years) to 8000 time periods have indicated that the decision rules have a runtime of less than ten seconds while the IP models cannot be solved due to the number of variables. These horizons are reasonable to consider given the life cycle of certain MEASC products, i.e., aircrafts.

We begin by exploring the impact of mitigation strategies on the recovery from the disruptions. The decision rule allows supply chain managers to examine not only the impacts of certain disruptions on the network but how different pre-event mitigation strategies can protect the performance of the network against such disruptions. As a demonstration, we examine the impact on the maximum tardiness and the time to recover over the five instances when each manufacturer carries a certain level of safety stock. In particular, we calculate the output of the manufacturer required to meet the demand during normal operating conditions for a demand cycle and then carry safety stock to be able to meet a certain percentage of this output (0%, 25%, 50%, and 100%) as shown in Figure 3 and Figure 4. It should be noted that the proposed scheduling rules is also optimal under any inventory policies. We adjust the proof for the optimality by noting that if a manufacturer i is holding quantity $\hat{q}$ of subcomponent j, then these q quantity of component j is removed from the local demand of manufacturer j. In other words, manufacturer j does not need to produce this q quantity so we can set $C_{r,q}^{j,k} = 0$ for the first $\hat{q}$ components.

An exploration of the return on investment of paying its manufacturers to carry this safety stock can now be conducted by supply chain managers. For example, for the Northridge event, Figure 3 shows that there is not a significant decrease in the maximum tardiness until the manufacturers carry 100% safety stock. The supply chain manager could now examine improvements in the recovery of the MEASC network over a variety of disruptive scenarios in order to determine if this improvement justifies the increases in the cost of operating the MEASC network. Additionally, these decision rules allow them to analyze the tradeoffs between the two resiliency objectives (maximum tardiness and time to recover) in their recovery process. This will allow managers to be more informed in making their mitigation strategies.

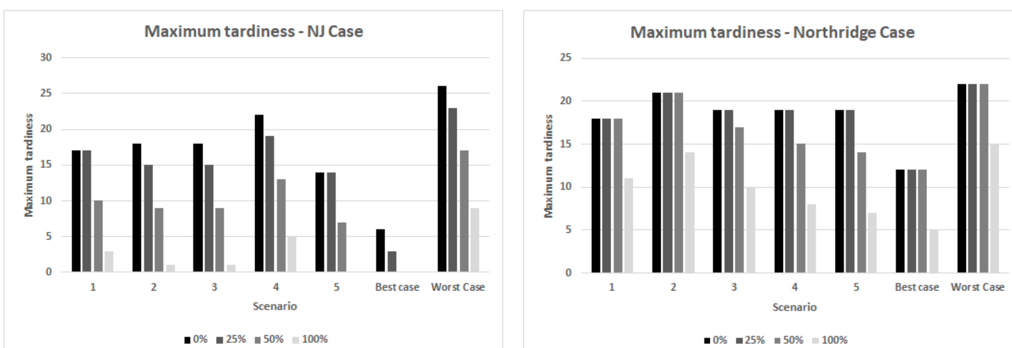

(a) (b)

Figure 3: The impact on the maximum tardiness objective through the safety stock mitigation strategies.

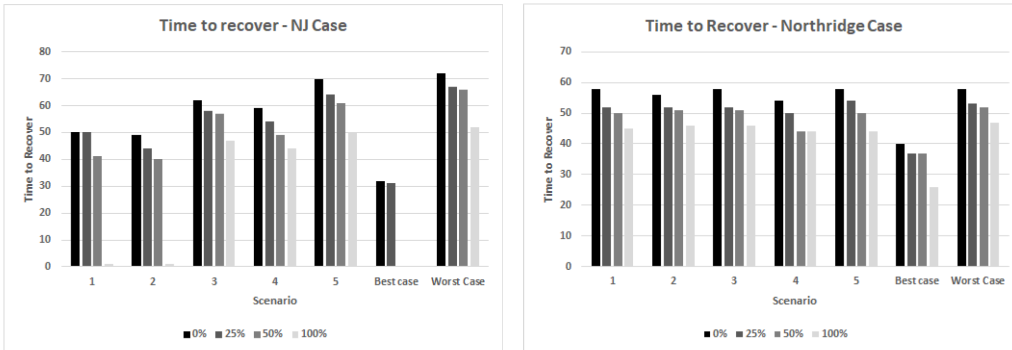

(a) (b)

Figure 4: The impact on the time to recover objective through the safety stock mitigation strategies.

6.2 Effects of stochastic disruption lengths

In this section, we consider the effects of stochastic disruption lengths, u_i, at individual manufacturers i on the optimality of the two decision rules. For the objective of minimizing maximum tardiness, u_i changing will not affect the ordering of EDD first decision rules as it is calculated solely based on the due dates of the final demands and the leadtimes. In the analysis of the EDD decision rules, we can see that changing u_i would only change the numeric value of the optimal maximum tardiness for the problem. but the optimality of decision rules remains, i.e. T^E_{dis} is affected by u_i but T^E_{out} remains unchanged. Therefore, even if there is uncertainty in the disruption lengths, it is globally optimal for the local manufacturers to schedule according to the EDD rule.

For the objective of minimizing TTR, we created a case study to see how varying noise level at u_i would affect TTR. We generated 100 scenarios from Case (i) hurricane disruptions. For each scenario, individual manufacturers have estimates of their own and others' disruption lengths. These estimates are subject to some levels of random noise ϵ , $u_i \pm \epsilon_i$. The noise ϵ_i is generated from the uniform distribution from 0 to the maximum noise level of $40\%, 80\%$, and 160% of the true disruption lengths. manufacturers calculate the cut-off based on the noisy disruption lengths and make their scheduling decisions. We then evaluate to see how they perform against the optimal plan with perfect knowledge of the disruptions by looking at the relative difference between the TTR of the plan with imperfect information and the optimal TTR, i.e. the optimality gap. Table 3 shows the summary statistics of the relative optimality gaps against varying level of noise. For noise level up to 40%, the optimality gap on average is about 2.9% and goes up to

13.1% for maximum level of noise 160%. These values are acceptable considering the high levels of errors in estimating the disruption lengths. For the maximum noise level of 40% and 80%, more than half the scenarios achieves optimal results using imperfect information.

Table 3: Statistics on the effects of stochastic disruption lengths on the optimality gap

Maximum noise level	Mean	Standard deviation	Minimum	25% percentile	Median	75% percentile	Maximum	No. of optimal results
40%	2.9%	4.4%	0.0%	0.0%	0.0%	7.3%	13.5%	69
80%	5.1%	6.7%	0.0%	0.0%	0.0%	10.6%	26.8%	52
160%	13.1%	15.4%	0.0%	0.0%	8.1%	20.0%	53.8%	35

Additionally, we consider a case study where the disruption lengths are stochastic but the manufacturers get better estimates as time goes on. At the start of the time horizon ($t = 0$), manufacturers will have noisy estimates of disruption time $u_{i,0} = u_i \pm \epsilon_i$, then make initial plans. After every few periods Δt, manufacturers will recalculate the cutoff k and make new plans with better estimates, $u_{i,t}$, where $u_{i,t}$ is generated randomly from uniform random distribution between $u_{i,t-\Delta t}$ and u_i. If the time t goes past the true disruption time for a manufacturer i, i.e. $t \geq u_i$, then $u_{i,t} = u_i$. All the production and shipping decisions that were carried out cannot be change, i.e. the shipping and production occurs from 0 to t are fixed. These updates continue for the remaining horizon. We want to evaluate how much improvement will be from the disruption lengths estimate updates relative to the previous case study. We found that setting short enough interval such as $\Delta t = 7$, which is equivalent of one demand cycle, actually gives the optimal results for all of the scenarios as if we have full knowledge of the disruption lengths. Our intuition is that if there is large enough gap between u_i and optimal TTR that once the true disruption length u_i is revealed, the algorithm still have enough flexibility to correct itself. In other words, once all the true disruption lengths are revealed, there is enough time before the actual TTR to 'catch up' with the unstarted'longest leadtime' shipments that are needed to achieve the optimal TTR.

6.3 Tradeoff between the two algorithms

To select between the two algorithms, the supply chain managers can analyze the trade-offs when using them to optimize different objectives. As a case study, we rerun 100 scenarios from the Case (ii) earthquake disruptions to compare EDD and DR-TTR with three relevant objectives: maximum tardiness, TTR and to-

Table 4: Comparison between EDD and DR-TTR algorithms against three resilience objectives (in weeks)

Decision rules	Objectives	Average	Standard Deviation	Optimality Gap
EDD	TTR	53.76	2.07	0.3%
	Total Tardiness	92.2	9.31	25.0%
	Max. Tardiness	19.05	1.12	0.0%
DR-TTR	TTR	53.6	2.03	0.0%
	Total Tardiness	87.63	9.38	18.8%
	Max. Tardiness	26.72	2.42	40.3%
	Optimal Total Tardiness	73.77	6.98	0.0%

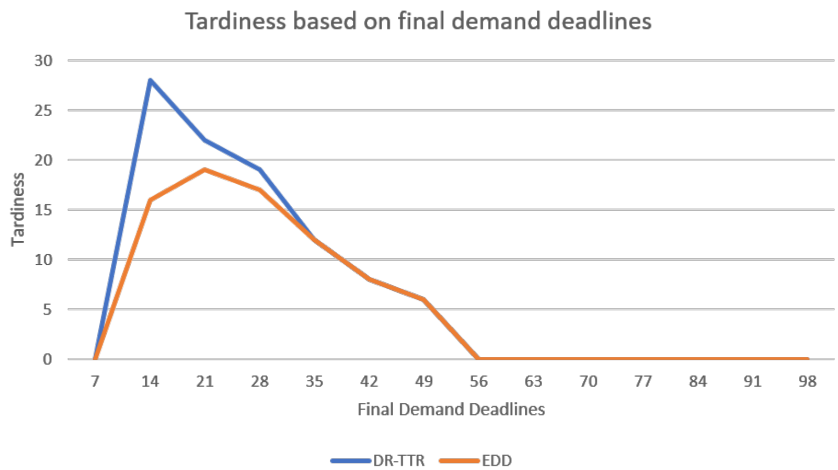

Figure 5: Order tardiness over demand deadlines comparison of between decision rules in one scenario

tal tardiness. As Table 4 shows, EDD outperforms DR-TTR quite significantly in minimizing the maximum tardiness while DR-TTR achieves slightly better results in minimizing total tardiness and TTR. The two algorithms shows similar patterns in their objectives to each other and to the optimal total tardiness objectives. In this particular disruption, the supply chain managers are likely to pick EDD as their scheduling policy. Similar trade-off analysis can be done for different type disruptions scenarios as well as performance metrics to give more managerial insights on selecting the right policy for the right type of disruptions. Furthermore, the manager can take a closer look at each particular scenario, to see how the decision rules differs in term of order tardiness. As an example, we take a look at a scenario from the 100 as shown in Figure 5. This further supports picking EDD as their scheduling policy for this particular disruptions as the tardiness of earlier orders is lower for EDD and equal for later orders.

6.4 Notes on runtime complexity and advantages of the decision rules

Theoretically, given a MEASC networks with $|N|$ of manufacturers, $|A|$ shipping arcs, M tiers and $|K|$ number of demands, implementing the EDD/LLT first decision rule requires applying the sorting algorithm for the local demand deadlines/leadtimes at $O(|N|)$ manufacturers where the number of local demands is equal to $|K|\times$the number of paths from the manufacturer to the final manufacturer. Although the number of paths could be exponential in terms of the number of manufacturers, in practice, it is often quite low(for example, in our test data, the max number of paths is 1 for tier 1, 4/4/5/2 for tier 2/3/4/5 manufacturers respectively, and the total number of unique paths in the network is 649). Theoretically, the number of local demand deadlines/leadtimes to sort depends on the number of paths from the manufacturer to the final assembler B. The maximum number of paths in a directed acyclic graph is $(N/M)^M$, by putting (N/M) nodes in each tier. For bounded M, we get a polynomial number of paths. In practice, M is small, e.g. Ford has up to 10 tiers Simchi-Levi et al. (2015). The EDD/LLT first decision rule is polynomial for a fixed number of (maximum) paths from the manufacturer to the final demand. Each node can calculate the local demand deadlines and leadtimes to B from the those values of its successors. Therefore, the time complexity of calculating local demand deadlines depends on number of demands $|K|$ and the out-degrees of manufacturers, i.e. number of arcs coming out of the manufacturers, $O(|A||K|)$. For reference, our network has 512 manufacturers, 705 arcs, with the maximum out degree of 4. As noted before, while the proofs use a pseudo-polynomial representation of the decision-making environment, we can implement our decision rules in polynomial time (based on the number of local demand deadlines) and formulate integer programs with a number of decision variables that is polynomial in terms of input into the entire scheduling problem. More specifically, the polynomial time representation for the the inputs is presented in the supplemental material. Thus, the EDD decision rule is polynomial for a fixed number of (maximum) paths from the manufacturer to the final demand.

The reverse EDD has the same complexity as the EDD rule. Implementing the LLT first decision rule at each manufacturers require applying the sorting algorithm for the leadtimes at $O(|N|)$ manufacturers where the total number of local demands is $|K|\times$the number of paths from the manufacturer to the final manufacturer. For DR-TTR, the complexity lies on the algorithm to determine $\bar{k}$ which involves applying the reverse EDD and LLT at $O(|N|)$ raw material manufacturers for $O(|K|)$ iterations each. Therefore, implementing the DR-TTR rule is also polynomial time. As the result, we see better computational time for both decision rules when compared to solving the integer programming models which, in general, is NP-hard.

Empirically, it is clear from the case study in Section 6.1 that the computational implementation of the decision rules are significantly quicker than the IP formulations. In addition to being computationally faster than the IP formulation, the decision rules are also more robust. Problems with longer time horizons are more demanding in terms of memory usage in solving the IP formulations than they are to the decision rules. The reason is that IP formulations require additional decision variables for each additional unit increase in the time horizon. Hence, MEASC networks with high variations in lead times as well as long planning horizons are likely too difficult to be solved with IP formulations. To demonstrate this limitation, we use the same aircraft production MEASC network data. The only difference is that we now use the estimates of the actual shipping lead times and production lead times provided in the data (in days as opposed to weeks in our previous analysis). The shipping lead times range from 2 to 68 days because of the discrepancies in ground and sea shipping. The production lead times range from 14 days to 578 days. It should be noted that since not all production lead time estimates are available in the data, we randomly sampled the lead time from the set of estimated lead times to assign the unknown production lead times. There is demand of 5 final product every day. Our disruption scenarios are based on the path of the 2016 Hurricane Matthew with 15 manufacturers are disrupted. The disruption lengths are randomly generated from 30 to 365 days.

Table 5 shows the comparison between the decision rules and the IP formulation with different length of planning horizon starting from 2000 days (or around five and a half years) to 8000 days. These planning horizons are not too extreme considering the life cycle in aircraft production. The IP formulations either take longer to run or simply not solvable due to memory limitation. On the other hand, the decision rules average under 10 seconds of runtime. This means that not only are the decision rules proven to be optimal, their short runtimes allow supply chain managers to run their network design against many possible disruption scenarios in a short amount of time. Thus, supply chain managers will be able to use the decision rules to evaluate and compare their pre-event mitigation strategies with a much larger set of scenarios. The ability to solve that problem will give broader insights into the recovery of these networks.

As mentioned earlier, one advantage of the approach is that decision rules are applied locally at a man-

Table 5: Runtime/solvability comparison between decision rules and integer programming

Objective	Time Horizon H(days)	Decision rules	Integer Programming Models
Maximum Tardiness	2000	1.24s	21m32s
	4000	3.28s	out of memory
	8000	8.78s	out of memory
Time to Recover	2000	1.22s	out of memory
	4000	3.26s	out of memory
	8000	8.95s	out of memory

ufacturer level. To minimize maximum tardiness, the centralized decision maker can induce the EDD first rule at each manufacturer by imposing local deadlines for each demand. DR-TTR requires a bit more information at the manufacturer level but does not require manufacturers to know the restoration plans of others in the network. Since final demands k are arranged sequentially in the order of deadlines D^k, the inclusion of either chronologically generated final demand ID numbers or final deadline information when ordering would be necessary. In practice, when ordering subcomponents from subtier manufacturers, every manufacturer needs to attach the final demand ID number or the final deadlines to its orders along with leadtime to final assembler (which can be added up as manufacturers movers down the tiers). When the disruption occurs, the calculation of $\bar{k}$ is done by the final assembler. At the local manufacturer level, the manufacturers only need to know the final demand cutoff $\bar{k}$ that decomposes the problem. The manufacturers then can follow LLT decision rules for with orders with ID up to $\bar{k}$ or final deadlines smaller than D^k and continue with reverse EDD after.

7 Conclusions and Future Work

We developed local decision rules at a manufacturer level to optimize the recovery for MEASC networks at the global level for two objectives: minimizing the maximum tardiness of any demand and minimizing the time to recover of the network. The decision rules allow each individual manufacturer to manage their own recovery process while focusing on the global objective of the entire supply chain network. Additionally, they allow managers to quantitatively test the impact of different pre-event mitigation strategies against disruption scenarios. This can be meaningful as the vast majority of resiliency analysis in industry are qualitative in nature and the quantitative impacts of the disruptive event on the network are not well-known. Therefore, the results of the decision rules can better inform supply chain managers of the exposure of each of their MEASC networks to various potential disruptions, especially when selecting manufacturers for the MEASC networks of future products, and understand the types of activities that should be undertaken (if economically feasible) to decrease this exposure of their current portfolio of MEASC networks. This understanding then helps to inform policies to improve the resiliency of these networks.

In addition to the theoretical analysis, the decision rules are shown empirically to be much faster and robus than solving IP formulations with commercial software. For future research, we would like to extend this work to create a resiliency index for their MEASC networks. This index will help to determine the vulnerability of their MEASC networks to different types of disruptive events as well as the true impact of such events. Another potential research direction is to examine the objective of minimizing total tardiness. While the problem at a manufacturer is NP-hard, it might be possible to create a decision rule in the special case where all requirements are unit (i.e., we need one sub-component from each incoming arc). Acknowledgments here

Acknowledgement

We would like to thank our industry partners for providing the data which our case study network is based on. The work of Huy Nguyen was partially supported by an RPI Presidential Fellowship. The work of Thomas C. Sharkey was supported by the U.S. National Science Foundation under award number CMMI-1254258.

References

Ang, E., D.A. Iancu, and R. Swinney. 2016. "Disruption risk and optimal sourcing in multitier supply networks." *Management Science* 63 (8): 2397–2419.

Atan, Z., and L V Snyder. 2012. "Disruptions in one-warehouse multiple-retailer systems." *Available at SSRN* .

Averbakh, I., and J. Pereira. 2015. "Network construction problems with due dates." *European Journal of Operational Research* 244 (3): 715 – 729.

Baxter, M., T. Elgindy, A.T. Ernst, T.Kalinowski, and M.W.P. Savelsbergh. 2014. "Incremental network design with shortest paths." *European Journal of Operational Research* 238 (3): 675 – 684.

Benjaafar, S., and M. ElHafsi. 2006. "Production and Inventory Control of a Single Product Assemble-to-Order System with Multiple Customer Classes." *Management Science* 52 (12): 1896–1912.

Benjaafar, S., M. ElHafsi, C.-Y. Lee, and W. Zhou. 2011. “Optimal Control of an Assembly System with Multiple Stages and Multiple Demand Classes.” *Operations Research* 59 (2): 522–529.

Bimpikus, K., D. Fearing, and A. Tahbaz-Salehi. 2013. “Multi-Sourcing and Miscoordination in Supply Chain Networks.” *Technical Report* .

Bollapragada, R., U.S. Rao, and J. Zhang. 2004. “Managing Inventory and Supply Performance in Assembly Systems with Random Supply Capacity and Demand.” *Management Science* 50 (12): 1729–1743.

Bruneau, M., S.E. Chang, R.T. Eguchi, G.C. Lee, T.D. O’Rourke, A.M. Reinhorn, M.Shinozuka, K.Tierney, W.A. Wallace, and D. von Winterfeldt. 2003. “A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities.” *Earthquake Spectra* 19 (4): 733–752.

Çelik, M. 2016. “Network restoration and recovery in humanitarian operations: Framework, literature review, and research directions.” *Surveys in Operations Research and Management Science* 21 (2): 47–61.

Celik, M., O. Ergun, and P. Keskinocak. 2015. “The Post-Disaster Debris Clearance Problem Under Incomplete Information.” *Operations Research* 63 (1): 65–85.

Ceryan, O., I. Duenyas, and Y. Koren. 2012. “Optimal control of an assembly system with demand for the end-product and intermediate components.” *IIE Transactions* 44 (5): 386–403.

CNN Library. 2011. “Japan Earthquake - Tsunami Fast Facts - CNN.com.” .

DeCroix, G.A. 2013. “Inventory Management for an Assembly System Subject to Supply Disruptions.” *Management Science* 59 (9): 2079–2092.

Dessouky, M.I., and S.K. Verma. 1998. “Flowshop scheduling with identical jobs and uniform parallel machines.” *European Journal of Operational Research* 109 (3): 620–631.

Diedrich, F., K. Jansen, U.M. Schwarz, and D. Trystram. 2009. “A survey on approximation algorithms for scheduling with machine unavailability.” In *Algorithmics of Large and Complex Networks. Lecture Notes in Computer Science*, Vol. 5515 LNCS, 50–64. Springer, Berlin, Heidelberg.

Djellab, H., and K. Djellab. 2002. “Preemptive hybrid flowshop scheduling problem of interval orders.” *European Journal of Operational Research* 137 (1): 37–49.

Gao, Y., D. Simchi-Levi, C.-P. Teo, and Z. Yan. 2016. “Disruption Risk Mitigation in Supply Chains - the Risk Exposure Index Revisited.” *SSRN Electronic Journal* .

Gerchak, Y., Y. Wang, and C.A. Yano. 1994. “Lot sizing in assembly systems with random component yields.” *IIE Transactions* 26 (2): 19–24.

Guirchoun, S, P Martineau, and J.-C Billaut. 2005. “Total completion time minimization in a computer system with a server and two parallel processors.” *Computers & Operations Research* 32: 599–611.

Gurnani, H., R. Akella, and J. Lehoczky. 2000. “Supply management in assembly systems with random yield and random demand.” *IIE Transactions* 32 (8): 701–714.

He, L. 2014. “Multi-echelon inventory optimization and demand-side management: Models and Algorithms.” PhD diss., Lehigh University.

Ho, W., T. Zheng, H. Yildiz, and S. Talluri. 2015. “Supply chain risk management: A literature review.” *International Journal of Production Research* 53 (16): 5031–5069.

Ji, Q., Y. Wang, and X. Hu. 2016. “Optimal production planning for assembly systems with uncertain capacities and random demand.” *European Journal of Operational Research* 253 (2): 383 – 391.

Jiang, H., D.and Li, T. Yang, and D. Li. 2016. “Genetic algorithm for inventory positioning problem with general acyclic supply chain networks.” *European J. of Industrial Engineering* 10 (3): 367.

Lee, C.Y., and G. Yu. 2008. “Parallel-machine scheduling under potential disruption.” *Optimization Letters* 2 (1): 27–37.

Lei, L., K. Lee, and H. Dong. 2016. “A heuristic for emergency operations scheduling with lead times and tardiness penalties.” *European Journal of Operational Research* 250 (3): 726 – 736.

Lu, Y., J.-S. Song, and D.D. Yao. 2003. “Order Fill Rate, Leadtime Variability, and Advance Demand Information in an Assemble-to-Order System.” *Operations Research* 51 (2): 292–308.

Nurre, S. G., B. Cavdaroglu, J. E. Mitchell, T. C. Sharkey, and W. A. Wallace. 2012. “Restoring infrastructure systems: An integrated network design and scheduling (INDS) problem.” *European Journal of Operational Research* 223(3): 794–806.

Ostrower, J. 2012. “Boeing 787 Profit Plan Strained by Order Cancellation - WSJ.” .

Pinedo, M. L. 2012. *Scheduling: Theory, Algorithms, and Systems*. Fourth ed. New York, NY: Springer.

Potts, C.N., S.V. Sevast’Janov, V.A. Strusevich, L.N. Van Wassenhove, and C.M. Zwaneveld. 1995. “The two-stage assembly scheduling problem: complexity and approximation.” *Operations Research* 43 (2): 346–355.

Ribas, I., R. Leisten, and J.M. Framiñan. 2010. “Review and classification of hybrid flow shop scheduling problems from a production system and a solutions procedure perspective.” *Computers & Operations Research* 37 (8): 1439–1454.

Rong, Y., Z. Atan, and L.V. Snyder. 2014. *Heuristics for base-stock levels in multi-echelon distribution networks with first-come first served policies*. Technical Report. Available at SSRN: http://ssrn.com/abstract=1475469.

Rong, Y., Z. Bulut, and L. V. Snyder. 2009. “Heuristics for base-stock levels in multi-echelon distribution networks with first-come first served policies.” *SSRN Electronic Journal* (July 2015): 1–36.

Rosling, K. 1989. “Optimal Inventory Policies for Assembly Systems Under Random Demands.” *Operations Research* 37 (4): 565–579.

Ruiz, R., and J.A. Vázquez-Rodríguez. 2010. “The hybrid flow shop scheduling problem.” *European Journal of Operational Research* 205 (1): 1–18.

Sarrafha, K., S.H.A. Rahmati, S.T.A. Niaki, and A. Zaretalab. 2015. “A bi-objective integrated procurement, production, and distribution problem of a multi-echelon supply chain network design: A new tuned MOEA.” *Computers & Operations Research* 54: 35–51.

Sawik, T. 2016. “Integrated supply, production and distribution scheduling under disruption risks.” *Omega* 62: 131–144.

Schmidt, C.P., and S. Nahmias. 1985. “Optimal Policy for a Two-Stage Assembly System under Random Demand.” *Operations Research* 33 (5): 1130–1145.

Schmitt, A. J., S. A. Sun, L. V. Snyder, and Z.J.M. Shen. 2015. “Centralization versus decentralization: Risk pooling, risk diversification, and supply chain disruptions.” *Omega (United Kingdom)* 52: 201–212.

Schmitt, A.J., and L.V. Snyder. 2012. “Infinite-horizon models for inventory control under yield uncertainty and disruptions.” *Computers and Operations Research* 39 (4): 850–862.

Simchi-Levi, D., W. Schmidt, and Y. Wei. 2014. “From Superstorms to Factory Tires, Managing Unpredictable Supply Chain Disruptions.” *Harvard Business Review* (January/February).

Simchi-Levi, David, William Schmidt, Yehua Wei, Peter Yun Zhang, Keith Combs, Yao Ge, Oleg Gusikhin, Michael Sanders, and Don Zhang. 2015. “Identifying Risks and Mitigating Disruptions in the Automotive Supply Chain.” *Interfaces* 45 (5): 375–390.

Snyder, L.V., Z. Atan, P. Peng, Y. Rong, A.J. Schmitt, and B. Sinsoysal. 2016. “OR/MS models for supply chain disruptions: A review.” *IIE Transactions* 48 (2): 89–109.

Song, J.-S., and D.D. Yao. 2002. “Performance Analysis and Optimization of Assemble-to-Order Systems with Random Lead Times.” *Operations Research* 50 (5): 889–903.

Song, J.-S., and P. Zipkin. 2003. “Supply Chain Operations: Assemble-to-Order Systems.” In *Supply Chain Management: Design, Coordination and Operation*, Vol. 11 of *Handbooks in Operations Research and Management Science*, 561 – 596. Elsevier.

Tang, C.S. 2006. “Perspectives in supply chain risk management.” *International Journal of Production Economics* 103 (2): 451 – 488.

The White House, Office of the Press Secretary. 2013. "Presidential Policy Directive - Critical Infrastructure Security and Resilience." .

Tomlin, B. 2006. "On the value of mitigation and contingency strategies for managing supply chain disruption risks." *Management Science* 52 (5): 639–657.

Tukamuhabwa, B.R., M. Stevenson, J. Busby, and M. Zorzini. 2015. "Supply chain resilience: Definition, review and theoretical foundations for further study." *International Journal of Production Research* 53 (18): 5592–5623.

Yang, Jaehwan. 2013. "A two-stage hybrid flow shop with dedicated machines at the first stage." *Computers & Operations Research* 40 (12): 2836–2843.

Yao, D.D. 1988. "Optimal Run Quantities for an Assembly System with Random Yields." *IIE Transactions* 20 (4): 399–403.

Yao, J., and L. Liu. 2009. "Optimization analysis of supply chain scheduling in mass customization." *International Journal of Production Economics* 117 (1): 197–211.

Yin, Y., Y. Wang, T.C.E. Cheng, W. Liu, and J. Li. 2017. "Parallel-machine scheduling of deteriorating jobs with potential machine disruptions." *Omega* 69: 17–28.

Supplementary Material

1 Summary of Relevant Notation

Parameters

Notation	Definition
N	set of manufacturers.
A	set of arcs. $(i,j) \in A$ implies that component i is required by manufacturer j.
B	the final assembly node.
$\mathcal{E}$	$\mathcal{E} = \{0, 1, 2, \ldots, M\}$: the set of tiers in the network.
$N(E)$	the set of nodes in Tier E, where $E \in \mathcal{E}$.
K	the set of demands for final assembled product.
H	the total number of time periods in the scheduling horizon.
D^k	the due date of demand $k \in K$ at the final assembly node.
$D_r^{i,k}$	the local due date that components i, which will be used for demand k, need to be sent out along path r.
$pred(i)$	the set of sub-component manufacturers for manufacturer i
$R(i)$	the set of paths r from a manufacturer i to B.
$p_{(i,j)}$	the shipping time from manufacturer i to manufacturer j.
p_i	the production time at manufacturer i.
cap^i	the capacity of raw material manufacturer i.
$cap^{i,1\to t}$	The cumulative capacity of manufacturer i, i.e., cumulative production up to t is bounded by it.
$Q_{(i,j)}$	the number of components of type i necessary in the production of *one* component of type j .
$Q_r^{i,k}$	the total components i needed to be on path r in a final product k.
p_r^i	the minimum lead time it takes a product from i to be processed and shipped along $r \in R(m)$.
u^i	the time manufacturer i is restored.

Variables

Notation	Definition
$C_{r,q}^{i,k}$	the time to start shipping a product q from manufacturer i to meet final product demand k through path r..
$t_{r,q}^{i,k}$	the tardiness of the product q from manufacturer i for demand k through path r based on the decision $C_{r,q}^{i,k}$.
$t^{i,k}$	the tardiness of the manufacturer i for demand k based on shipping decisions at i alone.

t^k	the tardiness of demand k at the final assembly node.
$T_{in}^{i,k}$	the tardiness of manufacturer i in meeting demand k based on the arrival of needed components.
T_{in}^{i}	the tardiness of manufacturer i across all demands based on the arrival of needed components.
T_{in}^{E}	the maximum tardiness of all manufacturers i in Tier $E \in \mathcal{E}$ based on the arrival of needed components.
T_{dis}^{i}	the tardiness of manufacturer i across all demands based on the disruption at i.
T_{dis}^{E}	the maximum tardiness of all manufacturers i in Tier $E \in \mathcal{E}$ based on their disruptions.
$T_{out}^{i,k}$	the maximum tardiness of manufacturer i in meeting demand k based on its shipping decisions.
T_{out}^{i}	the maximum tardiness of manufacturer i across all demands based on its shipping decisions.
T_{out}^{E}	the maximum tardiness of all manufacturers i in Tier $E \in \mathcal{E}$ based on their shipping decisions.
$\alpha_r^i(\phi)$	the earliest arrival time a product from manufacturer i to be processed and shipped at time ϕ to the final product through path $r \in R(i)$, where $r = \ i - j_1 - j_2..j_n - B$, considering disruptions at the manufacturers on the path.
$\tau^{i,k}$	the tardiness of manufacturer i for demand k based on decisions at i and the disruptions on paths from i to B.
$MS^{i,K}$	$MS^{i,K} = \max_{r,q,k \in K} \alpha_r^i(C_{r,q}^{i,k})$: the makespan of manufacturer i in scheduling the set of demand K.
TTR_r^i	$TTR_r^i = \max(\max_{k,q:t^k>0}(C_{r,q}^{i,k} + p_r^i), 0)$: the time to recover of manufacturer i through path r.
TTR^i	$TTR^i = \max_r TTR_r^i = \max_k TTR^{i,k} = \max_{r,k,q:t^k>0} \alpha_r^i(C_{r,q}^{i,k})$: the time to recover of manufacturer i.
TTR^E	$TTR^E = \max_{i \in N(E)} TTR^i$: the time to recover of Tier E.

2 Polynomial time representation for inputs

Table 2: Polynomial time representation for $C^{i,k}_{r,q}$

Production Completion Time	1	$\dots$	u^i	$u^i + p_i$	$u^i + 2p_i$	$\dots$	$u^i + \left\lceil \frac{\lvert K\rvert}{cap^i} \right\rceil \times p_i$
Production Capacity	0	0	cap^i	cap^i	cap^i	cap^i	$\lvert K\rvert - \left\lfloor \frac{\lvert K\rvert}{cap^i} \right\rfloor \times cap^i$
Cumulative Production Capacity	0	0	cap^i	$2cap^i$	$3cap^i$	$\dots$	$\lvert K\rvert$

Table 3: Polynomial time representation for D^k

Time	1	2	$\dots$	H
Local demand quantity	$\sum_{k,r:D^{i,k}_r=1} Q^{i,k}_r$	$\sum_{k,r:D^{i,k}_r=2} Q^{i,k}_r$	$\dots$	$\sum_{k,r:D^{i,k}_r=H} Q^{i,k}_r$
Cumulative demand	$\sum_{k,r:D^{i,k}_r\le 1} Q^{i,k}_r$	$\sum_{k,r:D^{i,k}_r\le 2} Q^{i,k}_r$	$\dots$	$\lvert K\rvert$

3 Minimizing the Maximum Tardiness: Proof of Optimality of the EDD Rule

These additional decision variables, based on our completion times, are necessary for this proof.

$T^{i,k}_{in} = \max_{s\in pred(i), r\in R(s):\, i\in r, q\in Q_{(s,i)}} t^{s,k}_{r,q}$: the tardiness of manufacturer i in meeting demand k based on the arrival of needed components, before i makes its production/shipping decisions.

$T^{i}_{in} = \max_k \max_{s\in pred(i), r\in R(s):\, i\in r, q\in Q_{(s,i)}} t^{s,k}_{r,q}$: the tardiness of manufacturer i across all demands based on the arrival of needed components, before i makes its production/shipping decisions.

$T^{E}_{in} = \max_{i\in N(E)} T^i_{in}$: the maximum tardiness of all manufacturers i in Tier $E \in \mathcal{E}$ based on the arrival of needed components, before the manufacturers make their production/shipping decisions.

$T^{i}_{dis} = \max_{k,r}(u^i + p_i - D^{i,k}_r, 0)$: the tardiness of manufacturer i across all demands based on the disruption at i.

$T^{E}_{dis} = \max_{i\in N(E)} T^i_{dis}$: the maximum tardiness of all manufacturers s in Tier $E \in \mathcal{E}$ based on their disruptions.

$T^{i,k}_{out} = \max_{r,q} t^{i,k}_{r,q}$: the maximum tardiness of manufacturer i in meeting demand k based on its shipping decisions.

$T^{i}_{out} = \max_k T^{i,k}_{out}$: the maximum tardiness of manufacturer i across all demands based on its decisions.

Hence, $T_{out}^{B}$is the maximum tardiness of the whole system.

$T_{out}^{E} = \max_{i \in N(E)} T_{out}^{i}$: the maximum tardiness of all manufacturers s in Tier $E \in \mathcal{E}$ based on their decisions.

By definition, the tardiness at Tier E does not account for any further disruptions up the MEASC network, i.e., from Tier $E-1, \ldots, 1, 0$. Therefore, we have $C_{r,q}^{i,k} \geq \max_{s \in pred(i)} \left(C_{r,q}^{s,k} + p_{(s,i)}, u^{i} \right) + p_i$, i.e., the completion time of a task k at i is greater than or equal to the arrival of any required sub-component for task k or the disruption time of i plus the lead time of i. In addition, we have $T_{out}^{E} = \max_{i \in N(E)} T_{out}^{i,k} = \max_{i \in N(E)} \max_{k \in K} \max_{r \in R(i), q} t_{r,q}^{i,k} = \max_{i \in N(E), k \in K, r \in R(i), q} \left(C_{r,q}^{i,k} - D_{r}^{i,k}, 0 \right)$. We then have that

$$
\begin{aligned}
& T_{out}^{E} \geq \max_{s \in N(E+1), r \in R(s), q} \left(\max \left(C_{r,q}^{s,k} + p_{(s,i)}, u^{i} \right) + p_i - D_{r}^{i,k}, 0 \right) \\
\rightarrow\ & T_{out}^{E} \geq \max_{s \in N(E+1), r \in R(s), q} \left(\max \left(C_{r,q}^{s,k} + p_{(s,i)} + p_i - D_{r}^{i,k}, u^{i} + p_i - D_{r}^{i,k} \right), 0 \right) \\
\rightarrow\ & T_{out}^{E} \geq \max_{s \in N(E+1), r \in R(s), q} \left(\max \left(C_{r,q}^{s,k} - D_{r}^{s,k}, u^{i} + p_i - D_{r}^{i,k} \right), 0 \right) \\
\rightarrow\ & T_{out}^{E} \geq \max \left(T_{out}^{E+1}, T_{dis}^{E} \right)
\end{aligned}
$$

This means that for any Tier E, the tardiness is lower bounded by both the schedule-caused tardiness of the previous tier, T_{out}^{E+1}, and its own disruption-caused tardiness, T_{dis}^{E}.

The system then has the following property for any schedule:

$$
\begin{aligned}
T_{out}^{M-1} &\geq \max \left(T_{out}^{M}, T_{dis}^{M-1} \right) \\
T_{out}^{M-2} &\geq \max \left(T_{out}^{M-1}, T_{dis}^{M-2} \right) \geq \max \left(T_{out}^{M}, \max_{E=M-2, M-1} T_{dis}^{E} \right) \\
\ldots & \quad \ldots \quad \ldots \quad \ldots \\
T_{out}^{1} &\geq \max \left(T_{out}^{M}, \max_{E=1 \ldots M-1} T_{dis}^{E} \right) \\
T_{out}^{B} &\geq \max \left(T_{out}^{M}, \max_{E=0 \ldots M-1} T_{dis}^{E} \right)
\end{aligned}
$$

Let T^{M*} be the optimal maximum tardiness at tier M. If we can show that the EDD rule can achieve $T_{out}^{B} = \max \left(T_{out}^{M*}, \max_{E=0 \ldots M-1} T_{dis}^{E} \right)$, then $T_{out}^{B} = T_{out}^{B*}$, i.e, the EDD rule is the optimal solution.

Lemma 3.1. *Applying the EDD rule at all raw material manufacturers (Tier M) minimizes the maximum tardiness for tier M, $T_{out}^{M} = T_{out}^{M*}$.*

Proof. At Tier M, consider a schedule O, which violates the EDD, that is optimal. We can apply the a similar interchange argument to this problem as for the proof of the one machine problem $1||T_{max}$ in the traditional scheduling literature (Pinedo 2012) : In this schedule there must be at least two adjacent shipments by a manufacturer $s \in N(M)$ such that $C^{s,k_1}_{r_1,q_1} > C^{s,k_2}_{r_2,q_2}$ and $D^{s,k_1}_{r_1} < D^{s,k_2}_{r_2}$ (either $r_1 \neq r_2$ *or* $k_1 \neq k_2$). This implies that

$$\rightarrow \begin{cases} C^{s,k_1}_{r_1,q_1} - D^{s,k_1}_{r_1} > C^{s,k_2}_{r_2,q_2} - D^{s,k_1}_{r_1} \\ \\ C^{s,k_1}_{r_1,q_1} - D^{s,k_1}_{r_1} > C^{s,k_1}_{r_1,q_1} - D^{s,k_2}_{r_2} \end{cases}$$

The maximum tardiness across in Tier M based on the current schedule is

$$\begin{aligned} T^M_{out} &= \max_{i\in N(M)} T^{i,k}_{out} = \max_{i\in N(M)} \max_{\text{k}} \max_{r\in R(i),q} t^{i,k}_{r,q} \\ &= \max\left(0, \max_{i\in N(M),(i,k,r,q)\neq(s,k_1,r_1,q_1),(s,k_2,r_2,q_2)} t^{i,k}_{r}, C^{s,k_1}_{r_1,q_1} - D^{s,k_1}_{r_1}, C^{s,k_2}_{r_2,q_2} - D^{s,k_2}_{r_2}\right) \end{aligned}$$

If we exchange shipment $C^{s,k_1}_{r_1,q_1}$ and $C^{s,k_2}_{r_2,q_2}$ in our schedule O, then the maximum tardiness is then

$$\hat{T}^{\,M}_{out} = \max\left(0, \max_{i\in N(2),(i,k,r,q)\neq(s,k_1,r_1,q_1),(s,k_2,r_2,q_2)} t^{i,k}_{r,q}, C^{s,k_2}_{r_2,q_2} - D^{s,k_1}_{r_1}, C^{s,k_1}_{r_1,q_1} - D^{s,k_2}_{r_2}\right) \leq T^M_{out}$$

Therefore, schedule O can be modified to follow the EDD rule without increasing the objective function value, i.e., applying the EDD rule gives the optimal value at tier M, i.e., $T^M_{out} = T^{M*}_{out}$. □

Lemma 3.2. *The maximum tardiness does not increase between shipping from Tier $E+1$ and receiving at Tier E, $T^{E*}_{in} = T^{E+1*}_{out}$*

Proof. The maximum tardiness of all manufacturers s in Tier $E \geq 1$ based on the arrival of needed components is:

$$T^E_{in} = \max_{s\in N(E)} T^s_{in} = \max_{s\in N(E)} \max_{k} \max_{i\in pred(s), r\in R(i):\, s\in r} \max_{q=1..Q_r} t^{i,k}_{r,q}$$

Based on our assumptions, each path entering Tier E must be from a manufacturer in Tier $E+1$, we have

$$T_{in}^{E} = \max_{s} \max_{k} \max_{i \in pred(s), r \in R(i):\ s \in r} \max_{q=1..Q_r} t_{r,q}^{i,k} = \max_{i \in N(E+1)} \max_{\mathrm{k}} \max_{r \in R(i),q} t_{r,q}^{i,k} = T_{out}^{E+1}$$

$$\rightarrow T_{in}^{E*} = T_{out}^{E+1*}$$

□

Lemma 3.3. *It is optimal to apply the EDD at any Tier $E,\ 1 \leq E \leq M-1$ to minimize the maximum tardiness T_{out}^{E}. The optimal maximum tardiness is equal to either the tardiness of previous tier T_{out}^{E+1*}, or the the disruption-caused tardiness T_{dis}^{E}. $T_{out}^{E*} = \max\left(T_{out}^{E+1*}, T_{dis}^{E}\right)$.*

Proof. At Tier $E,\ 1 \leq E \leq M-1$, we have

$$T_{out}^{E} = \max_{i \in N(E)} T_{out}^{i,k} = \max_{i \in N(E)} \max_{\mathrm{k}} \max_{r \in R(i),q} t_{r,q}^{i,k} = \max_{i \in N(E)} \max_{\mathrm{k}} \max_{r \in R(i),q} \left(C_{r,q}^{i,k} - D_{r}^{i,k}, 0\right)$$

Using the EDD and applying the same proof as we did for the raw material manufacturers in Tier M, $T_{out}^{E,EDD} = T_{out}^{E*}$. For any schedule, we also have

$$T_{out}^{E} \geq \max\left(T_{out}^{E+1}, T_{dis}^{E}\right) \rightarrow T_{out}^{E} \geq \max\left(T_{in}^{E}, T_{dis}^{E}\right)$$

We will show $T_{out}^{E*} = \max\left(T_{in}^{E*}, T_{dis}^{E}\right)$, by showing $T_{out}^{i*} = \max\left(T_{in}^{i*}, T_{dis}^{i}\right) \forall i \in N(E),\ 1 \leq E \leq M-1$. We have

$$T_{in}^{i} = \max_{k} \max_{s \in in(i), r \in R(s):\ i \in r, q \in 1..Q_r} \left(C_{r,q}^{s,k} - D_{r}^{s,k}, 0\right)$$

$$T_{out}^{i} = \max_{k} \max_{r \in R(i), q \in 1..Q_r} \left(C_{r,q}^{i,k} - D_{r}^{i,k}, 0\right)$$

$$T_{dis}^{i} = \max_{k,r}\left(u^{i} + p_{i} - D_{r}^{i,k}, 0\right)$$

By way of contradiction, suppose $\exists i,\ T_{out}^{i*} > \max\left(T_{in}^{i*}, T_{dis}^{i}\right)$.

Let $\left\{\hat{s}, \hat{k}, \hat{r}, \hat{q}\right\} = \mathrm{argmax}_{k, s \in pred(i), r \in R(s):\ i \in r, q \in 1..Q_r} \left(C_{r,q}^{s,k*} - D_{r}^{s,k}, 0\right)$, i.e., this quadruple has the largest tardiness into manufacturer i. There must exist $\left\{\tilde{k}, \tilde{r}, \tilde{q}\right\}$ such that $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} - D_{\tilde{r}}^{i,\tilde{k}} > C_{\hat{r},\hat{q}}^{\hat{s},\hat{k}*} - D_{\hat{r}}^{\hat{s},\hat{k}}$ and $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} - D_{\tilde{r}}^{i,\tilde{k}} > u^{i} + p_{i} - D_{\tilde{r}}^{i,\tilde{k}}$.

Using the fact that we start production as soon as i is restored and all components are available,

$$C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} = \max\left(\max_{s\in pred(i)}\left(C_{\tilde{r},\tilde{q}}^{s,\tilde{k}*} + p_{(s,i)}\right), u^i\right) + p_i$$

If $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} = u^i + p_i$, then $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} - D_{\tilde{r}}^{i,\tilde{k}} = u^i + p_i - D_{\tilde{r}}^{i,\tilde{k}}$ contradicting $T_{out}^{i*} > T_{dis}^{i}$

If $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} > u^i + p_i$, let $\tilde{s} \in pred\,(i)$ be the manufacturer that has $C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} = C_{\tilde{r},\tilde{q}}^{\tilde{s},\tilde{k}*} + p_{(\tilde{s},i)} + p_i$. We then have

$$\begin{aligned} C_{\tilde{r},\tilde{q}}^{i,\tilde{k}*} - D_{\tilde{r}}^{i,\tilde{k}} &= C_{\tilde{r},\tilde{q}}^{\tilde{s},\tilde{k}*} + p_{(\tilde{s},i)} + p_i - D_{\tilde{r}}^{i,\tilde{k}} = C_{\tilde{r},\tilde{q}}^{\tilde{s},\tilde{k}*} - D_{\tilde{r}}^{\tilde{s},\tilde{k}} \\ \rightarrow C_{\tilde{r},\tilde{q}}^{\tilde{s},\tilde{k}*} - D_{\tilde{r}}^{\tilde{s},\tilde{k}} &> C_{\hat{r},\hat{q}}^{\hat{s}\hat{k}*} - D_{\hat{r}}^{\hat{s},\hat{k}}, \end{aligned}$$

which contradicts that $\left\{\hat{s}, \hat{k}, \hat{r}, \hat{q}\right\} = \text{argmax}_{k,s\in pred(i),r\in R(s):\; i\in r,q\in 1..Q_r} \left(C_{r,q}^{s,k*} - D_r^{s,k}, 0\right)$. Therefore, $T_{out}^{E*} = \max\left(T_{in}^{E*}, T_{dis}^{E}\right)$, which means $T_{out}^{E*} = \max\left(T_{out}^{E+1*}, T_{dis}^{E}\right)$. □

Theorem 3.4. *Applying the EDD rule for each individual manufacturer minimizes the maximum tardiness of the whole MEASC network, i.e., T_{out}^{B}.*

Proof. Lemma 3.1 shows that if the EDD rule is applied for all raw material manufacturers (Tier M) then the maximum tardiness for Tier M is optimal. Lemma 3.2 shows that the maximum tardiness does not increase between shipping from Tier $E+1$ and receiving at Tier E. Lemma 3.3 shows that as we move up the tiers and continue to apply the EDD rule, the maximum tardiness T_{out}^{E} at each Tier E is optimal, since it is either the optimal maximum tardiness of the previous Tier T^{E+1*}, or the disruption-caused tardiness T_{dis}^{E} of that tier. Consequently, the maximum tardiness out of the final assembly T_{out}^{B} is optimal as applying the EDD rule for each manufacturer in the whole system gives the best possible scheduling based on the disruptions at all the tiers, i.e., we have that

$$T_{out}^{B} = \max\left(T_{out}^{B+1*}, T_{dis}^{B}\right) = \max\left(T_{out}^{B+2*}, T_{dis}^{B+1}, T_{dis}^{B}\right) = \cdots = \max\left(T_{out}^{M*}, \max_{E=0\ldots M-1} T_{dis}^{E}\right) = T_{out}^{B*}.$$

□

4 Minimizing the Time to Recover: Proof of Optimality of the Decision Rule

We have the following additional decision variables that are necessary for this proof.

$\tau^{i,k} = \max_{r,q}(\alpha_r^i(C_{r,q}^{i,k}) - D^k, 0)$: the tardiness of manufacturer i for demand k based on the production/shipping decisions at i and the disruptions at upper tier manufacturers on paths from i.

$TTR_r^i = \max(\max_{k,q:t^k>0}(C_{r,q}^{i,k} + p_r^i), 0)$: the time to recover of manufacturer i through path r.

$TTR^{i,K} = \max_{r,q,k\in K:t^k>0}\alpha_r^i(C_{r,q}^{i,k})$: the time to recover of manufacturer i for a particular set of final demands K.

$TTR^i = \max_r TTR_r^i = \max_{r,k,q:t^k>0}\alpha_r^i(C_{r,q}^{i,k})$: the time to recover of manufacturer i.

$TTR^E = \max_{i\in N(E)} TTR^i$ the time to recover of Tier E.

Properties of any schedules: For any schedule, we have the following properties for the TTR^E:

$$TTR^E = \max_{i\in N(E)} TTR^i$$

$$TTR^E = \max_{i\in N(E),r,k,q:t^k>0} \alpha_i^r(C_{r,q}^{i,k})$$

$$TTR^E = \max_{i\in N(E),r,k,q:t^k>0} \left(C_{r,q}^{i,k} + p_r^i, \max_{m\in r}(u^m + p_r^m)\right)$$

It should be noted that the term $\max_{m\in r}(u^m + p_r^m)$ ensures TTR^E accounted for the disruptions up stream. This differs from how the maximum tardiness is defined. We also have

$$C_{r,q}^{i,k} \geq \max\left(\max_{s\in pred(i)}\left(C_{r,q}^{s,k} + p_{(s,i)}, u^i\right) + p_i\right)$$

$$\rightarrow TTR^E \geq \max_{s\in N(E+1),r\in R(s),q:t^k>0} \max\left(C_{r,q}^{s,k} + p_{(s,i)} + p_i + p_r^i, \max_{m\in r}(u^m + p_r^m)\right)$$

$$\rightarrow TTR^E \geq \max_{s\in N(E+1),r\in R(s),q:t^k>0} \max\left(C_{r,q}^{s,k} + p_r^s, \max_{m\in r}(u^m + p_r^m)\right)$$

$$\rightarrow TTR^E \geq TTR^{E+1}$$

This means that for any Tier $E \leq M-1$, the time to recover is lower bounded by that of the time for recover of Tier $E+1$. The system then has the following property for any schedule:

$$TTR^M \leq TTR^{M-1} \leq \cdots \leq TTR^2 \leq TTR^1 \leq TTR^B$$

We will show $TTR^B = TTR^M = TTR^{M*}$, thus implying $TTR^B = TTR^{B*}$, i.e., the decision rule is optimal.

4.1 On the Optimality of the Rule for a Single Raw Material Manufacturer

Theorem 4.1. *For a set of demands $\tilde{K}$, let $MS^{i,\tilde{K}} = \max_{r,k\in\tilde{K},q}\left(C_{r,q}^{i,k} + p_r^i, \max_{m\in r}(u^m + p_r^m)\right) = \max_{k\in\tilde{K}}(D^k + \tau^{i,k})$ be the makespan of manufacturer i if we assume that the makespan of the system only depends on the scheduling of i. Applying the LLT rule at a manufacturer i minimizes the makespan for i for the set of demands $\tilde{K}$.*

Proof. We have: $MS^i = \max_{r,k,q}\left(C_{r,q}^{i,k} + p_r^i, \max_{m\in r}(u^m + p_r^m)\right)$. At i, consider a schedule O, which violates the LLT rule. In this schedule there must be at least two adjacent shipments such that $C_{r_1,q_1}^{i,k_1} > C_{r_2,q_2}^{i,k_2}$ and $p_{r_1}^i > p_{r_2}^i$

$$\rightarrow \begin{cases} C_{r_1,q_1}^{i,k_1} + p_{r_1}^i > C_{r_2,q_2}^{i,k_2} + p_{r_1}^i \\ \\ C_{r_1,q_1}^{i,k_1} + p_{r_1}^i > C_{r_1,q_1}^{i,k_1} + p_{r_2}^i \end{cases}$$

The makespan of manufacturers in Tier M based on schedule O is

$$MS^{i,\tilde{K}} = \max_{r,k,q}\left(C_{r,q}^{i,k} + p_r^i, \max_{m\in r}(u^m + p_r^m)\right)$$
$$MS^{i,\tilde{K}} = \max_{r,k,q}\left(C_{r_1,q_1}^{i,k_1} + p_{r_1}^i, C_{r_2,q_2}^{i,k_2} + p_{r_2}^i, \max_{(s,k,r,q)\neq(i,k_1,r_1,q_1),(i,k_2,r_2,q_2)} C_{r,q}^{s,k} + p_r^s, \max_{m\in r}(u^m + p_r^m)\right)$$

If we exchange shipment C_{r_1,q_1}^{i,k_1} and C_{r_2,q_2}^{i,k_2} in our schedule O, the makespan is then

$$\widehat{MS^{i,\tilde{K}}} = \max_{r,k,q}\left(C_{r_2,q_2}^{i,k_2} + p_{r_1}^i, C_{r_1,q_1}^{i,k_1} + p_{r_2}^i, \max_{(s,k,r,q)\neq(i,k_1,r_1,q_1),(i,k_2,r_2,q_2)} C_{r,q}^{s,k} + p_r^s, \max_{m\in r}(u^m + p_r^m)\right)$$

which is less than or equal to $MS^{i,\tilde{K}}$, since $C_{r_1,q_1}^{i,k_1} + p_{r_1}^i > C_{r_2,q_2}^{i,k_2} + p_{r_1}^i$ and $C_{r_1,q_1}^{i,k_1} + p_{r_1}^i > C_{r_1,q_1}^{i,k_1} + p_{r_2}^i$. Therefore, schedule O can be modified to follow the LLT rule without increasing the objective function value, i.e., the LLT rule gives the optimal value $MS^{i,\tilde{K}} = MS^{i,\tilde{K}*}$. □

Lemma 4.2. *For a subset of demands $\bar{K}$ that must be late, applying the LLT rule to a manufacturer i minimizes the time to recover $TTR^{i(\bar{K})}$ at i for group $\bar{K}$.*

Proof. We know that $\bar{K}$ must be late. Therefore, $TTR^{i,(\bar{K})} > \max_{k \in \bar{K}} D^k$ meaning that there exists $k \in \bar{K}, t^{i,k} > 0$ for any schedule. We have that $TTR^{i,\bar{K}} = \max_{k\in\bar{K}}(D^k + t^{i,k}) = MS^{i,\bar{K}}$ From Theorem 4.1, we know that applying the LLT minimizes the makespan. Hence, the LLT rule minimizes the time to recover at manufacturer i, $TTR^{i,\bar{K}*} = MS^{i,\bar{K}*}$ for a subset of demands $\bar{K}$ that must be late. □

Preliminary Results for the Reverse EDD Rule

Theorem 4.3. *For any set of demands K, if there exists a schedule with maximum tardiness of 0, then the reverse EDD schedule has the maximum tardiness of 0.*

Proof. We know that applying the EDD rule minimizes maximum tardiness (with value 0), as shown in Theorem 3.4. From the EDD schedule, if we exchange production completion time $C^{i,\tilde{k}}_{\tilde{r},\tilde{q}}$ assigned to the last demand $\tilde{k}$ (i.e., the demand with largest deadline) with largest production completion time $C^{i,\bar{k}}_{\bar{r},\bar{q}}$ that is less than or equal to $D^{i,\tilde{k}}_{\tilde{r}}$. Then the tardiness for $\tilde{k}$ is still 0. Since $C^{i,\bar{k}}_{\bar{r},\bar{q}} \leq C^{i,\tilde{k}}_{\tilde{r},\tilde{q}}$, the maximum tardiness for the remaining demands in K is still 0. Hence, the EDD schedule can be transformed into the reverse EDD schedule while maintaining the maximum tardiness of 0. □

Determining the Time to Recover of the Raw Materials manufacturers

Theorem 4.4. *For each raw material manufacturer i, decomposing the subproblems at $k(i)$ generated by the algorithm in Subsection 5.1 gives the optimal local time to recover $TTR^{i*} = \mathcal{C}(i)$, where TTR^{i*} is the optimal time to recover considering just manufacturer i and $\mathcal{C}(i)$ is the time to recover given by the algorithm.*

Proof. We have $K2 = \{k(i)+1, k(i)+2, \ldots |K|\}$ is the on-time subset of demand since there is a reverse EDD schedule with maximum tardiness of 0 (Theorem 4.3). Let $K1 = \{1, 2, \ldots, k(i)\}$ be the subset of remaining demands. We have two situations:

Case 1: $K1 = \emptyset$. Then, $TTR^{i*} = TTR^{i,K1*} = 0 = \mathcal{C}(i)$ since all jobs can be completed on time.

Case 2: $K1 \neq \emptyset$ We want to show that $K1$ must be late, i.e., $MS^{i,K1*} > D^{k(i)}$. Suppose $MS^{i,K1*} \leq D^{k(i)}$. This implies that $MS^{i,\{1,2,\ldots k(i-1)\}*} \leq D^{k(i)}$. Because the iteration did not break at the previous iteration, that means it is not possible to find a schedule with maximum tardiness of 0 for the subset of demands $\{k(i), k(i)+1, \ldots |K|\}$.Since the maximum tardiness of $K2 = \{k(i)+1, k(i)+2, \ldots |K|\}$ is 0, this implies that the tardiness of $k(i)$ must be positive. Hence, for the subset $K1 = \{1, 2, \ldots, k(i)\}$, $MS^{i,K1*} > D^{k(i)}$, or $K1$ must be late. Since $\mathcal{C}(i)$ is generated by the LLT rule and $K1$ must be late, $TTR^{i,K1*} = MS^{i,K1*} = \mathcal{C}(i)$ from Theorem 4.1 and Lemma 4.2. We have $TTR^{i*} = TTR^{i,K1*} = \mathcal{C}(i)$. □

Theorem 4.5. *Let $\bar{\mathcal{C}} = \max_{i \in N(M)} \mathcal{C}(i)$ be the maximum optimal individual TTR across raw material manufacturers i calculated by the algorithm in Subsection 5.1. Decomposing the main problem into 2 subproblems based on $\bar{k}$ and applying the decision rules at a raw material manufacturer i gives an updated time to recover, $TTR^i \leq \bar{\mathcal{C}}$*

Proof. Let $K1 = \bar{K} = \{k \in K | k \leq \bar{k}\}$ and $K2 = \{k \in K | k \geq \bar{k} + 1\}$. If $k(i) = \bar{k}$, from Theorem 4.4, we have that $TTR^{i*} = \mathcal{C}(i) \leq \bar{\mathcal{C}}$ since the sets $K1$ and $K2$ will not change.
Now, consider the case where $k(i) < \bar{k}$. Applying the LLT rule for $K1$ gives $TTR^{i,K1} = MS^{i,\{k \in K1 | t_k > 0*\}} \leq MS^{i,K1*} \leq D^{\bar{k}} < \bar{\mathcal{C}}$ where the second to last inequality comes directly from the definition of $k(i)$ (i.e., $k(i) < \bar{k}$ implying that $MS^{i,K1*} \leq D^{\bar{k}}$). In other words, the schedule that was created for manufacturer i initially for K1 had all tasks being completed by $D^{\bar{k}}$ and applying the LLT rule can only decrease the makespan for this set of jobs at i. □

4.2 The decision rules minimize TTR for the whole system: Proofs

Subproblem 1: For subset of final demands $k \leq \bar{k}$

We need to show that the LLT rule minimizes the time to recover of the whole MEASC network, i.e., TTR^B for the subset of final demands with deadline before cutoff $\bar{k}$ determined by the algorithm in Subsection 5.1, i.e. the subset of final demands that must be late. We will show that applying the LLT rule for all raw material manufacturers gives the optimal value TTR^{M*} (Theorem 4.6). Then, Theorem 4.7 will show that as we move up the tiers and continue to apply the LLT rule, the time to recover at a Tier E does not increase and equals the optimal value of the previous tier time to recover, $TTR^{E*} = TTR^{E+1*}$. Combining these two theorems, we can show that the LLT rule gives optimal value of the time to recover at the final assembly node $TTR^{B*} = TTR^{M*} = \bar{\mathcal{C}}$.

Theorem 4.6. *For $SP1$ where the set of demands that must be late are those such that $k \leq \bar{k}$, applying the LLT rule for all raw material manufacturers gives the optimal value $TTR^M = TTR^{M*} = \bar{\mathcal{C}}$.*

Proof. From Theorem 4.5, we know that applying the decision rule gives $TTR^i \leq \bar{\mathcal{C}} \; \forall \, i \in N(M)$ with $\bar{\mathcal{C}} = \max_i TTR^i$. Thus, $TTR^{M*} = \max_{i \in N(M)} TTR^i = \bar{\mathcal{C}}$. Therefore, $TTR^{M*} = \bar{\mathcal{C}}$ is the lowest time to recover that can be obtained at Tier M since TTR^i is a lower bound on the overall time to recover. □

Theorem 4.7. *For SP1, as we moves up the tiers, applying the LLT rule does not increase the time to recover, i.e.,* $TTR^{E*} = TTR^{E+1*}$

Proof. At Tier $E,\ 1 \leq E \leq M-1$, we have

$$TTR^E = \max_{i \in N(E),r,k,q} \left(C^{i,k}_{r,q} + p^i_r, \max_{m \in r} \left(u^m + p^m_r \right) \right)$$

Using the LLT rule and applying the same proof as in Theorem 4.4, $TTR^E = TTR^{E*}$. We want to show that $TTR^{E*} = TTR^{E+1*}$ by showing that $TTR^{i*} = \max_{s \in pred(i)} TTR^{s*} \forall\, i \in N(E),\ 1 \leq E \leq M-1$. We have $TTR^i = \max_r TTR^i_r = \max_{r,k,q:t^k>0} \alpha^i_r (C^{i,k}_{r,q})$. By way of contradiction, suppose $\exists i$, where $TTR^i > \max_{s \in pred(i)} TTR^s$. This implies that

$$\max_{r,k,q} \left(C^{i,k}_{r,q} + p^i_r, \max_{m \in r} \left(u^m + p^m_r \right) \right) > \max_{s \in pred(i), r \in R(s): i \in r,k,q} \left(C^{s,k}_{r,q} + p^s_r, \max_{m \in r} \left(u^m + p^m_r \right) \right)$$

$$\rightarrow \begin{cases} \max_{r,k,q} \left(C^{i,k}_{r,q} + p^i_r \right) > \max_{s \in pred(i), r \in R(s): i \in r,k,q} \left(C^{s,k}_{r,q} + p^s_r \right) \\ TTR^i = \max_{r,k,q} \left(C^{i,k}_{r,q} + p^i_r \right) \\ \max_{r,k,q} \left(C^{i,k}_{r,q} + p^i_r \right) > \max_{m \in r} \left(u^m + p^m_r \right) \end{cases}$$

Let $\left\{ \hat{k}, \hat{r}, \hat{q} \right\} = \operatorname{argmax}_{r,k,q} \left(C^{i,k}_{r,q} + p^i_r \right) = TTR^i$

$$\rightarrow \begin{cases} C^{i,\hat{k}}_{\hat{r},\hat{q}} + p^i_{\hat{r}} > \max_{s \in pred(i), r \in R(s): i \in r,k,q} \left(C^{s,k}_{r,q} + p^s_r \right) \\ C^{i,\hat{k}}_{\hat{r},\hat{q}} + p^i_{\hat{r}} > \max_{m \in r} \left(u^m + p^m_r \right) \end{cases}$$

Because we start production as soon as all components are available or when the manufacturer is restored, $C^{i,\hat{k}}_{\hat{r},\hat{q}} = \max \left(\max_{s \in pred(i)} \left(C^{s,k}_{r,q} + p_{(s,i)} \right), u^i \right) + p_i$ if the ordering $C^{s,k}_{r,q}$ and $C^{i,k}_{r,q}$ are identical. Note that since $\forall s \in pred(i), r \in R(s), p^s_r = p^i_r + p_{(s,i)}$, the order of the schedule in the LLT rule of s and i based on r should be identical because if $p^s_{r_1} > p^s_{r_2}, p^s_{r_1} + p_{(s,i)} > p^s_{r_2} + p_{(s,i)} \rightarrow p^i_{r_1} > p^i_{r_2}$. Therefore, we have a contradiction. Hence, $TTR^{E*} = TTR^{E+1*}$. □

Theorem 4.8. *The LLT rule minimizes the time to recover of the whole system for demands before or at cutoff $\bar{k}$.*

Proof. Combining Theorems 4.6 and 4.7 results in the LLT rule achieving $TTR^{B*} = TTR^{M*} = \bar{\mathcal{C}}$. □

SP2: For the demands after $\bar{k}$

Applying the EDD rule achieves optimal the maximum tardiness of 0 as shown in Theorem 3.4 since the process of finding $\bar{k}$ includes verifying that there exists a schedule with tardiness of 0 for every demand after $\bar{k}$. Therefore, $TTR^{B} = \bar{\mathcal{C}}$ remains.

5 Integer Programming Models

This section shows the formulations of the IP models that we use as a comparison with the decision rules. We begin with the required additional notation:

N': the set of nodes that need to be restored.

$BigM$: a sufficiently large number

qd^{t}: the quantity of final products with demand deadline of t

$ship_{(i,j),t}, (i,j) \in A$: the variable presents amount part i is shipped on arc (i,j) at time t.

$pro_{i,t}, i \in N$: the variable presents amount of part i produced at node i at time t where t is the time production ends.

$inv_{i,z,t}, i \in N$: the variable presents amount of part z stored at node i at time t.

5.1 Minimizing Maximum Tardiness

Model specific Notations:

$fd^{d,t}$: the variable presents the quantity of the final products received at B at time t, allocated to meeting the final demands with deadline d

$ud^{d,t}$: the variable presents the quantity of the final demands with deadline d remaining unmet at time t

$tbo^{d,t}$: the variable presents total amount of back-ordered final demands with deadline d at time t.

$ed^{d,t}$:the variable presents total amount of excess final demands with deadline d at time t.

$pbo^{d,\delta}$: the binary variable indicates whether there is a positive amount of back-ordered final demands with deadline d at time $t+\delta$. i.e., δ is the tardiness.

$$\text{Minimize } T^B_{out} \tag{1a}$$

$$\text{s.t.} \quad pro_{i,t} + inv_{i,i,t-1} = \sum_{(i,j)\in A} ship_{(i,j),t} + inv_{i,i,t}, \qquad \forall i \in N, t = 1\dots H \tag{1b}$$

$$ship_{(s,i),t-p_{(s,i)}} + inv_{i,s,t-1} = Q_{(s,i)} pro_{i,t+p_i} + inv_{i,s,t} \quad \forall i \in N, s \in pred(i), t = 1\dots H \tag{1c}$$

$$\sum_{(i,B)\in A} ship_{(i,B),t-p_{(i,B)}} = \sum_{d=1}^{H} fd^{d,t} \qquad \forall t = 1\dots H \tag{1d}$$

$$fd^{t,t} + ud^{t,t} = qd^t + ed^{t,t} \qquad \forall t = 1\dots H \tag{1e}$$

$$fd^{d,t} + ud^{d,t} = ed^{d,t} \qquad d = 1\dots H, \forall t = 1\dots H, t \neq d \tag{1f}$$

$$tbo^{d,t} = tbo^{d,t-1} + ud^{d,t} - ed^{d,t} \qquad d = 1\dots H, t = 2\dots H, \tag{1g}$$

$$pbo^{d,\delta} BigM \geq tbo^{d,d+\delta} \qquad \forall d = 1\dots H, \delta = 0\dots H-d \tag{1h}$$

$$T^B_{out} \geq \sum_{\delta=1}^{H-\delta} pbo^{d,\delta} \qquad \forall d = 1\dots H \tag{1i}$$

$$pro_{i,t} \leq cap^i \qquad \forall i \in N \setminus N(M), t = 1\dots H \tag{1j}$$

$$pro_{i,t} = 0 \qquad \forall i \in N', t = 1\dots u^i \tag{1k}$$

$$ship_{(i,j),t}, pro_{i,t}, inv_{i,z,t}, fd^{t,t}, ud^{d,t}, tbo^{d,t}, ed^{d,t} \geq 0 \tag{1l}$$

$$pbo^{d,\delta} binary \tag{1m}$$

5.2 Minimizing the time to recover

Model specific Notations:

ud^t: the variable presents unmet final demand at time t.

ed^t:the variable presents excess final demand at time t.

tbo^t: the variable presents total amount of back-ordered product at time t.

pbo^t: the binary variable indicates whether there is a positive backorders(1 if yes) at time t, i.e. there is a late demand

dis^t: the binary variable indicates (with value 1) if the supply chain has not recovered at time t

$$\text{Minimize } TTR^B \tag{2a}$$

$$\text{s.t.} \quad pro_{i,t} + inv_{i,i,t-1} = \sum_{(i,j)\in A} ship_{(i,j),t} + inv_{i,i,t}, \qquad \forall i \in N, t = 1 \dots H \tag{2b}$$

$$ship_{(s,i),t-p_{(s,i)}} + inv_{i,s,t-1} = Q_{(s,i)} pro_{i,t+p_i} + inv_{i,s,t} \quad \forall i \in N, s \in pred(i), t = 1 \dots H \tag{2c}$$

$$\sum_{(i,B)\in A} ship_{(i,B),t-p_{(i,B)}} + ud^t = qd^t + ed^t \qquad t = 1 \dots H \tag{2d}$$

$$tbo^t = tbo^{t-1} + ud^t - ed^t \qquad t = 2 \dots H \tag{2e}$$

$$pbo^t BigM \geq tbo^t \qquad t = 1 \dots H \tag{2f}$$

$$pbo^t \leq dis^t \qquad t = 1 \dots H \tag{2g}$$

$$TTR^B \geq t\, dis^t \qquad t = 1 \dots H \tag{2h}$$

$$pro_{i,t} \leq cap^i \qquad \forall i \in N \setminus N(M), t = 1 \dots H \tag{2i}$$

$$pro_{i,t} = 0 \qquad \forall i \in N', t = 1 \dots u^i \tag{2j}$$

$$ship_{(i,j),t}, pro_{i,t}, inv_{i,z,t}, ud^t, tbo^t, ed^t \geq 0 \tag{2k}$$

$$pbo^t, dis^t \quad \text{binary} \tag{2l}$$