



Department of Management Science  
Otto von Guericke University Magdeburg

**Scientific Project**

**Minimizing the production and logistics costs for a microchip manufacturer by optimal facility and carrier assignment**

**Author:** [REDACTED], [REDACTED]

**Email address:** [REDACTED], hamzauddin.siddiqui@st.ovgu.de,  
[REDACTED]

**Matriculation number:** [REDACTED]

**Study program:** Betriebswirtschaftslehre M.Sc., Operations Research & Business Analytics M. Sc.

**Examiner:** [REDACTED]

**Supervisor:** [REDACTED]

*Submission date: 23.01.2026*

## List of Figures

<b>Figure 1:</b> Delivery window size for orders (own illustration).....	8
<b>Figure 2:</b> Delivery window size for orders (own illustration).....	12
<b>Figure 3:</b> Delivery window size for orders (own illustration).....	12
<b>Figure 4:</b> Service level breakdown (own illustration).....	13
<b>Figure 5:</b> Average shipping Cost DTP (own illustration).....	32
<b>Figure 6:</b> Average shipping Cost DTD (own illustration).....	32
<b>Figure 7:</b> Plant daily capacity vs cost per unit (own illustration).....	14
<b>Figure 8:</b> Graphical representation of the outbound supply chain network (own illustration).....	18
<b>Figure 9:</b> Total cost breakdown across three components (own illustration).....	23
<b>Figure 10:</b> Plant capacity utilization heatmap by day with color intensity indicating utilization (own illustration).....	24
<b>Figure 11:</b> Distribution of capacity penalty costs across three progressive tears.....	24
<b>Figure 12:</b> Temporal demand concentration pattern across the planning horizon (own illustration).....	25
<b>Figure 13:</b> Order distribution across carriers (own illustration).....	26
<b>Figure 14:</b> Cost structure comparison showing the baseline solution's incomplete cost accounting (own illustration).....	27
<b>Figure 15:</b> Order distribution comparison across manufacturing plants, demonstrating load redistribution from extreme PLANT03 concentration (own illustration)...	28

---

## List of Tables

<b>Table 1:</b> Formulation of sets (own illustration).....	17
<b>Table 2:</b> Formulation of parameters (own illustration).....	18
<b>Table 3:</b> Optimized solution cost breakdown (own illustration).....	23
<b>Table 4:</b> Comparison of Baseline and Optimized Solutions (own illustration).....	26

# Table of Contents

List of Figures .....	2
List of Tables .....	3
1. Introduction.....	6
2. Data Exploration and Analysis.....	7
2.1 Data Description.....	8
2.1.1 Customer Requirements.....	8
2.1.2 Logistics Costs .....	9
2.1.3 Manufacturing Capabilities .....	11
2.2 Data Analysis .....	12
2.2.1 Customer Requirements.....	12
2.2.2 Logistics Costs .....	13
2.2.3 Manufacturing Capabilities .....	14
3 Problem Setting .....	15
3.1 Problem Description.....	15
3.2 Visualization Through A Network .....	16
3.3 Mathematical Notation .....	16
3.4 Mathematical Model .....	18
4. Results and Benchmarking .....	19
4.1 Computational Setup .....	20
4.2 Model Evolution and Infeasibility Resolution.....	20
4.2.1 Initial Hard Capacity Model and Infeasibility Discovery.....	20
4.2.2 Soft Capacity Constraint Formulation.....	21
4.2.3 Three-Tier Progressive Penalty Structure.....	22
4.3 Computational Performance .....	22
4.4 Optimized Solution Analysis.....	22
4.4.1 Overall Cost Structure .....	22
4.4.2 Capacity Penalty Tier Distribution .....	23
4.4.3 Temporal Distribution Patterns.....	25
4.4.4 Carrier Selection and Transportation Mode .....	25
4.5 Benchmarking Against Baseline Solution.....	26

4.5.1 Baseline Solution Characteristics .....	26
4.5.2 Comparative Analysis and Cost Reconciliation.....	27
4.5.3 Operational Value and Strategic Implications .....	28
4.6 Summary of Key Findings .....	29
5. Conclusion and Future Outlook.....	29
Bibliography .....	30
Appendix A: AI Usage Declaration.....	31
Appendix B.....	32
Distribution of Work .....	33

## 1. Introduction

The importance of an efficient, coordinated production and distribution network is increasing with its crucial role in performance of supply chains. These global and complex networks are challenged by productions with varying conditions such as costs, capacities, product availabilities and their connected outbound logistics. Especially in microchip manufacturing, where high-value components are being handled with time-sensitive deliveries, that have a desire for strict shipment requirements. Even small mismatches between demand timing and production capabilities can cause costly incidents. It is in the interest of the microchip supplier to try to save costs on production of the goods, their delivery and avoid poor customer service. Further, the customer pursues the goal that the ordered microchips are delivered in the requested time window to not cause any delay in their own value creation. A wide range of electronic industries relies on a steady supply of those parts. This creates a certain responsibility on the model, which optimizes the supply chain and plays a big role in gaining a competitive advantage through delivery on time.

Current real-world supply chains are characterized by multiple products and distribution centres that need to be managed. Therefore, the focus is on developing models that incorporate that. To measure the performance of a supply chain, the total cost minimization is a commonly used dimension (Fahimnia et al., 2013). The survey study by Mönch et al. (2018) shows that production planning and demand fulfilment are part of the most challenging decisions in microchip supply chains. Especially, capacity constraints for machines at plants specialized on certain products can emerge as binding bottlenecks. The formulation of assignment and routing decisions in an integrated production distribution model allows for optimizing the plant selection with its capacity usage and the linked logistics costs. Integrated production distribution models are efficient in recording cost trade-offs along the supply chain stages, particularly when production and distribution decisions are tightly linked (Fahimnia et al., 2013). Geopolitical events like trade conflicts and industrial shifts can have an impact on the microchip supply chain, resulting in volatile demand patterns (Xiong et al., 2025). While other models like the one proposed by Dzalbs and Kalganova (2020) often assume sufficient production capacity, real-world production often violates these assumptions and therefore have infeasible solutions. Past research commonly states due dates as the latest time for the finished deliveries. Additionally earliest delivery dates are used to construct a time window. A common practice is to introduce soft constraints to penalize any violation. Simultaneously, the supplier should maximize their capacity usage. Both aims compete in a trade-off and often even near optimal solutions conflict with a constraint (Ulrich, 2013).

This paper proposes a mixed-integer linear program-based model to minimize the production and logistics costs for a microchip manufacturer by optimal facility and carrier assignment, incorporating real-world constraints. While minimizing the total costs, we need to ensure feasibility for managing the assignment of each order to a plant and its further transportation to the customer. In the scope of this computational study, the key contribution is a reformulation of the hard plant capacity limits with a three-tier penalty structure to allow a controlled exceedance. Feasibility is even then ensured, when the customer demand exceeds the nominal plant capacity, while making the overtime costs visible instead of hiding them. This modelling approaches a closer step to economic realism while pursuing the clear objective to optimize costs. Computational results demonstrate an improvement in warehouse and freight costs. Additionally, and more importantly, is the quantification of the hidden costs from the plant capacity exceedance, while having a feasible solution that supports more transparent planning decisions.

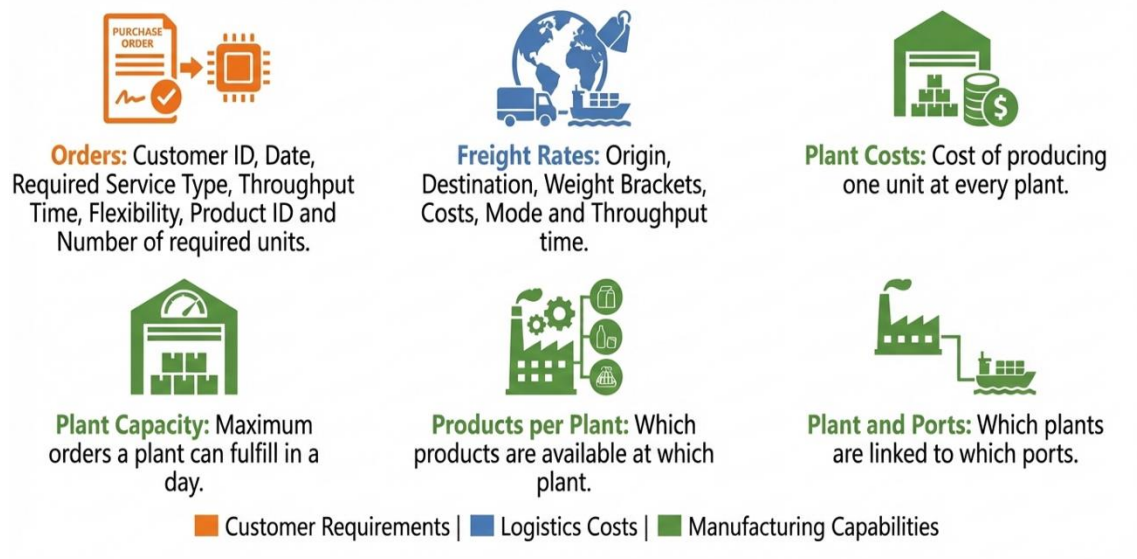
The remainder of this paper is structured as follows. Chapter 2 reviews the given dataset with a description and shares insights into the analysis. Then in Chapter 3, we first describe the problem verbally, followed by a graphical representation of the network and lastly, we display the mathematical model. Chapter 4 proceeds the computational study by firstly stating the computational setup, then by expanding the model with a penalty structure and thirdly by stating the computational performance. Afterwards, the optimized solution is analysed and benchmarked against the baseline solution of Dzals and Kalganov (2020). It ends with a summary of the key findings.

## 2. Data Exploration and Analysis

In this study, we utilize a real-world dataset from a global microchip manufacturer, published by Brunel University of London, representing a complex logistics optimization problem. The operational scope of this optimization problem involves a supply chain network which aims to fulfill 9,215 customer orders. These orders are processed through a decentralized manufacturing infrastructure with 19 plants. Once the products are completed, they are routed through a logistics funnel consisting of 11 origin ports and 9 unique carriers. The entire network converges at a destination port. The data also provides the limitations of this network by sharing the capacity limits of all plants and varied freight costs depending on weight brackets and service types.

To provide a clear picture of the available data, the following figure 1 maps out the key variables. It categorizes them into three main domains: customer requirements, manufacturing capabilities,

and logistics costs. In the next two subsections, we explore these three domains, first by explaining all the information that is available and then by analyzing and drawing insights from it.



**Figure 1:** Delivery window size for orders

## 2.1 Data Description

In this section, we explain all three domains of the data and discuss wherever we needed to correct or complete the available details. The customer requirements domain gives us the demand side of the problem and logistics costs along with manufacturing capabilities provide the supply side.

### 2.1.1 Customer Requirements

The customer requirements help us understand the scale of the task we are dealing with. As already mentioned, we have a total of 9,215 orders, but to understand the customer requirements better, we provide a short explanation of all the relevant elements.

**Order ID:** Every order has a ten-digit unique ID that differentiates it from other orders.

**Order Date:** The order date is also available, it is useful because to determine whether an order was delivered timely, we need to keep track of when the order was placed.

**Throughput time:** Throughput time is assumed to be the time (in days) after the order date, after which the customer wants the order.

**Service Level:** The dataset has three different types of service levels, which are Customer Referred Freight (CRF), Door to Door (DTD), and Door to Port (DTP). For CRF orders, the

customer arranges the freight themselves. In case of DTD orders, the customer receives the products at their doorstep, and for DTP, the customer picks up the order from the destination port.

**Ship ahead day count:** Ship ahead day count is assumed to be the flexibility of the customer regarding when they want the order. It represents the number of days **before** *Order date + TPT* in which the customer will accept an order.

**Ship late day count:** Similar to ship ahead day count, this also represents flexibility, but it represents the number of days **after** *Order date + TPT* in which the customer will accept an order.

**Customer ID:** The customer ID is also unique for every customer, some of the customers can only be served from their specified warehouses. These special customers are referred to as Vendor-Managed Inventory (VMI) customers.

**Product ID:** Every order also has a product ID, and it plays a vital role when selecting a plant to serve the order because not every product is available at every plant.

**Unit Quantity:** This tells us how many units of a product an order needs.

**Weight:** Weight of every order and is crucial because the logistics costs are based on weight.

The information on customer requirements is detailed. However, we still need to calculate the time windows because we want to determine when every order needs to be delivered to prioritize the manufacturing and shipping of the orders that are to be delivered earlier. We calculate the time windows and obtain the Earliest Delivery Date, indicating the earliest day the customer will accept the order, and the Latest Delivery Date, showing the last date the customer will accept an order.

$$\text{Earliest Delivery Date} = \text{Order Date} + \text{TPT} - \text{Ship Ahead Day Count}$$

$$\text{Latest Delivery Date} = \text{Order Date} + \text{TPT} + \text{Ship Ahead Day Count}$$

After adding these datapoints, we had everything that we needed from the demand side.

### 2.1.2 Logistics Costs

Moving onto supply side of the available data, first we have the logistics costs through which we understand the price of transporting an order from an origin port to the destination port. We now describe all the elements of the logistics cost like we did for customer requirements.

**Carrier Type:** We only have two carrier types; we assume these to be broader service like DHL or DPD which offer different types of services for different prices e.g. they can have two services,

a cheaper one and an expensive one. The cheaper service takes longer to deliver a product and offers little to no damage insurance of the contents but on the other hand, the expensive service delivers the product in a shorter timeframe and guarantees the safety of the transported goods.

**Carrier:** There are nine carrier options available in our data, two of them belong to one carrier type and the other seven belong to the other one. These can be different services of DHL such as economy, standard, express or letter shipping, all with a different service and price structure.

**Origin and Destination Port:** As mentioned earlier, the logistics information provides us with rates for moving orders from several origin ports to the destination port. We have 10 unique combinations of origin and destination port.

**Minimum and Maximum Weight Quantity:** Another factor that determines logistics costs is the weight of the shipment. For example, a carrier might charge one rate for an order weighing between 10kg and 50kg, but a different rate (often lower per kg) if the order is heavier and falls into the 50kg to 100kg bracket.

**Service Type:** Both the relevant service types have different rates, the rates for DTD being higher.

**Minimum Cost:** Even for light shipments, carriers incur fixed handling and administrative costs. To account for this, the dataset includes a **Minimum Cost** for every weight bracket. This is a "base fee" or a starting fare; regardless of how small the order is, the cost will never drop below this.

**Rate:** While the brackets determine which "tier" we fall into, the **rate** is the variable cost assigned to that specific tier. We calculate the exact shipping cost using this rate using the total weight of microchips in an order.

**Mode of Transport:** This describes the physical method used to move the goods. In our dataset, the predominant mode is **AIR**.

**Logistics Throughput Time (TPT):** Not all carriers travel at the same speed. Each shipping lane has a logistics TPT, which represents the estimated number of days the product spends in transit from the origin port to the destination port. The optimization model must balance this transit time against the customer's delivery deadline to ensure we don't choose a cheap but slow carrier that results in a late delivery.

The data for logistics costs was almost complete except for one instance which was that even though we have plants that can send their products to Port 1, we did not have the costs from this

port to the destination Port 9. Since all the connections between other ports and destination Port 9 were different in nature (different weight brackets, service types, mode of transport, etc.), taking an average was not feasible. We decided to leave these costs blank because the plants connected to this port had connections to other ports as well.

Now that we have explained how much it costs and how long it takes to ship products, we look at the manufacturing side to see where and how the products are made.

### 2.1.3 Manufacturing Capabilities

Having discussed one half of the supply side by looking at the logistics costs, we move towards the other side which is manufacturing. To talk about the production facilities, we use the terms warehouse and plants interchangeably throughout this paper. This section answers a few important questions for us such as: How much can we produce? When can we produce it? How much does it cost to produce the required quantities? Having answers to these helps us understand the ability to fulfill demand. We describe the available information on this section below:

**Warehouse Costs:** Firstly, we have the cost of producing one unit at every warehouse. While the plants maintain a uniform production cost for every product, these costs are different for every plant. The lowest cost of producing one unit is 37 cents while the maximum is 2.04 dollars.

**Warehouse Capacities:** All the warehouses have different daily capacities, and we cannot produce everything at the warehouses with lower costs. The lowest capacity is 7 and the highest capacity for a warehouse is 1,070. We assume that this capacity is for orders not units because these numbers are very low. Assuming these numbers to be units makes the problem infeasible.

**Products Per Plant:** This provides us with information on which products can be manufactured at which plants. Not all products are available at all plants. The highest number of unique products available at a plant is 781 and the lowest number is 1.

**Plants Port Connection:** Here we get the information regarding the connections between ports and plants. Once the products have been manufactured at plants, they can be moved to the ports (origin ports) linked to them. Only 3 plants are linked to two different ports, the rest of them only have one port. Since we do not have the costs of transporting products from plants to ports, we assume that it is free of cost.

Reflecting on the completeness and correction for data on manufacturing capabilities, we found one instance of incomplete data which was in the Products Per Plant section. We had a plant

(CMD9) in this data, which was different from our other plants, and one of our plants was missing from this table. Therefore, we replaced CMD9 with our missing plant to complete the information.

In this section, we also discussed how we corrected some of the missing or confusing information to prepare the data for the model. Since the entire dataset is defined now, the next section analyzes these numbers to highlight potential challenges and bottlenecks which we need to solve.

## 2.2 Data Analysis

In this section, we aim to identify the factors that can potentially impact the optimization of the logistics network and share the key insights that we were able to discover during our exploration of the data. We begin by analyzing the three domains: customer requirements, logistics costs, and manufacturing capabilities individually to find potential bottlenecks. Afterward, we look at how these domains overlap to see where the most difficult trade-offs occur for our model.

### 2.2.1 Customer Requirements

In manufacturing and logistics problems, flexibility is of great importance when it comes to customer requirements. The problem does have flexibility on a surface level when it comes to the deadlines of delivering orders because we have been given time windows in which to deliver, but when we inspected it further, we found out that there are only about 50% of the orders that have delivery window size of one or more days, the other half of orders do not have flexibility and need to be delivered on a specified date. The following figure 2 breaks down the flexibility of orders in terms of their window sizes. Because 95% of orders must be delivered within a 0–2-day window, the model has very little 'temporal flexibility.' This means we cannot easily delay an order to wait for cheaper production at some other plant, making the optimization much more challenging.





Figure 2: Delivery window size for orders

Figure 3: Delivery window size for orders

Apart from this, a few minor insights we discovered about customer requirements are:

- 1) 67.4% of the orders are for service level DTP, 23.3% for DTD and 9.3% for CRF, indicating that we will incur transportation costs for nearly 90% of the orders. Figure 4 visualizes this below.
- 2) Even though the warehouses can produce 1,541 unique products, we receive orders for only 772 of them. Orders for 83 unique products make up 50% of the order, indicating a skewness towards certain types of microchips.
- 3) In total, we need to deliver nearly 30 million units which weigh around 183 kilograms.
- 4) All the orders were placed on the same day.



**Figure 4:** Service level breakdown

### 2.2.2 Logistics Costs

In terms of logistics costs, we couldn't uncover any critical insights like flexibility, the important points that we discovered are:

- 1) 98% of the shipping lanes use AIR as mode of transport.
- 2) We discovered that DTP is around 45% cheaper than DTD on average, however only 14% of the available shipping brackets are for DTP.
- 3) Delivery time of AIR is 2.21 days on average while it is 0.93 days for GROUND.
- 4) As we move from lighter to heavier weight brackets, the cost/kg to ship a product reduces.

In addition to this, we used the average cost to ship from each port for orders weighing 1.41kg, 4.44kg, and 13.33kg for DTD as well as DTP. We chose these weights because these represent first quartile, median, and the third quartile of the weight of the 9,215 orders that need to be shipped, hence giving us a detailed overview of the shipping costs. These numbers are summarized in Appendix B in Figure 5 and 6. Figure 5 summarizes DTP shipment costs and it can be seen that


since there are fewer brackets to ship, in most of the cases it costs the same to ship an order irrespective of its weight. Port 4 is the cheapest to ship from while Port 8 is the most expensive.

In Figure 6, we see a summary of DTD shipping costs. Since we have more shipping brackets, we see more fluctuation in prices for three different weights that we have used. Port 3 is the most expensive to ship from and Port 5 is potentially the cheapest one.

### 2.2.3 Manufacturing Capabilities

Moving on to manufacturing capabilities, we realized that we run into a problem with flexibility once again, illustrated in Figure 3 above. Even though we have 19 manufacturing plants to manufacture the products, when we looked at the data for orders, we found out that nearly 70% of the orders can only be fulfilled by one plant. The products that these orders required were only available at one plant. In only 0.01% of the cases, we had the chance to decide between 5 plants to serve an order. This also indicated a potential bottleneck that we could face when trying to optimize the problem.

Apart from this, we plotted the capacities of the plants along their costs to try to understand which of our plants might get overburdened. The following figure 7 shows this chart. We observed that the plants in the top left side of this chart which are plant 1 and plant 3 might get overburdened because they are producing cheaper and also have a higher capacity.



**Figure 7:** Plant daily capacity vs cost per

From the supply side of things these were all the insights that we uncovered. A critical overlap between two of these domains is that since 95% of orders have almost no delivery flexibility and 70% can only be made at one plant, the model is often 'trapped'. It cannot change the production location or the delivery date, indicating a possibility of getting an infeasible solution.

Another thing that we observed is that the cost of transporting an order is significantly lower than cost of producing an order e.g. the largest order in terms of units ordered will only cost 25.23\$ to transport but the production cost is a staggering 290K\$. We calculated these costs because we wanted to explore the possibility of bundling orders while shipping to benefit from lower rates, but this discovery highlighted that transportation costs reflect only a fraction of the total costs.

To summarize, this chapter has provided a foundation of manufacturer's capabilities and objectives. We started by identifying the key variables across all domains, ensuring all missing information was completed. Through the analysis, we uncovered some significant challenges which included the fact that 95% of the orders have almost no delivery flexibility and 70% are restricted to a single manufacturing site. We also identified that while transportation costs are much lower than production costs, the tight overlaps between these domains create a highly constrained environment. These insights confirm that a manual solution is not feasible, and they provide the necessary foundation for developing the optimization model in the next chapter.

## 3 Problem Setting

Our handled scenario covers the production and further transportation of microchips for a global microchip producer. Therefore, we take a look at an outbound supply chain network. In section 3.1, we provide a problem description. Further in section 3.2, we display the network in a graphical representation and explain its content in more detail. In section 3.3, we then introduce the mathematical notation for the problem. Lastly, we proceed with the formulation and explanation of the objective function and constraints of our mathematical model in 3.4.

### 3.1 Problem Description

For the supply chain network, we consider four different kinds of nodes: firstly, the plants, secondly, the origin ports, thirdly, the carriers and lastly the destination port. Plants serve as a product supplier and provide orders ready for shipment. Origin ports work as transshipment hubs. From there, different carriers can serve the orders. The final destination is the last port to which each order is shipped to and where we consider the delivery as done. The plants are physically linked via infrastructure to at least one origin port. From these ports, the orders are shipped one more time to the final destination. This last shipment is done by carriers, which are different shipping agencies. They can either ship the orders via air or via ground transportation by trucks. Each plant has a daily capacity level of orders it can fulfil in a day. Additionally, each plant offers just a certain variety of products. Each order has its own ID number and comes with a customer

ID to match the recipient, the date of the order, the wanted products with their quantities, their weight and the required service level. Either DTD, DTP or CRF. Two different types of costs occur. The first one are the warehouse costs, which occur each time a plant fulfils an order. If the wanted service level is CRF, only the warehouse costs occur due to the customer handling the freight. Therefore, no transportation costs incur for this particular order. The second type of costs are the freight costs of each order for being further transported from the origin port by a carrier. The carriers offer different freight rates based on weight brackets. We have to assign a path for each order, which consists of a routing from the plant through the origin ports and selecting the transportation carrier for the delivery to the final destination. We have to decide which of the plants we choose that have the demanded product in stock. If the plant has multiple origin ports as connections, we have to decide which origin port to assign. Lastly, we decide on the delivering carrier depending on different order weights and prices. Our developed model minimizes the total costs, which consist of the freight and warehousing costs.

### 3.2 Visualization Through A Network

In this section, we display the visualization of the entire outbound supply chain network in Figure 8 to provide a better understanding of the existing connections of the nodes. The network contains a total of 19 plants, 11 origin ports, 9 carriers and one last destination port. The entire network consists of three stages. In the first stage of the network, each order is assigned to a plant. In the second stage of the network, each plant is linked to one or multiple origin ports. For example, plant  $j_2$  can only serve origin port  $p_3$  and plant  $j_{12}$  can only serve origin port  $p_4$ , while plant  $j_{10}$  can serve origin ports  $p_1$  and  $p_2$ . In the third stage, the origin ports are being served by different carriers with alternating service types. Origin port  $p_2$  can be served by carriers  $k_1, k_2, k_3, k_4, k_5$  and  $k_8$ . Port  $p_7$  is the only port that is just being served by a single carrier,  $k_3$ . All other ports are at least being served by 2 different carriers. Port  $p_1$  is not linked to any carriers due to missing freight rates in the dataset. Therefore, no transportation mode is displayed. All other origin ports offer air mode transportation, but only origin ports  $p_3$  and  $p_9$  offer ground mode transportation additionally. Finally, all carriers deliver to one final destination  $p_9$ , where the order delivery ends. The destination port  $p_9$  happens to be an origin port  $p_9$  at the same time.

### 3.3 Mathematical Notation

Now we use the aforementioned information from the problem description to construct the mathematical model. Therefore, we start with the notation. We define the following sets, parameters and variables. At first, we take a look at the sets. We denote by  $I$  the set of all orders.

To display all plants, we use  $J$  as another set. The set of all carriers is displayed with a  $K$ .  $P$  displays the set of all origin ports. The set of the different service levels is denoted by  $S$  and lastly,  $D$  denotes the set of planning days. All sets are listed in Table 1 below.

Sets	Description
$I = \{1, \dots, 9215\}$	Orders
$J = \{1, \dots, 19\}$	Plants
$K = \{1, \dots, 9\}$	Carriers
$P = \{1, \dots, 10\}$	Origin ports
$D = \{1, \dots, 9\}$	Planning days
$S = \{DTD, DTP, CRF\}$	Service levels

**Table 1:** Formulation of sets

Secondly, we continue with our parameters. To display the warehouse costs for an order  $i$  at the plant  $j$ , we introduce  $WC_{ij}$ . The freight cost for an order  $i$ , which is being served by a carrier  $k$  on an origin port  $p$  on day  $d$ , is denoted by  $FC_{ikpd}$ .  $Cap_j$  indicates that each plant  $j$  has a daily capacity of orders. The parameter  $Prod_{ij}$  signals, if a product for an order  $i$  is available at plant  $j$ . If the product is available,  $Prod_{ij}$  has the value of 1 and if it is not available, the value is 0. We use the parameter  $Net_{jp}$  to display the plant and origin port connections. If plant  $j$  connects to the origin port  $p$ , the value is 1 and if no link exists, the value is set to 0. Another bivariate parameter we introduce is the  $Route_{kps}$ . It is set to 1, if the carrier  $k$  serves the route from the origin port  $p$  with the service level  $s$ . If the carrier  $k$  does not serve that route, the value is 0. The earliest allowed shipment date for an order  $i$  is denoted by  $ES_i$ . For the latest allowed shipment date for an order  $i$ , the parameter  $LS_i$  is used. The last parameter  $SL_i$  indicates the required service level for an order  $i$ . All parameters are listed in Table 2 below.

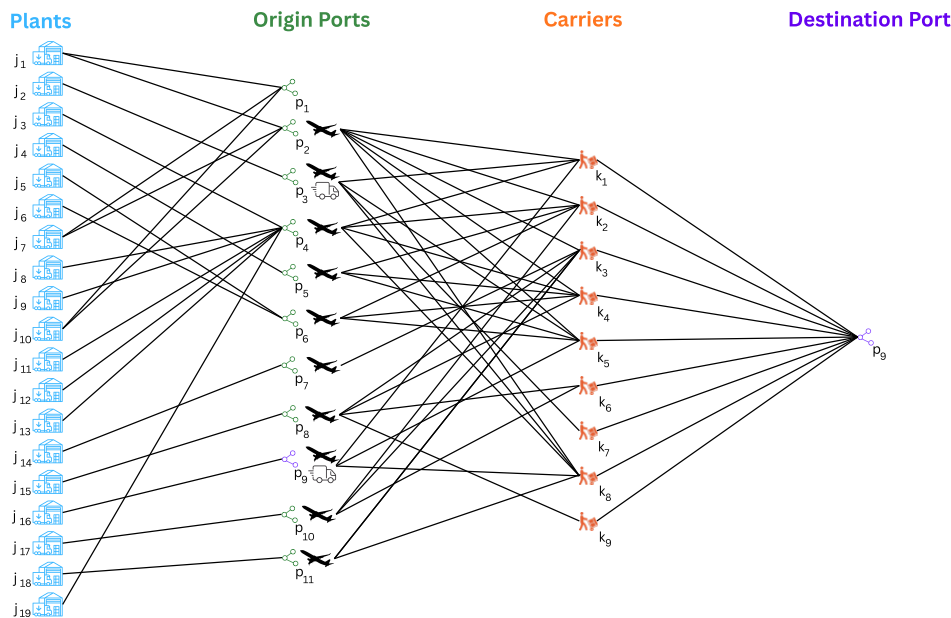
Parameters	Description
$WC_{ij}$	Warehouse cost for order $i$ at plant $j$
$FC_{ikpd}$	Freight cost for order $i$ via carrier $k$ and port $p$ on day $d$
$Cap_j$	Daily capacity of plant $j$
$Prod_{ij}$	$= \begin{cases} 1, & \text{if product for order } i \text{ is available at plant } j \\ 0, & \text{otherwise} \end{cases}$
$Net_{jp}$	$= \begin{cases} 1, & \text{if plant } j \text{ connects to port } p \\ 0, & \text{otherwise} \end{cases}$
$Route_{kps}$	$= \begin{cases} 1, & \text{if carrier } k \text{ serves rote from port } p \text{ with service level } s \\ 0, & \text{otherwise} \end{cases}$
$ES_i$	Earliest allowed shipping date for order $i$

$LS_i$	Latest allowed shipping date for order $i$
$SL_i$	Required service level for order $i$

**Table 2:** Formulation of parameters

Lastly, we introduce our bivariate decision variable. The variable  $x_{ijkpd}$  is set to 1, if an order  $i$  is assigned to the plant  $j$ , the carrier  $k$ , the origin port  $p$  on the planning day  $d$ . If the order is not assigned, the value of the variable is set to 0.

$$x_{ijkpd} \in \{0,1\}$$



**Figure 8:** Graphical representation of the outbound supply chain network

### 3.4 Mathematical Model

Now we conclude the model with the objective function and the constraints. Since we have a minimization problem, we want to minimize the total costs. These consist of the warehouse and freight costs. These appear if the decision variable with the corresponding indices has the value of 1. So, these costs are only added for that order  $i$ , plant  $j$ , carrier  $k$ , origin port  $p$  and planning day  $d$  combinations that are assigned. Therefore, we consider all orders, all plants, all carriers, all origin ports and lastly all planning days. The formulation of the objective function is displayed below in (1).

$$\text{Min } Z = \sum_i \sum_j \sum_k \sum_p \sum_d (WC_{ij} + FC_{ikpd}) \cdot x_{ijkpd} \quad (1)$$

The objective function is supplemented by the following constraints:

$$\sum_j \sum_k \sum_p \sum_d x_{ijkpd} = 1, \forall i \in I \quad (2)$$

The first constraint is the assignment constraint, where each order  $i$  can be assigned exactly once.

$$x_{ijkpd} = Prod_{ij}, \forall i, j, k, p, d \quad (3)$$

The second constraint ensures that an order  $i$  can only be assigned to a plant  $j$  if that plant has the required product in stock.

$$x_{ijkpd} \leq Net_{jp}, \forall i, j, k, p, d \quad (4)$$

(4) ensures that orders can only use existing physical connections between a plant  $j$  and an origin port  $p$ .

$$\sum_i \sum_k \sum_p x_{ijkpd} \leq Cap_j, \forall j \in J, d \in D \quad (5)$$

The next constraint limits the daily plant capacity, where the total orders assigned to plant  $j$  on day  $d$  cannot exceed its daily processing capacity.

$$x_{ijkpd} = 0, \text{ if } d < ES_i \text{ or } d > LS_i, \forall i, j, k, p, d \quad (6)$$

The fifth constraint forces order  $i$  to be shipped within its allowed date range.

$$x_{ijkpd} \leq Route_{kp,SL_i}, \forall i, j, k, p, d \quad (7)$$

(7) ensures order  $i$  can only use carrier  $k$  on origin port  $p$ , if that carrier offers the required service level on that route.

$$x_{ijkpd} \in \{0,1\}, \forall i, j, k, p, d \quad (8)$$

Lastly, the binary variable definition makes sure that all decisions are either 0, if not assigned or 1, if assigned.

Having defined the problem and completed the mathematical model, we now move to our results and benchmarking in the next section.

## 4. Results and Benchmarking<sup>1</sup>

The results from the model developed in Chapter 3 will be presented in this chapter. We will start by analyzing the solution we got from our soft capacity model, investigating the cost structure,

carrier assignments and plant utilization patterns. We then benchmark our solution against the baseline solution provided in the dataset to evaluate the model's performance and practical value.

## 4.1 Computational Setup

The dataset and model described in chapters 2 and 3 are the basis for all computational experiments presented in this chapter. We will compare our optimized solution (using soft capacity constraints with three-tier progressive penalties from Chapter 3's model) against the solution already provided in the dataset.

The model was solved using Gurobi 12.0.3 on a laptop with Intel i7 1165G7 chip, 16GB RAM with Linux Mint 22.1 as the OS). The model had approximately 54,000 binary decision variables and over 9,200 constraints, which tells us the full complexity of the problem. Our analysis focuses on 3 primary dimensions: (1) cost structure and economic implications, (2) capacity utilization and penalty distribution patterns and (3) operational improvements compared to the baseline solution. Details of these results are presented in Section 4.2.

## 4.2 Model Evolution and Infeasibility Resolution

### 4.2.1 Initial Hard Capacity Model and Infeasibility Discovery

When the model was initially deployed, it had traditional hard capacity constraints, requiring the total number of orders assigned to any plant on any day strictly not exceed that plant's processing capacity. This formulation is the standard approach in facility assignment optimization, where capacity limits are treated as inviolable.

However, when this model was run on the full dataset of 9,215 orders, Gurobi returned an infeasible status, indicating no solution was possible which satisfied all constraints simultaneously. Analysis of the Irreducible Inconsistent Subsystem (IIS) and some trial and error investigation revealed the root cause: on Day 3 (May 28, 2013) more than 1,700 orders were forced to ship due to their temporal window constraints and service level requirements. Unfortunately, PLANT03, which manufactures the products for most of those orders, has a daily capacity of 1,013 orders only, which created a deficit of 775 orders (76% over capacity).

To resolve this, we had to consider how we could model capacity violations in a manner that keeps the model realistic. Companies often (specially during peak season) employ overtime operations, weekend shifts and temporary workforce to meet demand commitments. Thus, we implemented

penalty multipliers (50%, 100% and 200%), at which point the modelling challenge became how to capture this operational reality while maintaining economic optimization.

#### 4.2.2 Soft Capacity Constraint Formulation

We reformulated the model using soft capacity constraints with explicitly economic penalties for exceeding capacity. This converts capacity limits from hard boundaries into economic trade-offs, allowing the optimizer to exceed capacity when the alternatives (such as refusing customer orders or using prohibitively expensive routing options) would be more costly

The soft capacity formulation introduces 3 new decision variables ( $y_1, y_2, y_3$ ) for each plant-day combination, representing the number of orders processed beyond nominal capacity in each penalty tier. Thus, the modified capacity constraint becomes:

$$\sum_{i,k,p} x_{ijkpd} \leq \text{Cap}_j + y_{1jd} + y_{2jd} + y_{3jd} \quad \forall j \in J, \forall d \in D \quad (9)$$

This allows the capacity to be exceeded through  $y$  variables, with each unit incurring penalty costs that create economic incentives to respect limits while maintaining feasibility

As such, the objective function is also modified to include penalty costs as a 3rd cost component:

$$\text{Minimize } Z = \sum(\text{warehouse costs}) + \sum(\text{freight costs}) + \sum(\text{penalty costs}) \quad (10)$$

Where the penalty cost component is formulated as:

$$\text{Penalty Costs} = \sum_{j \in J, d \in D} [y_1[j, d] \times \bar{q} \times c[j] \times 0.5 + y_2[j, d] \times \bar{q} \times c[j] \times 1.0 + y_3[j, d] \times \bar{q} \times c[j] \times 2.0] \quad (11)$$

Here,  $\bar{q}$  represents average order size and  $c[j]$  is the base warehouse cost per unit at plant  $j$ . The penalty variables ( $y_1, y_2, y_3$ ) are decision variables optimized simultaneously with the assignment variables ( $x$ ). When assignments exceed capacity at a plant-day combination, the optimizer will increase corresponding  $y$  values to satisfy the capacity constraint, which automatically increases the total cost. The optimizer balances this trade-off by either reassigning orders to alternative plants (if cheaper) or accepting the penalty cost (if the alternatives are more expensive). Each tier multiplies the base cost by its premium (50%, 100%, 200%). This ensures that the optimizer will only exceed capacity when the penalty costs are justified by savings in the warehouse or freight costs elsewhere in the solution.

The simultaneous modification of the constraint and objective creates a feasible model with economic realism, helping us quantify operational costs when demand exceeds capacity.

### 4.2.3 Three-Tier Progressive Penalty Structure

The penalty structure uses 3 progressive tiers designed to reflect the increasing marginal cost of capacity expansion through operational measures, with penalties calibrated as multiples of the base warehouse cost:

- Tier 1 (0 - 25% over capacity): 50% penalty premium (1.5x base warehouse cost)
- Tier 2 (25 - 50% over capacity): 100% penalty premium (2x base warehouse cost)
- Tier 3 (over 50% excess): 200% penalty premium (3x base warehouse cost)

We bound each tier to prevent arbitrary penalties. Tier 1 and 2 are each limited to 25% of nominal capacity ( $0.25 \times \text{Cap}[j]$ ), while Tier 3 is unbounded to ensure feasibility regardless of demand concentration.

## 4.3 Computational Performance

As described in section 4.1, the model has approximately 54,000 binary decision variables and over 9,200 constraints. Despite this scale, the solver achieved proven optimality in 0.33 seconds with a MIP gap of 0.00%, demonstrating that this optimization approach is practical for operational deployment.

The optimization successfully assigned all 9,215 orders (100% assignment rate) while satisfying all product availability, network connectivity, service level constraints and temporal window. The solution showed no constraint violations except the explicit capacity exceedance managed via the soft penalty mechanism.

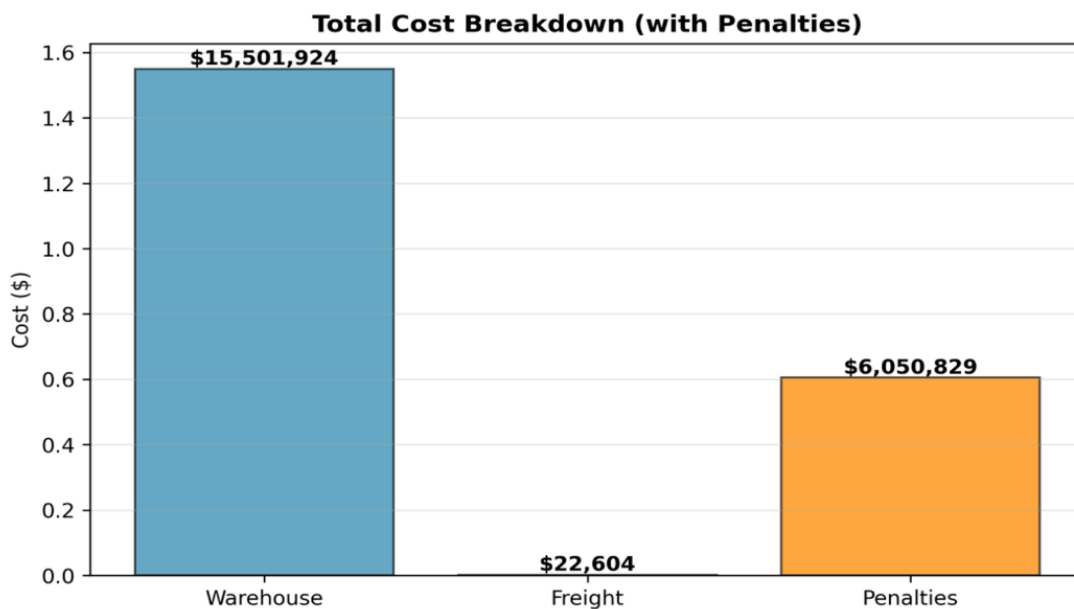
## 4.4 Optimized Solution Analysis

### 4.4.1 Overall Cost Structure

The total optimized cost came out to be \$21,575,357, revealing a 3-component cost structure that reflects the operational realities of this supply chain. Table 3 presents the comprehensive breakdown, with the relative scale of each component visualized in Figure 9:

Cost Component	Amount (\$)	Share (%)
Warehouse Costs	15,501,924	71.9
Freight Costs	22,604	0.1
Capacity Penalty Costs	6,050,829	28.0
<b>Total Cost</b>	<b>21,575,357</b>	<b>100.0</b>

**Table 3:** Optimized solution cost breakdown

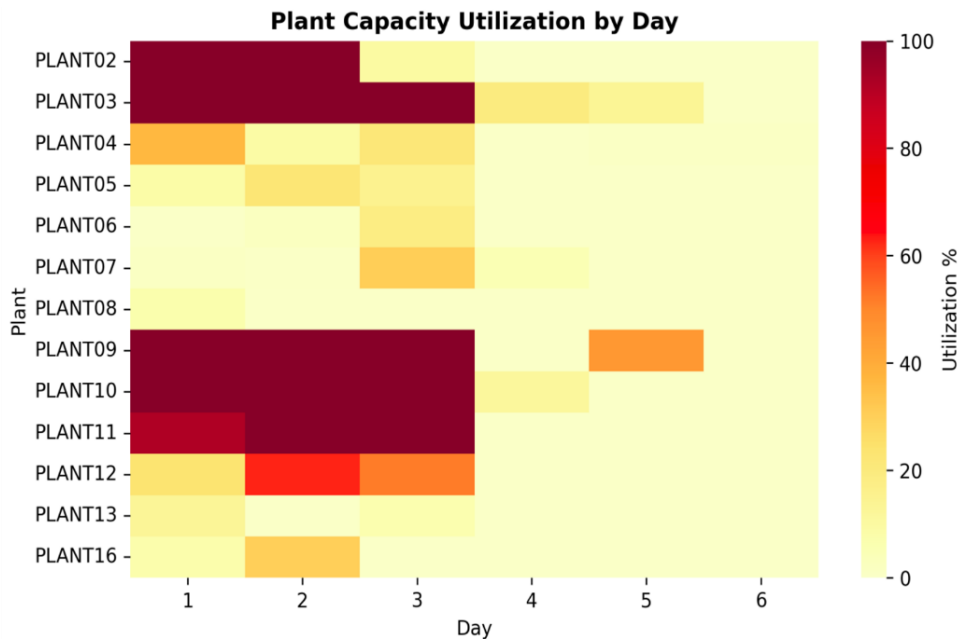


**Figure 9:** Total cost breakdown across three components

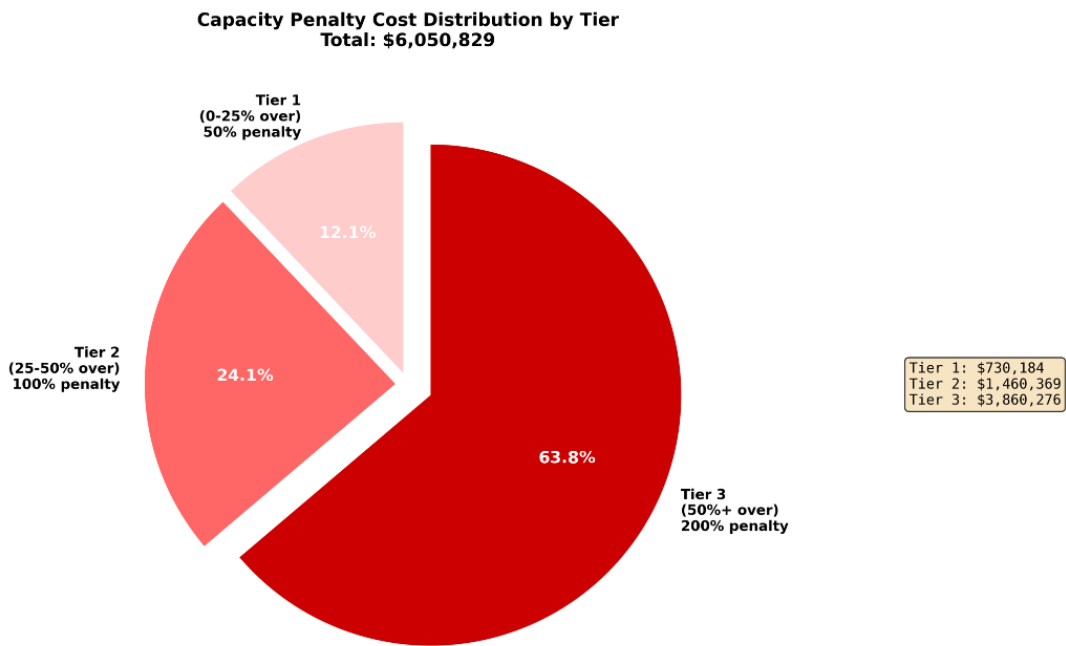
The significant penalty costs reflect overtime labor, temporary workforce, and weekend operations when plants exceed capacity. The model determined that \$6.05M in penalties was economically preferable to alternatives.

#### 4.4.2 Capacity Penalty Tier Distribution

Eight plant-day combinations exceeded nominal capacity during the planning horizon, as shown in the capacity utilization heatmap (Figure 10). PLANT03 accounts for three of these violations (Days 1, 2 and 3), with the deep red coloring indicating near or above 100% utilization, representing approximately \$4 million of the total penalty costs.



**Figure 10:** Plant capacity utilization heatmap by day with color intensity indicating



**Figure 11:** Distribution of capacity penalty costs across three progressive tiers

Figure 11 reveals that Tier 3 violations dominate with \$3.86M (63.8%), indicating a lot of capacity violations beyond the 50% threshold. Tier 2 contributes \$1.46M (24.1%), while tier 1 represents \$730K (12.1%).

This Tier 3 dominance tells us that the severe bottlenecks in the dataset did require substantial measures, which ultimately justifies the model's load distribution strategy.

#### 4.4.3 Temporal Distribution Patterns

Over 95% of the orders are concentrated on days 1-3, as visualized in Figure 12, reflecting customer service requirements which necessitated the changes to the model to make it feasible.




**Figure 12:** Temporal demand concentration pattern across the planning horizon

The area chart shows peaks of 3,168 and 3,123 orders on Days 1 and 3, which is followed by a decline to less than 300 orders per day after Day 3. This temporal concentration is the driving factor behind the capacity violations at multiple plants. Days 4-6 show substantially lower demand.

#### 4.4.4 Carrier Selection and Transportation Mode

Since the share of freight costs in the total cost is negligible (0.1%), carrier selection decisions have minimal impact on overall optimization. As shown in Figure 13, Carrier V444\_0 dominates with 5,146 orders (55.8%), more than 2.5 times the volume of the second-largest carrier (V444\_1 with 1,786 orders).



**Figure 13:** Order distribution across carriers

## 4.5 Benchmarking Against Baseline Solution

### 4.5.1 Baseline Solution Characteristics

The baseline solution provided in the dataset represents the company's historical assignment decisions for the same set of 9,215 orders. Analysis of this baseline reveals a total cost of \$15,695,280, distributed as \$15,633,222 (99.6%) in warehouse costs and \$62,058 (0.4%) in freight costs. Table 4 compares both solutions:

Metric	Baseline	Optimized
Total Cost (\$)	15,695,280	21,575,357
Warehouse Cost (\$)	15,633,222	15,501,924
Freight Cost (\$)	62,058	22,604
Penalty Cost (\$)	Not modeled	6,050,829
PLANT03 Orders	8,541 (92.7%)	~5,800 (63%)
Capacity Violations	Yes (implicit)	Yes (explicit)

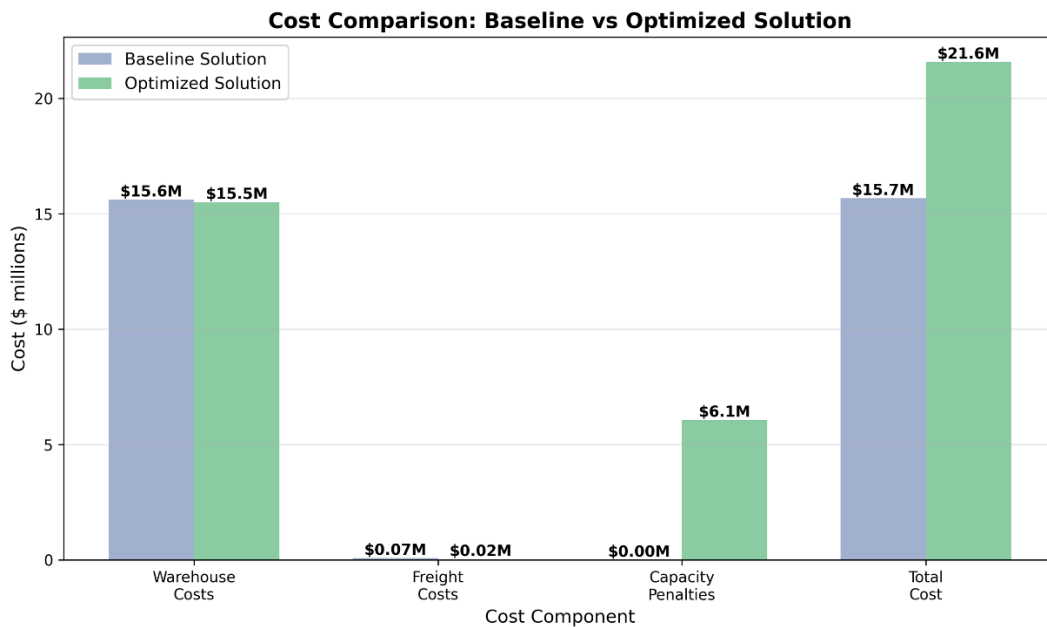
**Table 4:** Comparison of baseline and optimized solutions

Diving deeper in the baseline solution reveals some serious problems. It shows extreme concentration with over 92% of all orders assigned to PLANT03, resulting in severe capacity violations, given that PLANT03's daily capacity limit is 1,013 orders. With no indication of any

temporal distribution, we are forced to conclude that the baseline solution does not reflect feasible operations.

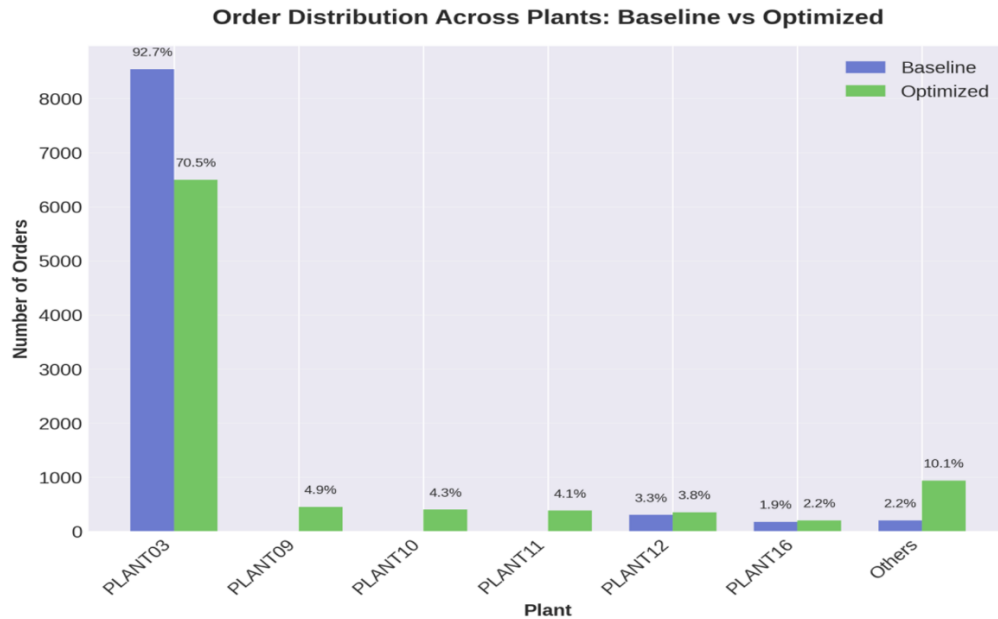
#### 4.5.2 Comparative Analysis and Cost Reconciliation

The apparent cost increase from baseline (\$15.7M) to our optimized solution (\$21.6M), requires nuanced interpretation. The baseline solution did not model or account for capacity penalties, treating all assignments as equally feasible regardless of capacity constraints. But we know that in reality that is not the case.



**Figure 14:** Cost structure comparison showing the baseline solution’s incomplete cost accounting

Figure 14 illustrates cost reconciliation. While optimized total cost (\$21.6M) appears higher than baseline, this does not indicate inferior performance. Warehouse costs decreased (\$15.6M → \$15.5M) and freight dropped (\$62K → \$23K). When comparing only the directly comparable cost components (warehouse and freight), our optimization achieves savings of \$170,752 (1.1%) through better facility and carrier selection.



**Figure 15:** Order distribution comparison across manufacturing plants, demonstrating load redistribution from extreme PLANT03 concentration

More significantly, the optimized solution distributes load more effectively across different plants. Figure 15 visualizes this redistribution, contrasting the extreme baseline against the optimized load spreading. While PLANT03 remains dominant due to product availability constraints, the optimized solution successfully engages 15 plants (vs 7 in baseline), with meaningful volume increases at PLANT09, PLANT10 and PLANT11. This redistribution reduces the severity of capacity violations on any single plant-day while maintaining all product availability and service level constraints.

#### 4.5.3 Operational Value and Strategic Implications

The optimization model demonstrates 2 primary contributions. First, it provides operational realism by incorporating capacity constraints and quantifying the economic costs of exceeding them, something the baseline approach ignored. Secondly, the model enables better resource planning by identifying which plant-day combinations will require overtime operations and quantifying their costs in advance.

Thus, the model's primary value lies not in cost reduction but in operational feasibility. This approach enables decision-makers to understand the true costs of current demand patterns and identify whether strategic interventions (such as capacity expansion or product line redistribution) would be economically justified.

## 4.6 Summary of Key Findings

The computational and benchmarking analysis reveal 3 critical insights about the problem:

**Infeasibility Resolution via Soft Constraints:** The initial hard capacity model proved infeasible due to Day 3's concentration of 1,788 orders at PLANT03, making strict adherence to the capacity limit operationally unrealistic. The soft capacity formulation with 3-tier progressive penalties transforms this infeasible problem into a solvable one while maintaining economic realism.

**Cost Structure Insights:** Analysis reveals that the warehouse costs (71.9%) and capacity penalties (28.0%) drive total expenditure, while freight costs are negligible. This suggests that interventions should focus on capacity management rather than transportation optimization.

**Benchmarking Interpretation:** The baseline solution's apparent cost advantage (\$15.7M vs \$21.6M) does not reflect superior performance. When comparing only directly comparable components (warehouse and freight), the optimization achieves \$171K in savings while providing superior load distribution (63% vs 93% concentration at PLANT03) and operational transparency regarding true fulfillment costs.

## 5. Conclusion and Future Outlook

This report looked at the facility and carrier assignment problem for a global microchip manufacturer managing over 9,200 orders across 19 plants, 9 carriers and 11 ports. Our primary contribution was demonstrating that soft capacity constraints with progressive tiers of penalties successfully transform this problem from being infeasible to solvable, while maintaining economic realism. While the initial hard capacity model proved infeasible due to concentrated demand on a particular plant-day combination, the 3 tier penalty framework achieved optimality in 0.33 seconds.

The optimized solution (\$21.6M total cost: 71.9% warehouse, 28.0% penalties, 0.1% freight) reveals that freight optimization offers minimal value while capacity management drives costs. Bench-marking against the baseline revealed extreme restriction in spatial flexibility, with 82.5% of the products manufactured at a single plant. Thus, our model's primary value lies not in cost reduction but in operational feasibility and strategic decision support for demand management. Future research should validate the penalty cost parameters empirically, extend to multi-period horizons and explore product line redistribution strategies to address manufacturing inflexibility.

## Bibliography

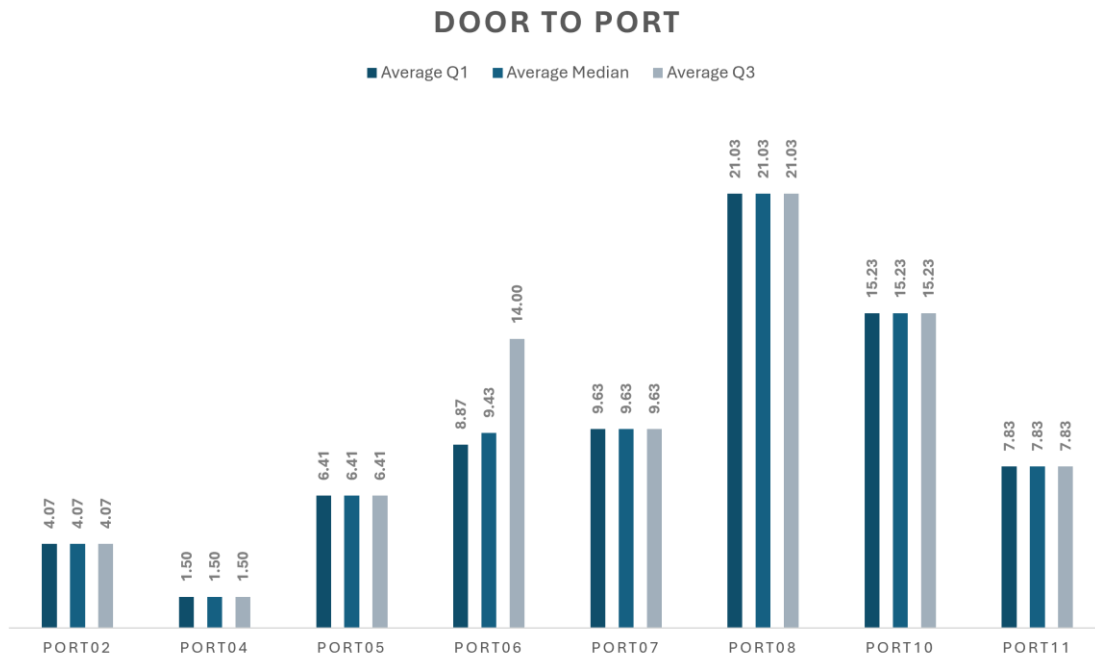
- Dzalbs, I., & Kalganova, T. (2020). Accelerating Supply Chains with Ant Colony Optimization Across a Range of Hardware Solutions. *Computers & Industrial Engineering*, 147, 106610.
- Fahimnia, B., Farahani, R. Z., Marian, R., & Luong, L. (2013). A Review and Critique on Integrated Production–Distribution Planning Models and Techniques. *Journal of Manufacturing Systems*, 32(1), 1–19.
- Mönch, L., Uzsoy, R., & Fowler, J. W. (2018). A Survey of Semiconductor Supply Chain Models Part III: Master Planning, Production Planning, and Demand Fulfilment. *International Journal of Production Research*, 56(13), 4565–4584.
- Ullrich, C. A. (2013). Integrated Machine Scheduling and Vehicle Routing with Time Windows. *European Journal of Operational Research*, 227(1), 152–165.
- Xiong, W., Wu, D. D., & Yeung, J. H. Y. (2025). Semiconductor Supply Chain Resilience and Disruption: Insights, Mitigation, and Future Directions. *International Journal of Production Research*, 63(9), 3442–3465.

## Appendix A: AI Usage Declaration

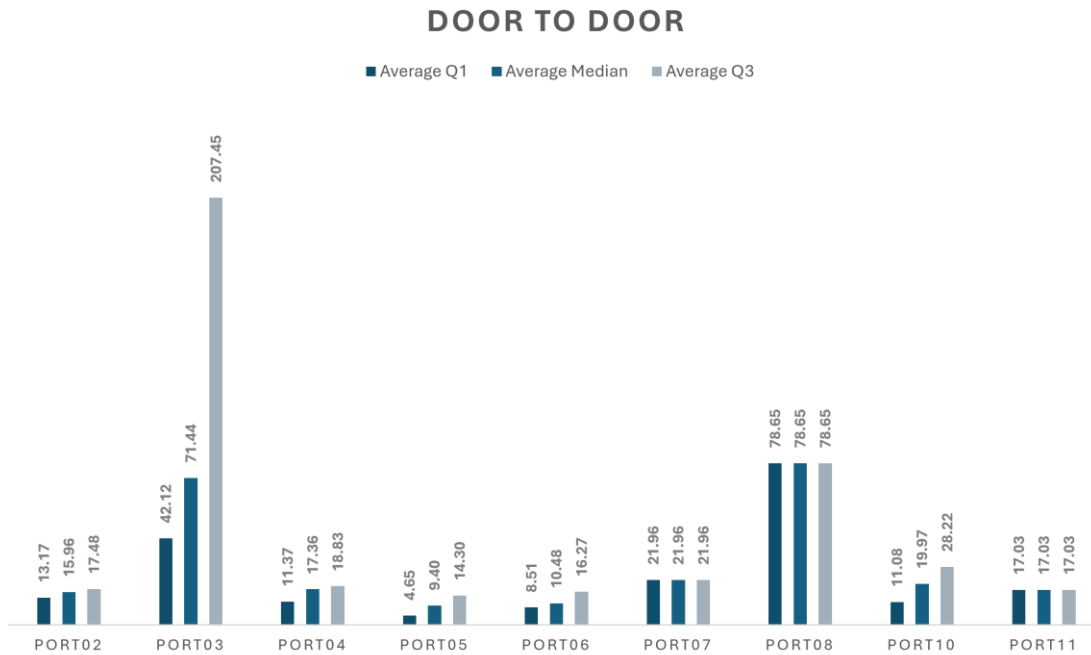
Name of AI Tool	Date of retrieval	URL	Prompt	Classification
Claude (Anthropic), Model: Claude Sonnet 4.5	January 10-22, 2026	claude.ai	Multiple iterative prompts including, but not limited to: "Create draft of Results and Benchmarking", "Enhance based on reference papers", "Add computational setup", "Create visualizations", "Reduce redundancy", "Address supervisor feedback"	INTERACTIVE

**Note on Methodology:** The above declaration represents a simplified summary of a more complex AI collaboration process. Throughout the semester (October 2025 - January 2026), we conducted approximately 30 separate AI sessions for various project tasks (data analysis, model development, implementation, coding, debugging). These were meticulously documented, along with all other relevant project files, creating a RAG-like system. Chapter 5 was then developed through extensive interaction where we provided strategic direction, reviewed drafts, verified technical accuracy, and incorporated supervisor feedback. Every technical decision, code implementation, and analytical insight in Chapter 5 originated from our semester-long work; the AI assisted in organizing and articulating these results coherently. We are prepared to demonstrate this documentation system and answer detailed questions about any aspect of the work to verify our understanding and ownership.

# 1Appendix B



**Figure 5:** Average shipping Cost DTP



**Figure 6:** Average shipping Cost DTD

## Distribution of Work

<b>Section</b>	<b>Name</b>
Introduction	[REDACTED]
Data Exploration and Analysis	[REDACTED]
Problem Setting	[REDACTED]
Results and Benchmarking	Hamzauddin Siddiqui
Conclusion and Future Outlook	Hamzauddin Siddiqui

## Declaration of Authenticity

I declare, Hamzauddin Siddiqui

- that I completed the present Scientific Project independently and without any unauthorized assistance
- that I did not use any aids other than those explicitly approved in advance and specified by me
- that that I have identified parts of the work whose wording or meaning has been taken from other documents (this also includes internet sources and AI-based tools), stating the source and
- that I have not yet submitted this work in the same or similar form to any other examination authority.

I agree,

- that my work can be checked for plagiarism using software. I am aware that plagiarism constitutes serious academic misconduct and that deceit will be punished in accordance with the study and examination regulations that apply to me.

I am aware,

- that in case of deceit the work will be officially declared “insufficient (5.0)”
- that I bear full responsibility for incorrect or distorted content generated by AI, incorrect references, violations of data protection and copyright law or plagiarism.

My intellectual performance was consistently at the forefront and I always managed the process myself.

..... [REDACTED], 23/01/2026)

(Signature) / (Place, Time)

## Declaration of Authenticity

I declare, [REDACTED]

- that I completed the present Scientific Project independently and without any unauthorized assistance
- that I did not use any aids other than those explicitly approved in advance and specified by me
- that that I have identified parts of the work whose wording or meaning has been taken from other documents (this also includes internet sources and AI-based tools), stating the source and
- that I have not yet submitted this work in the same or similar form to any other examination authority.

I agree,

- that my work can be checked for plagiarism using software. I am aware that plagiarism constitutes serious academic misconduct and that deceit will be punished in accordance with the study and examination regulations that apply to me.

I am aware,

- that in case of deceit the work will be officially declared “insufficient (5.0)”
- that I bear full responsibility for incorrect or distorted content generated by AI, incorrect references, violations of data protection and copyright law or plagiarism.

My intellectual performance was consistently at the forefront and I always managed the process myself.

[REDACTED], 23/01/2026)

(Signature) / (Place, Time)

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]