

Optimising the Frozen Flow

Buffer Design for the Dispatch Area in Automated Cold Storage Warehouses

SEN2331: Master Thesis
Merle van Beek

Optimising the Frozen Flow

Buffer Design for the Dispatch Area in Automated Cold Storage Warehouses

by

Merle van Beek

to obtain the degree of Master of Science in
Complex Systems Engineering & Management
at Delft University of Technology,
to be defended publicly on Thursday September 24

Student number: 5314461
Thesis committee: dr. S. Fazi - TU Delft, first supervisor
Prof. dr. M.E. Warnier - TU Delft, chair and second supervisor
P. de Breet - NewCold, external supervisor

Cover: Highbay with operating crane at a NewCold warehouse in
Wakefield, England (Modified) - NewCold Cooperatief U.A.

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

In 2020 I started my bachelor in systems engineering, policy analysis and management at the faculty of Technology, Policy and Management. Little did I know that this would also be the start of a growing interest in complex socio-technical systems and the eagerness to learn about these systems and to improve them. Before you lies my master's thesis, a result of this interest which marks the end of my studies in Delft.

In my bachelor's, I chose the specialisation in transport and logistics, which I continued in my master's. In addition to this specialisation, I also developed a growing interest in applying this knowledge on a less abstract level and in practice. I therefore decided to do a minor abroad in business administration. I gained more understanding of businesses, to even better understand complex challenges that exist in practice. Therefore, I started in 2023 with two master's; a master in Business Administration at Erasmus University and the one behind this thesis, Complex Systems Engineering and Management at Delft University of Technology.

All of these interests came together in my graduation internship at NewCold where innovation, automation and logistics result in complex systems, where technology, society, and management cannot be studied independent of each other resulting in a thesis combining these lenses.

I would like to thank my supervisors from the TU Delft for their support and contribution in shaping my thesis into research that is relevant both academically and for the industry. With confirmation, different perspectives, and honest feedback, they helped me put this project within the bigger picture, a picture that sometimes kept getting smaller for me.

I'm also very grateful for the opportunity and support I received at NewCold during this project. I have learned a lot about the company and the innovative, but down-to-earth way of working in a field that is developing fast. Seeing a warehouse in operation was a very nice kick-off of my internship and immediately put my assignment in perspective. A big thank you to all my colleagues for helping me collect all the data I needed, answer all my questions, and treat me like a team member.

To my family and friends, a big thank you for the encouragement, conversations, and optimism. Your support helped me through the process and helped me finish this chapter happy and proud.

Not only does my research zoom in on the outbound staging area of automated cold-storage warehouses, it also zooms out and puts the challenge and the proposed alternative in a bigger context. I hope this thesis contributes to NewCold's next projects, and the whole industry. May you enjoy reading this thesis as much as I enjoyed this research.

*Merle van Beek
Delft, September 2026*

Confidentiality Statement

This thesis uses, discusses and analyses internal documentation and operational data of NewCold. Two versions of this thesis exist. This is the public version deposited in the TU Delft Repository and is redacted relative to the version held by the thesis committee. This public version is approved solely for publication in the TU Delft Repository. Certain operational figures have been scaled, and sensitive information relating to the current and proposed Piacenza sites has been generalised or omitted. The relative relationships between figures have been preserved, ensuring that the analysis, findings and recommendations remain valid. Raw operational data are not reproduced; only derived and calculated data are included in this report.

Use of AI Tools

During this thesis trajectory, AI-based tools are used for writing and structuring information. Claude (Anthropic) assisted with structuring the chapters, content and modelling results, and readability of the text next to brainstorming and consistency checks. LaTeX formatting issues and translations were discussed, and supported with troubleshooting during modelling. DeepL was used for occasionally translating English to Dutch and Dutch to English.

The author is responsible for the content of this thesis. The research design, data analysis, interpretations, modelling and simulation belong to the author. No confidential or company-internal information of NewCold was entered into these tools. Their use was limited to structuring, language and general modelling questions.

Executive Summary

Automated cold storage warehouses increasingly rely on automation to manage growing throughput demands, labour complexities and sustainability restrictions in the constantly evolving cold-storage logistics sector. Yet the design of outbound buffer designs remains under-explored compared to long-term storage and retrieval processes. In automated cold-storage warehouses, fixed buffer lanes stage pallets hours ahead of truck pickup, creating a structural mismatch with the stochastic nature of truck arrivals. When trucks are delayed, buffer lanes fill up, disrupting outbound flow and forcing operational exception handling. An automated storage and retrieval system (AS/RS) could be a potential solution to this problem, in particular shuttle based storage and retrieval systems (SBS/RS) are identified as the most suitable technology given their fast-retrieval, small-footprint, and flexible release characteristics, well-suited to short-term buffering. This research investigates whether a shuttle-based storage and retrieval system (SBS/RS) can serve as an alternative to traditional buffer lanes. There is a limited empirical and case-study basis for evaluating such an alternative. This leads to the research question: *How can a shuttle-based dispatch buffer be designed and controlled to minimize crane congestion and outbound delays, improving outbound performance in an automated, cold-storage warehouse?*

The research combines a literature review and expert consultation with a thorough system analysis and discrete-event simulation model, executed in Simio. NewCold's warehouse in Piacenza, Italy is used as a case study with an existing warehouse and a to-be-built expansion. Three scenarios were compared: a validated baseline reflecting the current Piacenza layout, a proposed layout with a shuttle based storage buffer where picked pallets return to the highbay before staging, and a variant where picked pallets route directly to the shuttle buffer. The simulation model was validated against real operational data, closely reproducing observed warehouse behaviour.

The results show that a shuttle-based dispatch buffer's main value lies in decoupling outbound operations from truck-arrival uncertainty and in creating room to operate with fewer cranes. This benefit does not come without cost: a smaller crane fleet operates closer to its capacity limit. This confirms an under-resourced buffer was shown to destabilise the system rather than absorb demand. This confirms that such a design must be dimensioned deliberately rather than assumed to work at any scale. For NewCold, this research provides a validated basis for the Piacenza 2 design decision, together with a reusable simulation model that can support further design evaluation beyond the scenarios tested here. This research also offers a way to evaluate a not-yet-built facility's design choices before construction, using a model validated against a comparable existing operation, and highlights the value of a full-system analysis when investigating area-bounded design adjustments. Academically, this research addresses a gap in the SBS/RS literature by applying shuttle-based storage and retrieval methodology to an outbound dispatch-buffering role. Validated against real operational data, the research identifies external crane capacity as a coordination dependency not previously characterised in studies that treat the shuttle system as self-contained.

Contents

Preface	i
Confidentiality Statement	ii
Use of AI Tools	iii
Executive Summary	iv
Abbreviations	x
1 Introduction	1
1.1 Problem introduction	1
1.2 Problem statement	1
1.3 Research question and deliverable	3
1.4 Research structure	4
2 Literature Review	5
2.1 Core concepts	5
2.1.1 Warehousing in modern supply chains	5
2.1.2 Warehouse activities and design	6
2.1.3 Storage systems	7
2.1.4 KPIs for warehouses	9
2.2 KPI selection	9
2.3 AS/RS	12
2.4 Warehouse layout design	14
2.5 Knowledge gap	16
3 Methodology	18
3.1 Research design	18
3.1.1 Research flow diagram	21
3.2 Case study	21
4 System Analysis	22
4.1 Warehouse processes and flows	22
4.2 Warehouse characteristics	24
4.2.1 Piacenza 1	24
4.2.2 Piacenza 2	25
4.3 Uncertainties and operational constraints	25
4.4 Scenarios	26
4.4.1 Scenario 1	26
4.4.2 Scenario 2	26
4.4.3 Scenario 3	29
5 Model	32
5.1 Model logic and processes	32
5.2 Model setup	33
5.2.1 Scenario 1	34
5.2.2 Scenario 2	34
5.2.3 Scenario 3	36
5.3 Model input	37
5.3.1 Piacenza 1	37
5.3.2 Piacenza 2	38
5.4 KPI's	40

5.5	Assumptions and simplifications	41
5.6	Verification and validation	43
6	Scenario design and Simulation	44
6.1	Scenario 1	44
6.1.1	Validation	44
6.1.2	Results	45
6.2	Scenario 2	46
6.2.1	Model configuration	47
6.2.2	Results	50
6.3	Scenario 3	54
6.3.1	Model configuration	54
6.3.2	Results	57
6.4	Scenario comparison	60
6.4.1	Expected outcomes	61
6.5	Sensitivity analysis	62
6.5.1	Results	62
6.6	Economic analysis	64
7	Discussion	66
7.1	Key findings	66
7.2	Implications	67
7.2.1	Academic contribution	67
7.2.2	Practical implications	68
7.3	Limitations	69
7.4	Future research	70
8	Conclusion	72
	References	75
A	IDEF0 Model	81
A.1	Scenario 1	81
A.2	Scenario 2	86
A.3	Scenario 3	88
B	Model structure	90
C	Model parameters	91
D	Model validation	92
E	Additional Results	97
E.1	Scenario 1	97
E.2	Scenario 2	99
E.2.1	Configuration 1	102
E.2.2	Configuration 2	102
E.2.3	Configuration 3	103
E.2.4	Configuration 4	105
E.2.5	Configuration 5	106
E.2.6	Configuration 6	107
E.2.7	Configuration 7	108
E.3	Scenario 3	108
E.3.1	Configuration 1	110
E.3.2	Configuration 2	111
E.3.3	Configuration 3	111
E.3.4	Configuration 4	112

List of Figures

1.1 Warehouse flow diagram	2
3.1 Research Flow Diagram	21
4.1 Flowchart general scenario 1	23
4.2 Piacenza 1 layout	25
4.3 Flowchart scenario 2	27
4.4 Flowchart scenario 3	30
5.1 Warmup stabilisation highbay scenario 1	34
5.2 Warmup stabilisation highbay scenario 2	35
5.3 Warmup stabilisation buffer storage scenario 2	35
5.4 Highbay occupancy stabilisation, scenario 3	36
5.5 Buffer storage occupancy stabilisation, scenario 3	36
6.1 Pallets in highbay over a representative day, scenario 1	46
6.2 Inbound trucks in system over a representative day, scenario 2	52
6.3 Outbound trucks in system over a representative day, scenario 2	52
6.4 Crane retrievals per hour over a representative day, scenario 2	53
6.5 Pallets in buffer storage over a representative day, scenario 2	53
6.6 Pallets in highbay over a representative day, scenario 2	54
6.7 Inbound trucks in system over a representative day, scenario 3	58
6.8 Outbound trucks in system over a representative day, scenario 3	59
6.9 Crane retrievals per hour over a representative day, scenario 3	59
6.10 Pallets in buffer storage over a representative day, scenario 3	60
6.11 Pallets in highbay over a representative day, scenario 3	60
A.1 A-0 diagram for an automated, cold-storage warehouse.	81
A.2 Control systems of MHE	83
A.3 A-1 diagram for an automated, cold-storage warehouse.	83
A.4 A-2 diagram of processing inbound pallets in an automated, cold-storage warehouse.	84
A.5 A-2 diagram for sending out orders in an automated, cold-storage warehouse.	85
A.6 A-3 diagram for sending out orders in an automated, cold-storage warehouse.	86
A.7 A-2 diagram for sending out orders in an automated, cold-storage warehouse with shuttle based buffer storage.	87
A.8 A-3 diagram for staging pallet in outbound buffer storage in an automated, cold-storage warehouse.	88
A.9 A-2 diagram for sending out orders in outbound buffer storage in an automated, cold-storage warehouse.	89
E.1 Inbound trucks in system over a representative day, scenario 1	97
E.2 Outbound trucks in system over a representative day, scenario 1	98
E.3 Crane retrievals per hour over a representative day, scenario 1	98
E.4 Inbound registration throughput over a representative day, scenario 1	99
E.5 Buffer storage occupancy stabilisation, scenario 2, configuration 1	102
E.6 Buffer storage occupancy stabilization, scenario 2, configuration 2	103
E.7 Buffer storage occupancy stabilisation, scenario 2, configuration 3	104
E.8 Buffer storage occupancy stabilisation, scenario 2, configuration 4	105
E.9 Buffer storage occupancy, configuration 5, scenario 2	106
E.10 Buffer storage occupancy, configuration 6, scenario 2	107

E.11 Buffer storage occupancy, configuration 7, scenario 2	108
E.12 Buffer storage occupancy stabilisation, configuration 1, scenario 3	110
E.13 Buffer storage occupancy stabilisation, configuration 2, scenario 3	111
E.14 Buffer storage occupancy stabilisation, configuration 3, scenario 3	111
E.15 Buffer storage occupancy stabilisation, configuration 4, scenario 3)	112

List of Tables

2.1	AS/RS systems	8
2.2	KPI literature overview	11
2.3	AS/RS literature overview	13
2.4	Warehouse layout and buffer literature overview	15
3.1	Research sub-questions and corresponding methods	19
4.1	Expected outcomes for scenario 2	28
4.2	Expected outcomes for scenario 3	31
5.1	Model inputs Piacenza 1	37
5.2	Model inputs that differ for Piacenza 2	40
5.3	Key KPIs and their applicability by scenario	41
5.4	Assumptions and simplifications	42
6.1	Scenario 1 validation	44
6.2	Scenario 1 KPI summary	45
6.3	Scenario 2 capacity dimensioning KPIs	48
6.4	Scenario 2 capacity dimensioning KPIs (continued)	48
6.5	Scenario 2 KPI summary final configuration	51
6.6	Scenario 3 capacity dimensioning KPIs	55
6.7	Scenario 3 KPI summary	57
6.8	Scenario comparison, final configurations	61
6.9	Scenario 2 expected outcomes	61
6.10	Scenario 3 expected outcomes	62
6.11	Sensitivity analysis scenario 3	63
6.12	Equipment and cost comparison across design configurations	64
D.1	Model validation variables	96
E.1	Scenario 2, full KPI detail by experiment	100
E.2	Scenario 2, full KPI detail by experiment (continued)	101
E.3	Raw per replication KPI values, scenario 2, configuration 1	102
E.4	Raw per replication KPI values, scenario 2, configuration 2	103
E.5	Raw per replication KPI values, scenario 2, configuration 3	104
E.6	Raw per replication KPI values, scenario 2, configuration 4	105
E.7	Raw per replication KPI values, scenario 2 configuration 5	106
E.8	Raw per replication KPI values, scenario 2 configuration 6	107
E.9	Raw per replication KPI values, scenario 2 configuration 7	108
E.10	Scenario 3, full KPI detail by experiment	109
E.11	Raw per replication KPI values, configuration 1, scenario 3	110
E.12	Raw per replication KPI values, configuration 2, scenario 3	111
E.13	Raw per replication KPI values, configuration 3, scenario 3	112
E.14	Raw per replication KPI values, scenario 3 configuration 4	112

Abbreviations

Abbreviation	Definition
AS/RS	Automated Storage and Retrieval System
ATL	Automatic Truck Loading
ATU	Automatic Truck Unloading
AVS/RS	Automated Vehicle Storage and Retrieval System
BRC	British Retail Consortium
DES	Discrete Event Simulation
MFCS	Material Flow Control System
OSA	Operating Services Agreement
SBS/RS	Shuttle Based Storage and Retrieval System
SKU	Stock Keeping Unit
T-Car	Transfer Car
WMS	Warehouse Management System

Introduction

1.1. Problem introduction

Warehousing plays a central role in modern logistics and supply chain management. Especially in a world where sustainability and innovation are becoming increasingly important (MHI & Deloitte, 2023). Warehouses support production and distribution by temporarily storing goods and ensuring continuity in supply chains. Consequently, the design and operation of warehouse systems have become relevant factors in achieving high warehouse performance and efficiency. Warehouse systems consist of multiple interconnected processes and areas that collectively determine the overall performance of the system. The processes include receiving goods (inbound area), storing products (highbay area), order picking (pickfloor), and shipping (outbound area). The coordination of these activities strongly influences warehouse throughput, operational costs, and service levels (Rouwenhorst et al., 2000). Warehouses differ in their level of automatisisation and innovativeness, depending on product characteristics, customer requirements and regulations. Throughput is defined by the capacity of systems and actual operational activities (Rouwenhorst et al., 2000).

Cold-chain logistics is a type of warehousing that prevents food spoilage, a global problem, where 144 million tons of food loss could be prevented (Han et al., 2021). The market of cold-chain logistics is growing as society demands more and more frozen products (Ritter et al., 2023). Cold-storage facilities have high levels of automation, as strict product safety regulations, controlled temperatures and harsh working environments make manual work less feasible and costly (Datex, n.d.). Automation technologies such as automated storage and retrieval systems (AS/RS) are widely used in cold-storage warehouses (Meneghetti & Monti, 2015). Automation improves storage density, throughput, and reliability.

From the different processes in a warehouse, outbound processes play a crucial role in coordinating warehouse operations with transportation activities, particularly when goods must be adjusted for outbound shipments (Rouwenhorst et al., 2000). Outbound processes involve consolidating orders, staging pallets and preparing shipments for loading on trucks (J. Gu et al., 2007). Efficient coordination of these activities is essential for maintaining smooth material flows and avoiding congestion within the warehouse. Bottlenecks regarding truck arrival uncertainty in the outbound processes, arise in resource and equipment utilisation in the whole warehouse flow and throughput. These direct and indirect consequences affect not only the warehouse, but also its customers, and equipment suppliers.

1.2. Problem statement

Warehouse efficiency is the ability of a warehouse to optimize its operations and resources to achieve the desired level of performance in terms of speed, accuracy, and cost-effectiveness (Ballou, 2007). Commonly used metrics for warehouse performance are travel or cycle time, waiting time, system response time, throughput, investment cost, relocation duration and energy consumption (Q. Liu et al., 2024; Marchet et al., 2013; Xu et al., 2025; Zou et al., 2016). Cold-storage requirements related to quality and safety, lead to additional challenges for warehouse design like strict temperature require-

ments, specialized equipment for frozen environments, and health and safety limitations for employees. Designing an efficient warehouse layout becomes more complex when automated technologies and specific operational constraints must be considered simultaneously.

The following flow diagram (figure 1.1) of an automated cold-storage warehouse shows the general processes and flows of a warehouse.




Figure 1.1: Warehouse flow diagram

Four main areas within an automated cold-storage warehouse are distinguished: the inbound-, storage-, picking- and outbound area. The storage area is located in the highbay, whereas the inbound, outbound, and picking area are all situated in the lowbay area. The highbay area has a temperature of -25°C and the lowbay area has a temperature of -18°C . Between these two temperature-controlled parts of the warehouse, airlocks are in place to limit temperature and energy loss by opening gates one-by-one. Airlocks are also in place at the inbound and outbound docks. The highbay is fully automated, and the lowbay also involves manual labour.

In automated cold-storage warehouses, trucks arrive for loading and unloading at corresponding docks. For a truck delivering goods, the frozen goods on pallets are transported inside the warehouse with forklifts or electric pallet trucks (EPTs) to the I-point where pallets are registered and checked for irregularities. On conveyors and a monorail, approved pallets are moved to the highbay of the warehouse where goods are stored densely with an AS/RS system with cranes. An uneven distribution of crane movements in the highbay storage area can result in differences in equipment utilisation, maintenance requirements and routing inefficiencies. The distribution of all activities in the warehouse, including crane activities, is done by the Warehouse Management System (WMS) which controls, coordinates, and optimizes operational activities in a warehouse (Zhang et al., 2025). Once an order for a specific pallet pickup comes in, the WMS assigns the bufferlane the pallet will move to. The pallet is transported by conveyors and the monorail which are assigned by the MFCS (Material Flow Control System). Sometimes, an order requires only part of a pallet's contents, rather than one or more full pallets. Then, the pallet is moved from the highbay to the pickfloor where the required quantity is picked manually. After the order picking, a pallet is moved back to the highbay storage.

Internal and external uncertainties like WMS decisions and truck arrival times, are continuing to influence daily operations, system performance and customer satisfaction. In practise, these uncertainties can never be completely eliminated. For the outbound flow arrival times are a crucial factor for performance as customers depend on these orders, inbound movements are not as time constraining as they

are not yet to be retrieved from the warehouse. However, inbound and outbound trips are often linked which makes an inbound arrival also an outbound arrival for a different load. Moreover, although warehouses operate 24 hours a day, truck arrivals typically cluster around predictable peak hours. These peaks largely determine the required capacity of a warehouse, including outbound buffer infrastructure.

Buffer lanes temporarily store pallets ahead of truck pickup, but when trucks are delayed, pallets remain staged in those bufferlanes longer than intended, disrupting the schedule for following orders and forcing operational staff into real-time exception handling. These disruptions can cascade into missed truck departure windows, additional wait time, and delayed deliveries to customers. All directly affecting service reliability and quality, and operational costs. This mismatch between fixed lane-based staging and stochastic truck arrivals is especially pressing during peak periods, yet outbound buffer design and control remain underexplored. Existing research focuses predominantly on storage and retrieval processes, rarely studying buffer design, truck arrivals, and crane activities as an integrated system interacting with other warehouse activities and stakeholders, particularly in highly automated cold-storage warehouses. Empirical case studies using real-life data are also scarce. As warehouse automation continues to evolve, understanding how automated storage systems can support outbound buffering and staging becomes increasingly relevant. In practice, these dynamics are operationally significant, outbound delays increase pressure on cranes that are a shared resource for storage activities as well as retrieval activities. Unreliable timing risks service-level commitments to customers of the warehouse.

1.3. Research question and deliverable

This research investigates whether an automated storage and retrieval system (AS/RS) can serve as an alternative or supplement to traditional buffer lanes in the dispatch area of an automated cold-storage warehouse. AS/RS is often applied in automated warehouses for long-term storage, in the highbay. A shuttle based AS/RS system is most suitable for fast retrieval and a relatively small amount of storage spaces for short-term storage in the limited space available in the lowbay (Kriehn et al., 2018; Lerher, 2015; Wang et al., 2015), suiting a short-term storage purpose in a warehouse with pallets as stock-keeping-units (SKUs). Rather than staging pallets in fixed buffer lanes hours in advance, a shuttle buffer storage stores retrieves pallets flexibly and releases them only when the corresponding truck arrives. This shuttle will move horizontally once a task is assigned and is moved vertically by lifts, a so called 2D shuttle. The buffer storage stores pallets, depending on capacity, up to 24 hours in advance. This alternative dispatch solution has the potential to reduce crane congestion and improve outbound reliability, benefiting the warehouse and its customers.

The performance of this shuttle-based dispatch buffer is evaluated through discrete event simulation using real operational data from NewCold's warehouse in Piacenza, Italy. This warehouse is used as case-study with an existing warehouse and a to-be-built expansion. Three scenarios are developed: the current warehouse layout as a baseline, the proposed layout for a new warehouse in Piacenza with a buffer storage where picked pallets return to the highbay before staging, and an extended storage where picked pallets route directly to the shuttle buffer. By comparing these scenarios by key performance indicators (KPIs), the research identifies if a shuttle based AS/RS used as buffer storage in the outbound area, improves the outbound performance under uncertain truck arrival patterns, and which configuration best supports the performance. Next to the simulation results, a full picture of technical, and economical costs and benefits is given.

This leads to the following research question:

How can a shuttle based dispatch buffer be designed and controlled to minimize crane congestion and outbound delays, improving outbound performance, in an automated cold-storage warehouse?

Addressing these challenges contributes to improved warehouse system performance and provides valuable insights for the design of future dispatch buffers in in automated cold-storage warehouses. Without design approaches that address this inflexibility, outbound performance remains structurally exposed to truck-arrival uncertainty as automation and throughput continue to grow. This leaves warehouses reliant on manual exception handling to absorb disruptions that an improved buffer system could prevent.

1.4. Research structure

The structure of this thesis is as follows. Chapter 2 provides a literature review, establishing core concepts, warehouse KPI frameworks, and AS/RS and buffer design literature, and identifies the knowledge gap this research addresses. Chapter 3 outlines the research methodology and research design. Chapter 4 presents the system analysis, characterizing the Piacenza warehouse's flows through a flowchart and detailed IDEF0 model, and describing the scenarios for the new Piacenza warehouse. Chapter 5 describes the discrete-event simulation model, including its logic, setup, inputs, KPIs, assumptions, and validation and verification. Chapter 6 reports the scenario design and simulation results, including model validation and sensitivity analysis on the most optimal configuration. Chapter 7 discusses the key findings, their academic and practical implications, the limitations of this research, and directions for future research. Finally, chapter 8 presents the conclusion, answering the main research question.

2

Literature Review

2.1. Core concepts

Within this research, there are a few concepts that are of importance. Definitions of these concepts are given below. These definitions are seen as starting point of this chapter and the research within the field of warehousing.

Warehousing: 'The major roles of a warehouse include the buffering of the material flow along the supply chain to accommodate variability caused by product seasonality or batching in production and transportation, consolidation of products from various suppliers and value-added-processing such as pricing and labelling' (J. Gu et al., 2007).

Automation: 'A system or method that uses, for example, self operating equipment or electronic devices to replace humans in doing routine or repetitive work' (Cooksey, 1986). The routine or repetitive task associated with warehousing of a product, is the storage and retrieval of goods. This leads to automated storage and retrieval systems (AS/RS).

Cold chain logistics: 'A combination of environments, tools, and sequence of actions in a supply chain to keep perishable good within a specified temperature range, to preserve quality and safety for consumption by the end user' (Aung & Chang, 2023).

Cold-storage warehouse: 'Buildings that can maintain a constant low temperature, are widely used in various fields like food, chemistry, pharmaceuticals, vaccines, plasma and scientific experiments' (Gould, 2019).

2.1.1. Warehousing in modern supply chains

At the moment, there is a growing need for efficient logistics due to globalization, customer expectations, rising labour costs, and availability of industrial land (Moghadam et al., 2025). This increasing demand has led researchers and engineers to focus on developing intelligent transportation systems to improve logistics efficiency (Blatnický et al., 2020)(Bozek et al., 2020). Within the supply chain, warehousing operations are an essential part (Albert et al., 2023) (J. Gu et al., 2007) (Moghadam et al., 2025). Warehousing is a significant and vital part of food production and distribution (Cooksey, 1986). There is a distinction to be made between distribution and production warehouses. Distribution warehouses fulfill the purpose of storing products and processing orders from external customers. Production warehouses fulfill the purpose of storing raw materials, work-in-progress-, and finished goods that are associated with manufacturing or assembly (Rouwenhorst et al., 2000). In this research, distribution warehouses are studied.

Productivity is critical in warehouses to maintain efficiency and competitiveness within their supply chain (Albert et al., 2023), also driven by market competition. Technologies such as automation and

communication make further performance improvements possible (J. Gu et al., 2007). 90% of the warehouses worldwide are still managed manually or with low-level automation (Kumar & Mahapatra, 2021). With technologies of Internet of Things (IoT) and Artificial Intelligence (AI) continuously improving, the development automated warehouses has improved. Smart warehouses are a concept resulting from smart technologies that drive industry transformation. The basic characteristics of smart warehouses can be classified into these perspectives: information interconnection, equipment automation, process integration, and environmental sustainability (Zhen & Li, 2022).

Beyond automation, workforce satisfaction and well-being also affects overall warehouse productivity (Trab et al., 2015). Order picking accounts for roughly 50% of warehouse activities and remains largely manual (Battini et al., 2016), as companies often favour human flexibility over automation in dynamic layouts (Roodbergen et al., 2007). The repetitive nature of this work can cause fatigue, discomfort, and injury, underscoring the need to incorporate human factors into warehouse design and operations (Nasir et al., 2024). Non-conventional environments such as cold-storage warehouses further affect employee behaviour (Trab et al., 2015).

Within food product distribution, temperature control is an essential factor affecting product quality and product safety (Meneghetti & Monti, 2015) (Taoukis et al., 2016). 40 % of the food requires refrigeration and 15 % of electricity worldwide is used for refrigeration (Meneghetti & Monti, 2015). Problems in cold supply chains are a known cause of food loss, the global loss of food in weight is estimated at up to 33 % of the worldwide production (Skawińska & Zalewski, 2022). Three types of food supply chains are identified: frozen (below -18 degrees Celsius), chilled (0 to 15 degrees Celsius), and ambient (Akkerman et al., 2010). The size of cold-chain logistics within the logistics industry, was valued at 364 billion US dollars in 2024 (Statista Research Department, 2025). The global cold chain logistics market has grown rapidly, however the level of automation remains low. Automation in cold storage logistics requires a balance between safety and efficiency. There is a need for unmanned operation and human-machine interaction in these warehouses (Zhao et al., 2019).

2.1.2. Warehouse activities and design

There are various processes and activities happening in a warehouse, see figure 1.1. These processes are defined as the following: receiving, storing, order picking, and shipping (Rouwenhorst et al., 2000) (Staudt et al., 2015) (J. Gu et al., 2007). Moghadam et al. (2025) mentions inventory management as well. Receiving is defined as the arrival of goods where trucks are assigned to docks, and scheduling and execution of unloading activities are made (J. Gu et al., 2010). Storing is defined as a reserve area for bulk storage, and a forward area for picking, goods are moved from the unloading area to a designated place (Staudt et al., 2015). A storage system can be integrated with a picking area (Roodbergen & Vis, 2009a). The order picking process is the retrieval from the storage and the transport of goods to be sorted and/or consolidated (Rouwenhorst et al., 2000). In the shipping area, orders are checked, packed and loaded in a carrier (Rouwenhorst et al., 2000). Here, scheduling and assignment of trucks to docks is taking place too (Staudt et al., 2015).

The purpose of warehouse management can be categorized in three main objectives: improving operational efficiency, reducing inventory costs, and improving customer service (Frazelle, 2016). Warehouse layout and operation should be designed to support these objectives, next to constraints as budget, space availability and regulatory requirements (Mohamud et al., 2023). Design decisions influence the implementation costs and have an impact on the maximum throughput capacity of the system (Gagliardi et al., 2012). Logistics costs are determined for a big part in the designphase of warehouses (Rouwenhorst et al., 2000). Mohamud et al. (2023) state that allocation planning and layout design are operations of the warehouse that are implicitly the most important in practice, these operations enhance the probability of success (de Koster et al., 2007) (Nof, 2007). The different design phases are: functional description, technical specification, equipment selection, and layout determination (Rouwenhorst et al., 2000).

Decisions in warehouse design happen on three different levels: strategic, tactical and operational level. Strategic decisions are decisions with long term impact, concerning mostly high investments. Examples are the design of the process flow and selection of the type of warehousing systems (Rouwenhorst et al., 2000)(Moghadam et al., 2025). Tactical decisions are medium term decisions, based on the outcome of strategic decisions. Tactical decisions often include the dimensions of resources, determination of lay-

outs and some organizational issues (Rouwenhorst et al., 2000). In the operational phase, processes have constraints resulting from decisions made on higher levels. Assignment and control problems of people and equipment are challenges in this phase (Rouwenhorst et al., 2000), just like order picking (Moghadam et al., 2025). Resources needed in a warehouse are storage units, a storage system, picking equipment, computer systems, material handling equipment and staff (Rouwenhorst et al., 2000). Buffer design decisions span these levels simultaneously: sizing and layout are tactical choices, but their consequences, such as staff handling schedule disruptions through manual exception handling, surface at the operational level.

2.1.3. Storage systems

Automated storage and retrieval systems (AS/RS) are warehousing systems that are used for the storage and retrieval of products in both distribution and production environments (Roodbergen & Vis, 2009a). AS/RS systems can store and retrieve loads in various settings, and are an innovation where aisle captive cranes serve unit load racks with single or multiple depth and operate completely unmanned (Roodbergen & Vis, 2009a). These systems can store inventory more densely than manual storage systems, which saves space, energy and light, especially when a lot of vertical space is used. Next to that, labour costs are limited, and reliability and accuracy are high with AS/RS systems (Roodbergen & Vis, 2009a). Gagliardi et al. (2012) state that automated storage and retrieval systems and other advanced technologies are used in modern distribution centers because they provide fast, accurate and efficient material handling 24/7. Popovic et al. (2012) mention advantages as integration into automated MHE systems and increased product safety. There are also drawbacks like greater investments costs and flexibility (Roodbergen & Vis, 2009a). Also, the system must first move potential blocking pallets to retrieve the target pallet. Moving blocking pallets from their original positions to target locations is called relocation (Ying et al., 2026). Relocation costs energy and time, and could also create new blocking issues. Resulting in a vicious cycle that significantly increases shuttle workload and reduces overall efficiency (C. Lu et al., 2020). AS/RS are often seen in high volume distribution centres where unit-loads are finished goods pallets (Popovic et al., 2012) (Gagliardi et al., 2012). Introduction of AS/RS was made in cold chain due to unfavourable working conditions (Meneghetti & Monti, 2015).

A basic AS/RS utilizes automated machines or stacker cranes to store and retrieve unitised product, this is a stacked product on a storage pallet (Cooksey, 1986). The system consists of racks that are served by cranes running through aisles between the racks. Racks are metal structures with locations that can store loads, like pallets. Next to the other main components, namely cranes, aisles, I/O (input/output) points, and pick positions. Cranes are machines that can autonomously move, pick-up and drop off loads. Aisles are formed by the empty spaces between racks, where cranes can move. An I/O point is a location where retrieved loads are dropped off and where incoming loads are picked up for storage (Cooksey, 1986).

There are various types of AS/RS that can be categorized into unit-load crane based, mini-load crane based, unit load autonomous vehicle and mini-load shuttle based (Lerher, 2015), described and shown below (2.1).

Type AS/RS	Loadtype	Technology	Description
Unit-load crane-based	Pallets	Stacker cranes	Robust for heavy pallets; high storage density; lower throughput than shuttle-based
Mini-load crane-based	Small SKUs, trays, boxes	Compact stacker cranes	Fast retrieval; high SKU density; less applicable for heavy items
Unit-load autonomous vehicle (AVS/RS)	Pallets	Autonomous vehicles and lifts	Parallel movements and high throughput; scalable and flexible; coordination complexity.
Mini-load shuttle-based	Small SKU's, trays, boxes	Shuttles per tier + lifts	Very high throughput; modular; high investment
Robotic compact storage and retrieval system (RCS/RS)	Totes, trays	Mobile robots on orthogonal grid	Flexible, high throughput, scalable

Table 2.1: AS/RS systems

Unit-load AS/RS are fully automated systems where a single storage or retrieval transaction involves only one unit-load (Gagliardi et al., 2012). Typically the unit-load is a pallet (Roodbergen & Vis, 2009a). The most basic option is the single unit-load aisle captive AS/RS where one unit-load is moved at a time, and racks are single deep which means that every load is accessible for a crane. There are systems where cranes are capable of changing aisles, there are also multi-shuttle cranes that can transport 2 or more loads at a time. A unit-load AS/RS can operate in a single or dual command cycle. Where a single cycle, only storage or retrieval tasks are performed before the crane returns. A dual cycle consists of two moves, one storage and one retrieval (Popovic et al., 2012). Roodbergen and Vis (2009a) mentions carousel systems are suitable for storing small and medium sized products at different levels. Robotic compact storage and retrieval systems are upcoming and applied with goods that are stored in standardised containers and operated from above. Advantages are seen in high throughput, rapid order fulfillment, enhanced redundancy through multiple operating robots and scalability across storage capacity, energy demand, costs. Autonomous vehicle storage and retrieval systems separate horizontal and vertical travel. Vehicles (shuttles) travel horizontally over rails through aisles while lifts are used to transfer loads vertically (Roodbergen & Vis, 2009a). Shuttle based storage and retrieval systems are often used when throughput demand is high. Shuttle based storage and retrieval systems include one or more shuttle carriers, at least one elevator, a racking structure and a control system (Kriehn et al., 2018). The shuttles have the disadvantage of having large investment costs (Popovic et al., 2012).

Shuttle carriers can be implemented aisle and tier, tier to tier, and tier captive. Systems with aisle and tier captive shuttle carriers are often used for high throughput demands. Aisle and tier shuttle systems use a shuttle carrier for every tier, the shuttle and elevator use buffer locations at each tier to store or retrieve totes. Horizontal and vertical transport are decoupled from each other, capable of higher throughput than tier to tier. Tier to tier shuttle carriers cannot change aisles but tiers. Here, the shuttle uses the elevator to change to another tier, reducing waiting times (Kriehn et al., 2018). Tier-to-tier has less shuttles than storage tiers (Wang et al., 2015). Tier-to-tier shuttle based storage and retrieval systems are recognized for their enhanced flexibility, reduced costs, and improved shuttle utilisation (Xu et al., 2025) (C. Lu et al., 2020). Tier-captive means that each shuttle is assigned to a corresponding storage tier in one aisle with an equal amount of shuttles and storage tiers (Wang et al., 2015).

Control policies

Within storage systems, task scheduling affects performance a lot. Task cycle time is measuring the time of all movements of the crane. A crane can spare travel time by combining an inbound- and an outbound job in a single cycle time in conventional AS/RS (Wang et al., 2015). The scheduling problem of AS/RS is similar to the traveling salesman problem .

Sequencing retrieval requests sorts retrieval requests with the aim of minimizing travel time. Sequencing of unit loads for storage can be according to various rules: first-come-first-serve, shortest completion

time, nearest-neighbor, shortest leg (Han et al., 2021), on-line asymmetric Travelling Salesman Problem (Ascheuer et al., 1999). Generally, storage requests come with first come first serve approach. Where the sequence of goods represent a physical flow where it is challenging to adjust the filling. While retrieval requests can be re-sequenced since they are just an information flow (Roozbeh Nia et al., 2017).

The storage location assignment problem is a logistics problem that aims at optimizing item storage location planning inside a warehouse or distribution center (José et al., 2019). The objective is achieved by minimizing the cost and/or time needed for storing and accessing inventory based on the available space, frequency of access to items and proximity of items to the picking area. Popovic et al. (2012) mention random, dedicated and class-based storage.

2.1.4. KPIs for warehouses

Performance analysis has the objective of helping managers evaluate the performance of the enterprise and make decisions in consequence. The warehouse manager uses performance analysis tools and techniques to assess the performance. Performance analysis is defined as the periodic measurement and comparison of actual levels of achievement with specific objectives, measuring the efficiency and the outcome of a corporation (C. S. Lu & Yang, 2010; Staudt et al., 2015). Warehouse efficiency is defined as the ability of a warehouse to optimize its operations and resources in order to achieve the desired level of performance in terms of speed, accuracy and cost-effectiveness (Ballou, 2007). To minimize costs of cold storage facilities, rack configuration, and surfaces and volumes of the cold cell should be optimized (Meneghetti & Monti, 2015). Every design phase has its own performance criteria according to Rouwenhorst et al. (Rouwenhorst et al., 2000). A performance indicator is a numeric value that represents a complex empirical phenomenon to evaluate performances, giving insights to the management office (Faveto et al., 2023). Performance criteria can be seen as design objective or constraint (Rouwenhorst et al., 2000). The data gathered by sensors or operators must be aggregated into valuable tools representing system performance. Indicator use is linked to either the current status of a process to compare it with a benchmark, or to continuously monitor the progress of a process in a specific time frame, or to evaluate the impact of a particular strategy or change by measuring the KPI in two moments (Faveto et al., 2023). There is a distinction made between hard, quantitative and easily measured, metrics and soft metrics, qualitative and measured with more sophisticated tools. The hard measures can be assigned to a single time dimension: time, quality, cost and productivity (Staudt et al., 2015). Automated warehouses represent a different reality from manual or semi-automated warehouses where costs are mostly linked to labor and other variable costs.

The true performance of AS/RS is influenced by other systems. The design of AS/RS consists of the physical design, system choice and configuration. These choices are made, based on historical and forecasted data, product characteristics, available budget, throughput and space. A typical design problem has a given capacity, which means that the number of aisles, rack height and length is constant (Roodbergen & Vis, 2009a).

2.2. KPI selection

Literature is reviewed to analyse KPIs for the measurement of warehouse performance and buffer behaviour, see table 2.2.

Author	Article title	Warehouse type	Method	KPIs	Contribution
(Rouwenhorst et al., 2000)	Warehouse design and control: Framework and literature review	General warehouses (distribution & production)	Framework and literature review	Investment costs, operational costs, volume and mix flexibility, throughput, storage density, response time, order fulfilment quality	KPIs differ depending on strategic, tactical and operational levels; design decisions constrain later KPI improvements.
(Staudt et al., 2015)	Warehouse performance measurement: a literature review	General warehouses	Literature review	Direct: Dock-to-stock time, picking time, order lead time, on-time delivery, fill rate, inventory cost, throughput, labour productivity. Indirect: employee performance, automation level, flexibility, customer perception, value-adding activities	4 KPI dimensions: time, cost, quality, productivity; distinction between direct and indirect indicators.

Author	Article title	Warehouse type	Method	KPIs	Contribution
(Faveto et al., 2023)	Review-Based Method for Evaluating KPIs in Warehouse Systems	General warehouses	ware- Framework for KPI identification, classification and ranking	Generic: bottleneck rate, capacity flexibility, critical WIP, inventory turnover, machine collision, number of failures, object misplacement, peak utilisation, picking accuracy, positioning accuracy, receptivity, resource utilisation, selectivity, shelf occupation, stock balance, throughput, travel distance, unoccupied space, unprocessed order, vehicle capacity, warehouse exposition ration, planning time, queue waiting time, retrieval time, storage time, task time, travel time, warehouse av, lead time, order elaboration time, packing time, cycle time, inventory time. Time: charging platform ev, charging time, Maintenance cost, management cost, Retrieval cost, Space cost, storage cost. Cost: direct labor cost, holding cost, indirect labor cost, inventory cost. ICT: algorithm reliability, bar code reliability, image rec speed, response latency, solver iterations. Environmental: barometric pressure, humidity, pollutant, roof tempreature, temperature. Emission: energy consumption, energy recovery, passive consumption, pollutant emmission, space occupation, vehicle autonomy. Labor: activity automation, human activity time, human error, human utilisation, machine safety, noise, operators per area, work safety	Unique KPI list clustered into economic, environmental and social dimensions with expert validation, measured on relative frequency, citation weighted frequency, uniqueness metric, time weighted frequency.
(Zou et al., 2016)	Modelling Parallel Movement of Lifts and Vehicles in Tier-Captive AVS/RS	Tier-captured shuttle AS/RS	Analytical modelling + simulation	System response time, waiting time, vehicle utilisation, lift utilisation, throughput	Proposes parallel processing policy; validated analytical & simulation model.
(Z. Liu et al., 2021)	Energy Consumption Model for SBS/RS	SBS/RS	Mathematical modelling	Throughput capacity, travel times shuttles and lifts, energy consumption	Quantifies relationship between throughput, travel time and energy use in SBS/RS
(Popovic et al., 2012)	Genetic algorithms for sequencing AS/RS with triple shuttle	Triple-shuttle AS/RS	Genetic algorithm approach	Cycle time, throughput, system efficiency	Improved sequencing algorithm for class-based storage.
(Roodbergen & Vis, 2009a)	Survey of automated storage and retrieval systems	AS/RS	Literature review	Travel time, throughput, waiting time, responsiveness, space utilisation, operational efficiency	Defines AS/RS structure; reviews design decisions and control policies.
(Xu et al., 2025)	Operation time & rack stability optimisation in tier-to-tier SBS/RS	Tier-to-tier SBS/RS	Multi-objective optimisation (NSGA-II, PKBIS)	Total operation time, rack stability	Integrates task sequencing, lane selection and stability as formal KPI.
(Marchet et al., 2013)	Design framework for AVS/RS	AVS/RS	Simulation & case study	Flow time, throughput, equipment costs, racking costs, floor space costs	Identifies bottlenecks for lifts and vehicles; recommends optimal rack configurations.
(Wu et al., 2020)	Modelling & design for SBS/RS	SBS/RS	Open queueing network & simulation	Order cycle time, waiting times per station, maximum throughput, subsystem throughput, facility cost, building cost, utilisation per subsystem	Analytical model comparing SBS/RS and robotic systems.
(Tokat et al., 2022)	Fuzzy C-means KPI design for loading operations	General warehouses (docks)	Fuzzy C-means & simulation	Unit price, monthly demand, products loaded, delayed volume, average delay time, delay ratio, cost of delay	Data-driven KPI clustering for loading operations.
(Bajec et al., 2020)	Warehouse social & environmental performance framework	General warehouses	ware- Multi criteria decision making fuzzy Delphi & BWM	Environmental: GHG emissions, waste discharge, water discharge, energy use, water use, packaging materials, land use, biodiversity effects. Social: safety, training, employee turnover, community engagement, negative local impact	Addresses lack of standardized definitions; proposes structured ESG KPI set.
(Legowo & Wijaya, 2022)	WMS performance using IT-BSC in FMCG	FMCG distribution warehouses	Quantitative case study	Business contribution: cost efficiency, value creation	
Operational excellence: process efficiency, system reliability, operational performance (Baruffaldi et al., 2020)	Performance improvement framework for 3PL picking	3PL warehouses	Decision-support framework	Picking time, picking travelling time, picking distance, SKU turnover, inventory mix, visited slots, inbound/outbound throughput, space utilisation	KPIs based on warehouse diagnosis rather than generic lists.
(Chiaraviglio et al., 2025)	Overall Warehouse Effectiveness (OWE)	General warehouses	ware- Conceptual & mathematical modelling	Overall warehouse effectiveness = availability × performance × quality	Extends overall equipment effectiveness (used in manufacturing) into warehousing, first aggregated warehouse KPI.
(Baglio et al., 2023)	How 3PL providers evaluate warehouse features	3PL warehouses	Multiple case study	Lead time, on-time delivery, flexibility, responsiveness, completeness, order processing cost, labour productivity, utilisation	Connects warehouse building features with 3PL service performance, building characteristics are most important of warehouse features, internal areas and utilities are important for value adding services.

Author	Article title	Warehouse type		Method	KPIs	Contribution
(Chen et al., 2024)	Integrated cost-based warehouse performance model	Retail houses	ware-	PCA & DEA & PIPRECIA & IFR MARCOS	Rental costs, amortisation costs, electricity costs, forklift costs, salaries, shrinkage, consumables costs, number of demands, total turnover, boxes shipped	Shows salaries most critical KPI; robust warehouse efficiency ranking.
(Abdul Rahman et al., 2023)	Decision analysis of warehouse productivity indicators	Distribution centres	cen-	Fuzzy analytical hierarchy process	Space: storage space utilisation, throughput, turnover, transport utilisation, staging area utilisation, building utilisation. Information system: WMS. Labor: picking productivity, manpower productivity, receiving productivity, shipping productivity. Equipment: picking productivity, put-away productivity and shipping productivity	Space most decisive dimension; WMS the single most important KPI.
(Ribino et al., 2017)	Agent-based simulation for improving warehouse layout	Automated warehouses	AGV	Agent-based modelling	Throughput, AGV efficiency, speedup factor, gate configuration impact	Layout & AGV strategy strongly affect performance, nearest neighbor policy is superior to priority policy.
(Mohammed et al., 2025)	When warehouses dream (Industry 5.0 KPIs)	General houses	ware-	Focus groups & G-DEMATEL & F-RAWEC	Traditional: accident rate, equipment downtime, inventory turnover, labour productivity, on-time delivery rate, order cycle time, order picking accuracy, space utilisation, warehouse operating cost. Sustainability: energy consumption, GHG, MHE power sources, waste generation, material used for packaging, percentage recycled material, work-life balance, employee safety. Resilience: risk management, visibility, supply chain desing, contingency planning, redundancy, agility, financial competencies, collaboration, recovery. Industry 5.0: innovation, interoperability, data sharing and transparency, technology based flexibility and adaptability, cognitive computing and AI, mutual cognitive human-robot collaboration, big data and advanced simulation, technology based employee safety and management, cyber physical systems, continuous improvement	Creates first integrated I5.0 performance framework.

Table 2.2: KPI literature overview

The academic literature on warehouse performance measurement shows a wide diversity of key performance indicators, reflecting various warehouse types, processes and design objectives. Early work shows that performance metrics are linked to strategic goals of the warehouse and design decisions impose constraints on performance levels (Rouwenhorst et al., 2000). According to Staudt et al. (2015) definitions for some warehouse performance metrics are missing. KPIs are classified into four categories: time, cost, quality and productivity. Also, a distinction is made between activity specific indicators and resource related indicators. Sustainability related KPIs are often missing in warehouse performance frameworks (Staudt et al., 2015).

In several articles, specialised performance metrics are studied, in particularly automated warehouses (with automatic vehicles, and automated storage and retrieval systems). Research on AS/RS often evaluates cycle time, waiting time, system responsiveness, throughput, and storage and retrieval handling rates (Popovic et al., 2012; Roodbergen & Vis, 2009a; Zou et al., 2016). Studies of Q. Liu et al. (2024), Wu et al. (2020), and Marchet et al. (2013) incorporate energy consumption, subsystem utilisation, relocation times and investment costs too. This indicates a growing need for the linkage of operational performance with energy efficiency.

Faveto et al. (2023) developed a framework for selecting, evaluating, and ranking KPIs. The evaluation is based on the frequency of appearance, citation-weighted relevance, uniqueness and expert assessments. The highest scoring indicators across academic and industrial contexts are: travel time, travel distance, picking time, lead time, cycle time and throughput. Inventory cost, space occupation, machine safety, and energy consumption are also ranked high (Faveto et al., 2023). Developments in performance measurement for WMS and logistics service providers expand the classical operational indicators of performance. System reliability, WMS contributions to business value, inbound/outbound throughput, SKU turnover, and service completeness are mentioned (Baruffaldi et al., 2020; Legowo & Wijaya, 2022). Recent work on industry 5.0 KPIs by Mohammed et al. (2025) shows an integration of traditional KPIs with indicators for sustainability, resilience and human-centricity. Those frameworks emphasise energy consumption, greenhouse gas emissions, waste management, employee safety, training, and system adaptability as essential components of warehouse performance.

KPIs should be incorporated in warehouse design, the best warehouse layout is one that makes most of the space while reducing the amount of time spent travelling and the number of points of contact (Mohamud et al., 2023), next to other KPIs depending on the context. It is recommended to use alternative sources next to academical literature for KPI definition, like industry sources (Faveto et al., 2023).

2.3. AS/RS

Looking at the selection of AS/RS in warehouses, there are several academical sources found that explore different systems. Especially shuttle based storage and retrieval systems (SBS/RS) are mentioned numerously. SBS/RS can have advantages compared to crane based automated storage and retrieval systems when looking at throughput, energy efficiency, robustness and flexibility (Wang et al., 2015) (Kriehn et al., 2018). SBS/RS can achieve higher throughput compared to AS/RS and also require less energy, and greater flexibility (Kriehn et al., 2018). In a comparison of (Ha & Chae, 2018) of tier to tier SBS/RS and tier captive SBSRS, concluded was that tier captive SBS/RS can achieve the desired throughput rate with fewer shuttles. (Küçükyaşar et al., 2021) concluded that certain tier to tier systems outperform the captive systems on initial investment costs, throughput rates and average energy consumption.

Research on AS/RS systems and their optimal design and configurations have been made, a literature overview is shown in table 2.3.

Author	Article title	Warehouse type	AS/RS system	Problem	Method	Key contribution
(Roodbergen & Vis, 2009a)	A survey of literature on automated storage and retrieval systems	General (distribution and production)	Unit-load AS/RS	Review of design and control problems in AS/RS: system configuration, storage assignment, batching, dwell-point positioning, sequencing	Literature review	Structured framework for AS/RS research, most studies focus on static and isolated problems. Lack of dynamic and stochastic modelin, limited integration with other warehouse processes, insufficient joint optimization of design and control decisions
(Wang et al., 2015)	Task scheduling for multi-tier shuttle warehousing systems	Distribution warehouses	Multi-tier shuttle warehousing system (MSWS)	Task scheduling of retrieval operations, coordination between multiple shuttles and single lift	Mathematical modelling, scheduling problem, multi-objective optimization	Captures system dynamics of parallel picking and sequential transport, trade off between shuttle and lift utilisation, pareto optimal solutions for scheduling decisions
(Kriehn et al., 2018)	Impact of class-based storage, sequencing of retrieval requests and warehouse reorganisation on throughput of shuttle-based storage and retrieval systems	Automated distribution warehouse	SBS/RS	Impact of storage management policies on system performance: class based storage, sequencing of retrieval requests, warehouse reorganization. Increase throughput and reduce processing times.	Discrete-event simulation	Storage policies strongly influence throughput, class based storage most impact. Balancing system components, bottlenecks shift depending on configuration. Travel time reduction is key driver of performance. Zoning reduces travel distance for shuttles and lifts.
(Xu et al., 2025)	Operation time and rack stability optimisation in tier-to-tier shuttle-based storage and retrieval systems with multiple retrieval locations	Automated warehouse	SBS/RS	Retrieval optimisation, task scheduling and location assignment. Minimizing total operation time and maximizing rack stability	Mathematical modelling	Introducing joint optimization of task scheduling and location assignment. Rack stability is new objective. Poor scheduling leads to waiting times and inefficiency.
(Wu et al., 2020)	Modelling and design for a shuttle-based storage and retrieval system	Automate warehouse (e-commerce)	SBS	System design and performance modelling	queuing model, design algorithm, validation and case study	First model of tier-captive SBS/RS with cross shuttles and dual command cycles. Analytical queuing model, cost optimisation, real case, simulation validation.
(Zou et al., 2016)	Modeling parallel movement of lifts and vehicles in tier-captive vehicle-based warehousing systems	Automated warehouse (tote-based storage)	Tier captive AVS/RS	Parallel and sequential operation of lift and vehicles effects on system performance	Queuing network, approximation method, validation	Parallel is better than sequential, especially in small systems. For large systems, reducing processing time but increasing waiting time.
(Z. Liu et al., 2021)	Energy consumption model for shuttle-based storage and retrieval systems	Automated warehouse (tote-based storage)	SBS/RS	SBS/RS configuration to minimize energy consumption	Analytical energy consumption model, throughput model, case study	Lift is main energy driver, shuttle acceleration is beneficial, velocity increases gives radical energy consumption. First energy model with cross shuttles, integrating throughput, travel time, and energy consumption.
(Carlo & Vis, 2012)	Sequencing dynamic storage systems with multiple lifts and shuttles	Automated warehouse (distribution and production)	Dynamic AS/RS	Scheduling of multiple lifts and shuttle operations	Heuristic scheduling approach, interference modelling and benchmark comparison	Performance heavily depends on smart scheduling and coordination.
(Ekren et al., 2015)	Warehouse design under class-based storage policy of shuttle-based storage and retrieval system	Automated warehouses	SBS/RS	Optimizing rack design under class based storage policy	Discrete-event simulation	High rise warehouses perform beter, lifts are main bottleneck. Linking storage policy with system design. Design decisions strongly affect performance.

Author	Article title	Warehouse type	AS/RS system	Problem	Method	Key contribution
(Azadeh et al., 2019)	Robotized and automated warehouse systems: review and recent developments	General (distribution and production)	AVS/RS, SBS/RS, RMFS	Design and operational control challenges of robotized warehouse systems	Literature review	Depth-to-width ratio of approximately 2 with centrally located lift performs best; class-based storage reduces cycle time; zone-based vehicle assignment shows diminishing returns beyond a threshold number of zones
(Battarra et al., 2024)	A design framework for shuttle based automated storage systems	Automated distribution	SBS/RS	Physical design features affecting system performance	Design framework/classification	Structures shuttle-system design around three layout dimensions: racks, material handling vehicles, and buffer areas; performance cannot be optimized considering any dimension in isolation
(Y. Li & Li, 2022)	Shuttle based storage and retrieval system: a literature review	Automated warehouse	SBS/RS (aisle- and grid-based)	Configuration of input/output points, lifts, and shuttles	Literature review	Double-command cycles reduce cycle time up to 20% compared to single-command; single-tier shuttle configurations generally most cost-effective
(Eder, 2020a)	Analytical model to estimate the performance of shuttle-based storage and retrieval systems with class-based storage policy	Automated warehouse	Tier-captive SBS/RS	Performance estimation under class-based storage policy	Analytical modelling	Not directly applicable to this research (no class-based storage used), but illustrates strong interaction between storage policy and tier configuration
(D'Antonio & Chiabert, 2019)	Analytical models for cycle time and throughput evaluation of multi shuttle deep lane AVS/RS	Automated warehouse	Multi-shuttle deep-lane AVS/RS	Cycle time and throughput evaluation	Analytical modelling	Storage and retrieval sequencing rule significantly affects system performance, independent of the physical layout
(Lerher et al., n.d.)	Method for evaluating the throughput of shuttle-based storage and retrieval systems	Automated warehouse	SBS/RS	Throughput evaluation	Analytical modelling	Elevators/lifts are typically the system bottleneck determining overall throughput

Table 2.3: AS/RS literature overview

The literature covers a variety of AS/RS configurations, including unit-load crane-based systems, shuttle-based storage and retrieval systems (SBS/RS), and autonomous vehicle storage and retrieval systems (AVS/RS). Across these studies, the main research focus is on improving internal system performance, most notably in terms of throughput, travel time, waiting time, cycle time, energy consumption, and equipment utilisation.

Automated storage and retrieval systems (AS/RS) have been extensively studied in warehouse research, particularly in relation to design, control, and performance optimisation. Early work in this field focuses on providing a structured overview of AS/RS decision problems. For example, Roodbergen and Vis (2009b) classify AS/RS research into design and control decisions, including storage assignment, batching, sequencing, and dwell-point positioning. More recent studies focus on specific AS/RS configurations, particularly shuttle-based storage and retrieval systems (SBS/RS), which are widely used in high-throughput distribution environments. Wang et al. (2015) analyses task scheduling in multi-tier shuttle systems and show that system performance is strongly influenced by the coordination between parallel shuttle operations and sequential lift movements. Similarly, Carlo and Vis (2012) studies scheduling in dynamic AS/RS systems with multiple lifts and demonstrate that system performance heavily depend on coordinated control policies.

Several studies investigate the impact of storage and sequencing policies on system performance. Kriehn et al. (2018) show that class-based storage and retrieval sequencing can significantly improve throughput in shuttle-based systems by reducing travel times. Ekren et al. (2015) further demonstrate that warehouse design and storage policies are strongly interrelated, and that high-rise configurations combined with class-based storage can improve system performance while reducing resource requirements.

Configuration and layout decisions within shuttle-based systems have also received dedicated attention. Battarra et al. (2024) proposes a design framework structuring shuttle-based AS/RS performance around three interacting layout dimensions: racks, material handling vehicles, and buffer areas. And argues that these cannot be optimised in isolation from one another. In a broader review of robotized and automated warehouse systems, Azadeh et al. (2019) similarly finds that a depth-to-width ratio of approximately 2, combined with a centrally located lift tends to produce the best system performance, while zone-based vehicle assignment strategies show diminishing returns once the number of zones exceeds a certain threshold. G. Li et al. (2024) extends this configuration-level perspective in a re-

view of shuttle-based systems, showing that double-command cycles can reduce cycle time by up to 20% compared to single-command operation, and that single-tier shuttle configurations tend to be the most cost-effective choice as adding shuttles per tier increases cost more than proportionally to the resulting throughput gain. Eder (2020b) develops an analytical model estimating shuttle-based system performance, specifically under class-based storage and tier-captive configuration. Although this storage policy is not adopted in the present research, the model illustrates how strongly storage policy and tier configuration interact in shuttle-based systems more generally. D'Antonio and Chiabert (2019) further show, in an analysis of multi-shuttle deep-lane AVS/RS systems, that the storage and retrieval sequencing rule adopted has a significant effect on system performance independent of the physical layout itself, reinforcing that configuration and control-policy decisions cannot be treated separately. Finally, Lerher et al. (n.d.) identify elevators and lifts as the typical bottleneck determining the throughput ceiling of shuttle-based systems, a finding also found in the identification of lifts as the main driver of energy consumption by Q. Liu et al. (2024).

Other research focuses on system design and performance modeling. Wu et al. (2020) develops an analytical queuing model for shuttle-based systems to estimate throughput and cycle times, while Gagliardi et al. (2012) proposes a simulation-based framework to analyse multi-aisle AS/RS configurations under realistic, dynamic conditions. In addition, Marchi et al. (2020) present a design framework for autonomous vehicle storage and retrieval systems (AVS/RS), highlighting the importance of balancing system configuration, throughput, and cost. At the operational level, several studies address coordination and system dynamics. Zou et al. (2016) models the interaction between lifts and vehicles in tier-captive systems and show that parallel operations can improve performance under certain conditions, while also introducing trade-offs related to waiting times. Xu et al. (2025) extends this by jointly optimising task scheduling and storage location assignment, demonstrating that coordinated decision-making is essential to reduce inefficiencies and avoid conflicts within the system. More recently, sustainability aspects have also been incorporated into AS/RS research. Q. Liu et al. (2024) develops an energy consumption model for shuttle-based systems and show that system performance is strongly influenced by the trade-off between throughput and energy usage, with lifts identified as the main contributors to energy consumption.

While AS/RS automation reduces manual retrieval labour, it does not eliminate human involvement in the system. Operational staff remain responsible for exception handling when automated processes encounter disruptions, such as delayed truck arrivals. This operational dimension is absent from the AS/RS literature reviewed, evaluating system performance almost exclusively in terms of throughput, energy, and configuration.

Taken together, this body of work establishes SBS/RS as a well-understood and optimised technology for internal throughput, energy, and configuration performance. However, this research is concerned not with general storage and retrieval performance, but specifically with the outbound dispatch-buffering role introduced in chapter 1. Where the storage needs to hold retrieved pallets flexibly until truck arrival, rather than long-term bulk storage. Viewed through that lens, the other AS/RS categories identified in table 2.1 are less suited to the problem at hand. Unit-load crane-based systems are optimised for dense, long-term bulk storage rather than fast, flexible release, and their comparatively lower throughput (Lerher et al., 2021) makes them a poor match for a buffer that must release pallets on short notice as trucks arrive. Mini-load crane-based and RCS/RS systems are designed primarily for small SKUs, trays, and totes rather than full pallets, which does not match the unit-load pallet flow of the dispatch area of this study. Unit-load AVS/RS offers throughput and flexibility comparable to SBS/RS, but introduces coordination complexity across autonomous vehicles that is not warranted at the relatively small buffer scale considered here. Given the fast-retrieval, small-footprint, and flexible-release characteristics required to decouple crane retrieval from outbound release (Kriehn et al., 2018; Wang et al., 2015), this research therefore scopes its technology choice to shuttle-based storage and retrieval systems (SBS/RS) as the basis for the proposed dispatch buffer design.

2.4. Warehouse layout design

Literature is reviewed to analyse and assess research done on warehouse layout, especially on outbound - and dispatch area layout design, shown in table 2.4.

Author	Article title	Warehouse type	Layout focus	Buffer design	Problem	Method	Key contributions
(Kalafat et al., 2025)	An Integrated Framework for Internal Replenishment Processes of Warehouses Using Approximate Dynamic Programming	Automotive spare parts warehouse	Forward area, reserve storage area, buffer area between forward and reserve storage	No physically designed area or seperated, created whenever needed	Forward reserve warehouse replenishment problem	Approximate Dynamic Programming	First model to include buffer areas in forward-reserve replenishment optimization, prioritizing picking from buffer. Buffer is tactical resource
(W. Gu et al., 2025)	Coordinated Scheduling for Zero-Wait RGV/ASR Warehousing Systems with Finite Buffers	Salt warehouse, automated	Internal flow between inbound and outbound	Entry and exit points for long-term storage	Joint scheduling of RGVs and ASRs	MILP scheduling model, AMHS algorithm, validation	Scheduling model for heterogeneous robots with finite buffers
(Y. Li et al., 2026)	Multi-objective simulation-optimization for a distribution center resource planning	Distribution center	Inbound and outbound buffer areas	Manual inbound and outbound buffer areas, no automation	Multi objective simulation optimization framework	Event driven discrete event simulation and multi objective optimization	Inbound buffer is identified as main bottleneck, workforce and forklift trade offs, case study
(Marolt et al., 2025)	Analytical model and swapping policy assessment of a vertical lift module – buffer integrated storage system	Automated warehouses	Vertical lift module (VLM) with automated buffer storage systems	Reservoir not placed in VLM, swapping cycles between BSS and VLM, back-end storage buffer for VLM. Extended storage capacity	Evaluating swapping policies, minimize unnecessary swaps, throughput insights	Simulation model	First research VLM with automated buffer system, dual command cycles, look ahead strategy as optimal as binary integer program
(Yu & Egbelu, 2008)	Scheduling of inbound and outbound trucks in cross docking systems with temporary storage	Cross docking distribution center	Inbound and outbound area	Temporary storage buffer at shipping dock, manual	Inbound and outbound truck docking sequence, product assignment and use of temprary buffer	Heuristic algorithm, complete enumeration, mathematical mixed-integer programming model	Minimum temporary storage and minimize products routed to storage
(Albert et al., 2023)	Trends and new practical applications for warehouse allocation and layout design: a literature review	General warehouses	The whole warehouse	Forward picking areas, storage capacity planning, dynamic and flexible storage assignment.	Real life constraints in planning and layout design, disconnection academic models and practice	Literature review	Mapping of allocation and layout models. Most existing layout research assumes homogenous SKU, lack of flexible layout design methods, insufficient practical case studies, no methods address temperature controlled environments, very few solutions support resilience
(Amjath et al., 2023)	Trends and new practical applications for warehouse allocation and layout design: a literature review	Manufacturing and logistics system	Inter-facility internal material logistics, outbound logistics	Buffer sizing and placement	Determining optimal buffer sizes, bottlenecks	GEM, finite queueing network model, simulation, case study	Integration of intra logistics and outbound logistics buffer optimization, buffer placement impacting throughput
(Tutam & White, 2019)	A multi-dock, unit-load warehouse design	Unit-load warehouse	Dock placement, warehouse shape factor	Not directly addressed (dock/shape design focus)	Minimizing expected travel distance between dock doors and storage locations	Analytical optimization	Shape factor of approximately 2 performs robustly across a wide range of dock-door configurations; more dock doors does not necessarily reduce expected travel distance
(Guo et al., 2016)	Impact of required storage space on storage policy performance in a unit-load warehouse	Unit-load warehouse	Storage zoning, warehouse shape factor	Not directly addressed (storage zoning focus)	Storage policy performance versus required storage space	Analytical modelling	Class-based storage outperforms full-turnover-based policy; optimal shape factor decreases as demand skewness increases
(Dukic & Opetuk, 2012)	Warehouse layouts	General warehouses	Conventional vs. innovative (fishbone, chevron) layouts	Not addressed	Travel distance reduction in storage and picking areas	Comparative literature review	Innovative layouts reduce travel distance for single-location picks but underperform the conventional layout for multi-location picks
(Yener & Yazgan, 2019)	Optimal warehouse design: literature review and case study application	General warehouses	Storage and order-picking efficiency	Not addressed	Overall warehouse design optimization	Literature review and case study	Design decisions should be evaluated jointly against investment cost, throughput, and order fulfilment quality
(Shahrudnejad et al., 2024)	A novel framework for automated warehouse layout generation	General warehouses	Automated layout generation	Not addressed	Multi-objective layout optimization under physical and functional constraints	Framework proposal	No single optimal layout exists given the multi-objective, constraint-dependent nature of the problem
(Huertas et al., 2007)	Layout evaluation of large capacity warehouses	Large capacity warehouses	General layout evaluation	Not addressed	Layout design under varying operational conditions	Review and evaluation framework	Layout design is inherently problem-dependent; no universally best design, methodology, or policy

Table 2.4: Warehouse layout and buffer literature overview

Studies differ in warehouse type and focus on layout area. They illustrate how layout decisions influence operational performance, congestion and the ability to support reliable inbound and outbound flows. This variety also shows that, while general warehouse layout and shape design has received sustained research attention, papers that integrate buffer design explicitly into that layout question remain limited.

Kalafat et al. (2025) introduces one of the few optimization models which incorporates intermediate buffer areas within a replenishment layout, focusing on using buffers as tactical resources that reduce delays between reserve and forward picking zones. Similar attention to the internal flow is seen in material transport between inbound and outbound points in automated warehouses, where the physical arrangement of entry and exit points constrains robot scheduling and buffer availability. In manual distribution centers, the inbound buffer area is identified as a structural bottleneck, showing that layout elements such as staging capacity, aisle allocation, and forklift travel zones shape not only throughput but also workforce and resource trade-offs. This is one of the few instances in the literature where buffer design is explicitly linked to its staffing consequences, rather than treated purely as a throughput or capacity problem (Li2026MultiobjectiveCenter).

In the study of Marolt et al. (2025), vertical lift modules are researched and the interaction between automated equipment and temporary storage areas are explored. Buffer locations are placed in automated systems or manually operated near dock doors, also confirmed by (Yu & Egbelu, 2008).

Other research addresses warehouse shape and dock configuration more generally, without a specific buffer focus. Tutam and White (2019) shows that a warehouse shape factor (width-to-depth ratio) of approximately 2 performs robustly across a wide range of dock-door configurations. And that adding more dock doors does not necessarily reduce expected travel distance. Guo et al. (2016) similarly find that class-based storage policies generally outperform full-turnover-based policies, and that the optimal shape factor decreases as demand skewness increases. This indicates that layout recommendations cannot be separated from the specific demand pattern of the system being designed. Yener and Yazgan (2019) argues more broadly that warehouse design decisions should be evaluated jointly against investment cost, throughput, and order fulfillment quality, rather than optimising storage or picking efficiency in isolation. This lack of a single transferable solution is made explicit by Shahroudnejad et al. (2024) and Huertas et al. (2007), who each conclude that automated layout generation is inherently a multi-objective, constraint-dependent problem with no universally optimal design. Reinforcing that layout decisions, including buffer placement, must be evaluated within the specific operational context of the warehouse under study rather than applied as a generic best practice.

Literature reviews were done by Albert et al. (2023) and Amjath et al. (2023) who highlight broader structural challenges in warehouse layout models, namely the homogeneous SKU characteristics and the lack of support for flexible, dynamic, or temperature-controlled environments. Models remain mostly through queueing abstractions rather than physical layouts (Amjath et al., 2023).

Layout, buffer placement, and flow design are critical but under-theorised components of warehouse performance. Most contributions either model internal flows in isolation, address general layout and shape questions without engaging with buffer design specifically, or treat buffers as capacity constraints rather than spatial or operational ones.

2.5. Knowledge gap

Warehouse operations are generally structured around three main components: warehouse design and technology selection, storage location assignment, and order picking. However, these components are often studied in isolation rather than as an integrated system (J. Gu et al., 2007; Rouwenhorst et al., 2000). This fragmentation is reflected in the broader warehouse literature, where research mainly focuses on specific subcomponents in the warehouse such as storage systems, order picking, and internal scheduling policies, while the overall flow of goods through the warehouse is less frequently considered (J. Gu et al., 2007).

A significant body of research has been dedicated to automated storage and retrieval systems (AS/RS), particularly in relation to internal system performance. Existing studies primarily focus on improving throughput, travel times, energy consumption, and resource utilisation through optimisation of storage policies, task scheduling, and system design. However, these studies typically analyse AS/RS as self-

contained long-term storage systems and do not explicitly consider interaction with other warehouse processes, such as order picking, outbound staging, and dock operations (Roodbergen & Vis, 2009a; Wang et al., 2015). In addition, the warehouse context in existing studies is often only partially specified. Although some research considers distribution centres or e-commerce fulfillment environments, cold-storage requirements and dispatch-oriented applications are rarely addressed. As a result, AS/RS are primarily positioned as technologies for storage and retrieval in high-bay or reserve areas (Kriehn et al., 2018; Roodbergen & Vis, 2009a; Wu et al., 2020), rather than as potential solutions for outbound staging or dispatch buffering. Consequently, the role of AS/RS in supporting outbound operations under uncertainty remains underexplored.

This gap is particularly relevant for outbound warehouse operations, as their efficiency directly affects the performance of the entire logistics system (Ying et al., 2026). Despite this importance, relatively limited attention has been given to the design and control of outbound buffer processes. In practice, warehouses often rely on buffer lanes in the dispatch area, where pallets are staged in advance of truck arrivals. However, these arrival times are uncertain, and improvements in forecasting accuracy are limited. Moreover, the operation of these buffer areas strongly interacts with upstream storage systems, making them sensitive to fluctuations in demand and system congestion. This limitation is also reflected in the lack of design models that incorporate supply chain-specific constraints, such as temperature control requirements in cold-chain logistics (Akkerman et al., 2010; Meneghetti & Monti, 2015; Zhao et al., 2019). In addition, there is a lack of empirical case studies that apply real-life data to the design and optimisation of warehouse systems, highlighting a disconnection between academic research and industrial practice (J. Gu et al., 2007).

Overall, the literature shows that there is limited attention for the design and optimisation of outbound buffer operations in automated warehouses. In particular, the interaction between staging processes in the dispatch area, stochastic truck arrival patterns, and storage technologies remains insufficiently studied. This gap is even more pronounced in highly automated cold-storage warehouses, where operational constraints and system dynamics differ significantly from conventional warehouse environments. There is not only a technical gap: although some studies acknowledge investment cost as a relevant consideration (Yener & Yazgan, 2019), few frame the technology choice itself as a managerial trade-off between operational flexibility and investment, rather than only an operational performance question. Shuttle-based AS/RS and other technology choices in the reviewed literature are predominantly evaluated on operational consequences, risking an incomplete picture for warehouse operators who must also weigh these trade-offs in adoption decisions.

These knowledge gaps lead to the following research question:

How can a shuttle based dispatch buffer be designed and controlled to minimize crane congestion and outbound delays, improving outbound performance, in an automated cold-storage warehouse?

3

Methodology

3.1. Research design

As introduced, this research investigates optimisation opportunities for the outbound buffer area of automated cold-storage warehouses. Specifically, it analyses the potential of an alternative short-term buffer solution as a replacement or supplement to traditional buffer lanes. An industry promising alternative, a shuttle based storage and retrieval system, will be researched given its suitability according to literature and its significant attention from industry practise. The SBS/RS has operational benefits compared to alternatives, because of the loading type of a pallet, just as a higher throughput than a crane-based system. For a buffer application, this is crucial, even though such a system has higher costs. Shuttles will move horizontally on racking levels and pallets will be vertically moved by lifts, a so called 2D shuttle storage system.

This research follows a simulation-based approach aimed at evaluating and improving outbound buffer configurations under stochastic and dynamic conditions. AS/RS design and control issues are difficult to deal with as the system evolves in highly dynamic, turbulent environments in which demands are frequently subject to change at short notice. Demand is also stochastic and can be seasonal, next to a system state that often changes (Gagliardi et al., 2012). Discrete Event Simulation (DES) is selected as the primary research method, because warehouse systems are characterised by stochastic, dynamic and highly interdependent processes, which are difficult to capture using purely analytical or mathematical optimisation models. To evaluate the performance of the proposed shuttle-based buffer system, a scenario-based experimental design is applied. Different scenarios are defined based on variations in routing flows in the warehouse, next to scenario specific configurations regarding resources and area sizing. Each scenario is evaluated using predefined key performance indicators (KPIs), such as throughput, waiting time, crane utilisation and loading time. This allows for a systematic comparison between the current buffer lane system and the proposed shuttle-based solution. To ensure reliability of results, verification and validation is executed.

The main research question is *'How can a shuttle based dispatch buffer be designed and controlled to minimize crane congestion and outbound delays, improving outbound performance, in an automated cold storage warehouse?'*

This research question will be further analysed and answered by the following sub-questions and methods (table 3.1), to be discussed below.

Table 3.1: Research sub-questions and corresponding methods

Sub-questions	Method(s)
What causes outbound delays in the dispatch area of automated cold storage warehouses, and which design alternatives exist to address these causes?	Literature review Expert consultation Desk research
What are the process flows and system characteristics of the Piacenza warehouse that determine outbound performance?	Desk research Expert consultation
Which system parameters, constraints, and key performance indicators (KPIs) are relevant for modelling and evaluating the dispatch buffer area of an automated, cold-storage warehouse?	Literature review Desk research Expert consultation
How can the context of a shuttle-based dispatch buffer be modelled?	Simulation
How does the shuttle based dispatch buffer perform compared to traditional buffer lanes under different operational conditions?	Simulation

What causes outbound delays in the dispatch area of automated cold storage warehouses, and which design alternatives exist to address these causes?

With the first sub-question, current bottlenecks and challenges that are related to outbound delays and performance, are investigated within literature and within practise at NewCold. This will be done with a literature review, expert consultation, and desk research. Next to adding context and background to this research, the literature review is used to establish a theoretical understanding of outbound processes, common performance issues, and known bottlenecks in warehouse operations. This provides a structured overview of existing knowledge and ensures that the identified challenges are grounded in academic research. Next to this, alternative solutions for outbound staging and layout are explored to broaden the understanding of potential system configurations and identify suitable design options. As highlighted in previous studies, warehouse research often lacks practical validation and may not fully capture the complexity of real-world operations (J. Gu et al., 2007). Therefore, expert consultation with NewCold practitioners is included and internal documentation is assessed to complement the literature findings. This allows for the identification of context-specific challenges, validation of theoretical insights, and a deeper understanding of operational constraints in automated cold-storage warehouses.

What are the process flows and system characteristics of the Piacenza warehouse that determine outbound performance?

In order to redesign and optimize the dispatch area in a warehouse, a thorough understanding of cold-storage warehouse design is crucial, next to the flows inside a warehouse. Information will be gathered through internal desk research and expert consultation. The warehouse flow and processes, including storage, retrieval, and outbound buffering activities, are demonstrated by a flowchart. And an IDEF0 model is made to represent key processes, system components, and material flows on a more detailed level. The purpose of the model is to provide a structured and simplified representation of the real system, which serves as a foundation for the simulation model. It ensures that relevant processes, interactions, and constraints are captured before translating the system into a computational model. Next to the real system, a flowchart and IDEF0 model supports to structure and visualize the adjusted warehouse layout and flows for the buffer storage scenarios.

Which system parameters, constraints, and key performance indicators (KPIs) are relevant for modelling and evaluating the dispatch buffer area of an automated, cold-storage warehouse?

To fully understand the stated problem and operations, system parameters and constraints are defined, explained, and rephrased or redefined where needed. This information is gathered by a literature review, internal data analysis, and expert consultation. Data includes information on truck arrival patterns, crane capacity, buffer capacity, processing times, and control rules governing storage and retrieval operations. Internal data provides quantitative input for model parameters, while expert consultation supports the interpretation of system behaviour, processes, and the identification of critical constraints that may not be directly observable from data alone. Next to that, expert consultation contributes to the decision making regarding modelling assumptions and simplifications. For researching KPIs, literature is assessed and gathered from real-life operations. The KPIs are used to assess system performance in terms of efficiency, responsiveness, and resource utilisation.

How can the context of a shuttle based buffer be modelled?

For this sub-question, the system parameters, constraints and KPIs are all used to form a simulation model that is able to test the outbound shuttle-based buffer design under certain circumstances and variation of system components. Although mathematical modelling and optimisation approaches are frequently applied in AS/RS research, these methods often rely on simplifying assumptions regarding system dynamics and stochastic behaviour. In contrast, simulation allows for a more realistic representation of uncertain truck arrivals, dynamic system states and interactions between storage and dispatch processes. Simulation models are used as decision support systems, allow for different scenarios, and support performance evaluation under different configurations (Tako & Robinson, 2012). Given the focus of this research on outbound performance, crane congestion and stochastic truck arrival patterns, a simulation-based approach is considered most suitable. Discrete Event Simulation (DES) models systems with processes, events, and queues. Individual entities are moving through a system and take into account timing and stochastic behaviour. DES is frequently applied in system performance analysis, order handling, and planning, making it applicable for this research. The simulation model is developed in Simio, a DES environment that supports object-oriented modelling and is well suited for modelling complex material flows, entity routing, and resource contention in warehouse systems. Simio allows process logic to be defined through built-in objects such as servers, paths, and resources, which map naturally into the warehouse components ("Discrete Event Simulation Software (DES) | Simio", n.d.).

Three scenarios are designed where two main alternatives of the shuttle based buffer are implemented into the warehouse layout. The model allows performance differences to be attributed to single changes at a time due to step-by-step adjusting of model configurations regarding resources and layout. Scenario 1 models the current layout of the Piacenza warehouse with traditional buffer lanes, serving the baseline reference and is validated by real-life data of the existing warehouse. Scenario 2 models a new proposed layout for the second warehouse in Piacenza, including the buffer storage for ordered pallets. Scenario 3 adds to scenario 2, to directly store picked pallets coming from the pickfloor, instead of passing the highbay again as in scenario 1 and 2. NewCold has proposed a draft layout for the second warehouse. This layout is analysed based on literature review and configuration changes to the model are applied to check for general layout optimisation results for the outbound area in a warehouse.

The simulation model will be validated and verified using system data and sensitivity analysis. Verification and validation are critical to ensure model credibility and support reliable decision-making based on simulation results (Law & Kelton, 2000). Verification focuses on ensuring that the simulation model is implemented correctly and behaves as intended. This includes debugging, logic checks and sensitivity testing. For verification, a few design phases are identified where the model is built with small extensions. Every time an adjustment is made, the model is checked for realistic values. Also, the entities in the model should move as expected through the model in terms of routing and speed. Very big deviations like a low arrival rate or a big rejection rate, will be tested and checked for to test if the model still behaves as expected. Validation is performed by comparing simulation outcomes with real system data and expert insights from NewCold (Law & Kelton, 2000). This ensures that the model provides a realistic representation of the actual warehouse operations. Model results as utilisation rates of systems, and waiting times for activities, are being compared with actual data. Multiple sessions

with involved NewCold experts are hosted for feedback on the model behaviour and (desired) output. Assumptions are made clear, and not affecting the research scope significantly.

How does the shuttle based dispatch buffer perform compared to traditional buffer lanes under different operational conditions?

Lastly, a comparison and evaluation of the alternative shuttle based buffer scenarios and the traditional buffer lane scenario is made based on performance indicators. The analysis is conducted using the output of the simulation model, in which three scenarios are defined to represent different circumstances. A sensitivity analysis on the scenario and distribution with the most optimal results is done. This analysis includes small, and various adjustments of three relevant parameters will be done to check for logical responses and robustness. These include variations in loading processing times, retrieval activation times, and the share of pick pallets. By systematically varying these parameters, the robustness and performance of each buffer concept can be assessed under both normal and stressed conditions. Based on these findings, a final and complete recommendation is formulated regarding the most suitable buffer design for improving the outbound performance in automated cold-storage warehouses.

3.1.1. Research flow diagram

With the definition of various sub-questions, all complementing to the main research question of this thesis, various phases can be defined. In the research flow diagram (figure 3.1), the sub-questions are connected to the different chapters of this thesis corresponding to the phases of the research. The following phases are defined: understanding and framing, system analysis, design requirements, scenario testing, and evaluation. These phases are spread over the following chapters: introduction, literature review, system analysis, model, scenario design and simulation, discussion, and conclusion.




Figure 3.1: Research Flow Diagram

3.2. Case study

The practical data analysed, researched, and used in this thesis assignment comes from NewCold, just as the experts assessed for various subjects. NewCold is one of the leading companies in automated, cold storage warehouses, operating 26 automated cold storage warehouses worldwide and continuously innovating with their design and technologies (Schuberg Philis, 2023). Their warehouses have a total capacity of 2 million pallet positions. Temperatures in the lowbay area are equal to -18°C and in the lowbay to -25°C degrees Celsius.

The warehouse and therefore the case that is studied for this research, is based in Piacenza, Italy. Here, NewCold is extending the warehouse by building another warehouse next to the existing one, and the company is looking for optimisations in their warehouse design. A shuttle based dispatch area is an idea that is being implemented in new designs of NewCold, showing their trust in the buffer alternative. However, there is no data yet on the performance of such a buffer solution, increasing the need for academic input for an optimal design of this area.

4

System Analysis

4.1. Warehouse processes and flows

The general flow of a pallet in a warehouse is discussed in chapter 1 (figure 1.1). Zooming in on this figure, there are six main processes and corresponding areas defined in a warehouse: inbound docks, I-point, highbay, pickfloor, bufferlanes, and outbound docks. From this, three main activities of a warehouse are defined, with a focus in this research on the last one: processing inbound pallets, storing pallets, and processing orders.

Based on internal conversations at NewCold, a flowchart is made to illustrate the general warehouse process with traditional buffer lanes for the dispatch area, shown in figure 4.1. The flowchart consists of the following components: start/end of the process (ovals), activities (rectangles), and decisions (diamonds). The components are connected with arrows representing the flows between the various components.




Figure 4.1: Flowchart general scenario 1

The start of the warehouse process is recognized as the arrival of an inbound truck, filled with pallets that have to be stored. When there is an inbound dock free, the truck is unloaded by a maximum of two employees that use a forklift or electric pallet truck (EPT) to remove the pallets from the truck and place them on the inbound conveyors. The inbound conveyors take the pallet to the I-point, where the pallet is checked for irregularities and standards. When a pallet gets rejected, it can be either a quick fix or a full rework. When a quick fix is required, the pallet transfers on the conveyor to a buffer lane where an employee performs the fix. After that, the pallet transfers on the conveyor lane to the I-point again. When a full rework is required, the pallet is taken off the conveyor by an employee and taken to the rework station. At the rework station, the pallet is fixed and taken to the inbound conveyor again by forklift or EPT. Once accepted, the pallet is officially registered in the WMS and gets assigned a storage place in the highbay. Based on this place, the pallet is moved to the corresponding aisle and crane, and stored in the highbay.

An order arrival is recognized as the signal to start retrieval of ordered pallets from the highbay. For picked pallets, this is from 30 hours before planned arrival of the outbound truck. For bulk pallets, this is from 6 hours before planned arrival of the outbound truck. Pallets that require picking, move to the pickfloor where picking is performed. Picked pallets move from the pickfloor to the highbay, to be stored until 6 hours before planned pickup. Then, the bulk or picked pallet moves to the bufferlanes. Once the corresponding truck arrives, it goes to the dock with it's corresponding order (being) staged. When the order is completed on the bufferlane, the truck is loaded by a maximum of two employees that use a forklift to remove the pallets from the bufferlane and place them into the truck. The process ends with the filled, outbound truck leaving the warehouse.

In addition to the flowchart, an IDEF0 model is also made. IDEF0 is a function modelling method to model decisions, actions, and activities of a system or organization, now and in the future ("IDEFØ – Function Modeling Method – IDEF", n.d.), giving more detail and information on controls and mechanisms for every activity in the warehouse. The IDEF0 model of the warehouse process is found in appendix A.

4.2. Warehouse characteristics

4.2.1. Piacenza 1

The NewCold warehouse in Piacenza, Italy, includes a highbay, pick floors, inbound docks, and manual outbound docks. In figure 4.2 a simplified layout of the warehouse is shown. Detailed capacity, equipment and layout information is not disclosed in this public version.




Figure 4.2: Piacenza 1 layout

4.2.2. Piacenza 2

The proposed second warehouse in Piacenza includes a highbay, pick floors, a dedicated buffer storage area and both manual and automated outbound handling. The model evaluates how this configuration affects operational performance. Detailed capacity, equipment and layout information is not disclosed in this public version.

4.3. Uncertainties and operational constraints

There are uncertainties that affect daily warehouse operations, each with a different degree of predictability and impact. Truck arrivals is a significant source of uncertainty for outbound performance. Most customers lack automated order systems, requiring the site staff to proactively contact carriers to confirm arrival windows. Next to that, traffic and other unforeseen circumstances make truck arrival

times complex to estimate. The I-point inspection results are also uncertain. Not all inbound pallets pass the I-point inspection on the first attempt. Rejections may result from pallet dimensions, weight, barcode or stacking conditions and may require correction before re-entry into the process. Order composition adds another layer of variability, the share of pick and bulk pallets in an order vary. This results in a picking workload that is unevenly distributed across orders.

Bufferlane behaviour compounds these uncertainties, buffer lanes are physically fixed, accumulating conveyors, meaning that pallets staged in a lane cannot be bypassed. The lane effectively enforces a first-in-first-out sequence and combined with the fact that lanes are rarely fully emptied under normal operations, a single truck that fails to arrive or arrives with delays does not just delay one order. That delay can block the bufferlane and its availability for orders staged.

Alongside the uncertainties coming with the warehouse flow and processes, operational constraints are in place that limit the range of feasible design choices. Temperature requirements are non-negotiable, the highbay must remain at -25°C and the lowbay at -18°C , with airlocks required at every transition point to limit temperature loss and to reduce fire risk through oxygen reduction. Health safety, and food safety regulations constrain equipment design and the automation extent of manual processes. Employee availability is similarly bounded where staff works in shifts with limited continuous time permitted in the cold environment. Commercial constraints also shape operational behaviour. Operating Service Agreements (OSAs) between the warehouse and its customers contain service requirements that influence operational prioritisation. The specific commercial terms are confidential and are not disclosed in this public version.

4.4. Scenarios

4.4.1. Scenario 1

Scenario 1 models the current operational layout of NewCold's Piacenza warehouse and serves as the reference against which all subsequent scenarios are evaluated. The system includes the existing warehouse resources and operating processes, with pallets staged in fixed buffer lanes ahead of truck arrival. Bulk and pick pallets are retrieved in advance of planned truck arrival, with pick pallets activated earlier where additional handling time is required for order picking. All model parameters are derived from real operational data from July 2025. The purpose of this scenario is to establish baseline performance values for the key performance indicators, including truck loading time, crane utilisation, trucks loaded per hour, and bufferlane fill time. No parameter variations are applied within this scenario. The baseline results provide the empirical foundation for assessing whether subsequent design changes lead to meaningful performance improvements or trade-offs.

4.4.2. Scenario 2

Scenario 2 models the layout for the second NewCold warehouse in Piacenza, referred to as Piacenza 2, implementing the buffer storage. The shuttle buffer stores only outbound ordered pallets, both bulk and picked pallets, in the period between retrieval from the highbay and truck arrival at the outbound dock. Pallets are only released from the shuttle buffer to the outbound dock once the corresponding truck has arrived and docked. This approach reduces the dependency of crane retrieval operations on the accuracy of truck arrival time estimates, which is a key source of buffer congestion identified earlier.

With the implementation of a shuttle based buffer storage, the warehouse flow and processes of the system explained in 4.1, changes for the outbound process. The warehouse flow for scenario 2 is visualised in the flowchart below (figure 4.3).




Figure 4.3: Flowchart scenario 2

In this scenario, bulk pallets are retrieved earlier than in the baseline scenario to allow staging within the shuttle buffer prior to truck arrival. Several capacity configurations were evaluated through simulation to identify a stable operating configuration in which no individual resource formed a persistent bottleneck.

Then, the pallet moves to the buffer storage where it will be stored until the corresponding truck arrives at a free dock. The arrival of an outbound truck activates retrieval from the highbay and the pallets are staged on short bufferlanes that are shared by two docks. A maximum of two employees transfer the pallets from the bufferlane into the truck. After all pallets are loaded into the truck, the truck departs.

An IDEF0 model for the changed activities compared to scenario 1, are shown and described in Appendix A.

The proposed, early-stage layout (figure ??) is taken as starting point of the model. The decision to base the layout on this proposed layout is made after careful consideration of literature in the field of warehouse layouts discussed in chapter 2. The layout reflects real site and construction constraints at the Piacenza location, that are taken as given restrictions in this research. A layout only taking academic literature into account limits the relevance to the industry for this research. The layout design literature does not converge on a single optimal configuration, findings on shape factor, dock placement and zoning, are all context dependent and vary with demand, storage policies, and warehouse types. Redesigning the full layout from scratch would require significant research effort, at the expense of the outbound area design focus of this research. Lastly, several of the general design principles identified in literature are derived primarily from ambient or generally automated warehouse contexts. Cold-storage requires strict separation between the areas, and general heuristics could create operational impracticalities.

The layout is tested by performing simulation runs. The proposed layout is adjusted due to computational reasons, as with the expected inbound truck arrival demand the original proposed amount of inbound docks was not enough to handle inbound arrivals. The layout is adjusted to an extra inbound dock. First, the number of operative cranes is varied to assess crane utilisation rates and the extent to which crane congestion constrains outbound throughput. Next to this, the buffer size and shuttles per level are analysed, and tweaked where needed. Lastly, the amount of outbound docks is altered. These parameters represent resource counts and capacities within the fixed layout, rather than structural or spatial elements, and can therefore be adjusted without conflicting with the site constraints described above.

Expected results

The expectations of the results of this scenario focus on the outbound staging area and are shown in table 4.1.

Table 4.1: Expected outcomes for scenario 2

Aspect	Expected outcome
Bufferlane occupancy	Improved compared to scenario 1
Highbay occupancy	Smaller compared to scenario 1
Loading time	Smaller scenario 1
Truck arrival-to-departure time	Smaller than scenario 1
Order activation to retrieval time	Lower than scenario 1

Firstly, the shuttle buffer storage is expected to improve bufferlane occupancy as pallets are only retrieved once the corresponding truck arrives. Highbay occupancy is expected to relatively decrease as pallets are divided between two storage spaces in the warehouse compared to one in scenario 1. Also, the average time between truck arrival and departure is expected to be smaller than the reference scenario due to in-time retrieval instead of scheduling retrieval based on planned arrival. Resulting in the same effect on loading time and in more trucks that can be loaded per hour at a dock in theory, depending on demand. Cranes are used for the same movements as in the reference scenario; therefore, the workload per pallet is expected to be unchanged relatively compared to scenario 1. However, there is more spread in retrieval from highbay, so the average time between order activation and retrieved from highbay is expected to be less than in scenario 1.

4.4.3. Scenario 3

Scenario 3 extends the shuttle-based buffer from scenario 2 to accommodate not only ordered pallets but also pallets after picking. The motivation for this extension is that short-term pallet storage in automated cold storage warehouses currently involves multiple separate physical areas, each consuming crane capacity for storage and retrieval. Consolidating also the flows from the pickfloor to the highbay, has the potential to reduce total crane movements and free highbay capacity for long-term storage operations. In scenario 3, the buffer storage serves two pallet types: ordered pallets awaiting truck arrival and picked pallets coming from the pickfloor directly once picking is completed. Structurally, this changes the routing of picked pallets compared to scenario 2. Rather than returning to the highbay after picking and being retrieved again later, picked pallets are treated identically to a bulk pallet, and move directly from the pickfloor into the buffer storage, skipping the highbay return entirely. This removes one full storage-and-retrieval cycle per picked pallet.

The warehouse flow and processes of the system change again for this scenario compared to the traditional outbound layout of scenario 1 and the proposed layout for scenario 2. Below, in figure 4.4 a flowchart is shown for scenario 3.




Figure 4.4: Flowchart scenario 3

The change in the warehouse process here is the fact that the picked pallet moves directly to the buffer storage after picking. Just as in scenario 2, the activation times for retrieval from the highbay are the same. The retrieval from buffer storage is triggered by outbound truck arrival, just as in scenario 2.

An IDEF0 model for the changed activities compared to Scenario 2 is described in Appendix A.

Scenario 3 initially adopts the same capacity configuration as scenario 2, including the number of shuttles, buffer storage size, operative cranes, and inbound and outbound docks, determined through experiments. Only the pallet-routing logic for picked pallets differs between the two scenarios. After this, capacity dimensioning is done for the buffer storage size, and corresponding shuttles, and highbay size and operative cranes.

Expected results

The expected results for scenario 3 are listed in table 4.2.

Table 4.2: Expected outcomes for scenario 3

Aspect	Expected outcome
Crane utilisation & workload	Substantially lower than scenario 2
Buffer storage occupancy	Higher than scenario 2
Highbay occupancy	Lower than scenario 2
Bufferlane fill time	Slightly smaller than scenario 2
Order activation to retrieval time	Smaller than scenario 2
Loading time	Slightly lower than scenario 2
Trucks loaded per hour per dock	Slightly higher than scenario 2

Compared to scenario 2, the expectation is that the crane utilisation and workload decreases significantly, as picked pallets (around 23% of an order) will not move to the highbay twice. Therefore, is also the crane retrieval time and order activation to retrieval time is expected to be less while throughput is expected to remain unchanged. The buffer storage occupancy is expected to increase as picked pallets stay in the buffer storage for a longer period of time, while the highbay occupancy is expected to decrease due to only once visited by picked pallets. The time between outbound truck arrival and departure (loading time) is expected to slightly decrease as for very early outbound truck arrivals, an order is completed in the buffer storage faster than scenario 2. The amount of trucks loaded per hour is expected to increase slightly compared to scenario 2 due to the improved readiness of orders at planned pickup time while throughput stays the same.

5

Model

5.1. Model logic and processes

Simio is used as discrete event simulation software. This paragraph describes the model logic on entities, objects and the process flow modelled in Simio.

The main physical entity flowing through the system are pallets. Other entities in the model are: inbound trucks, outbound trucks, orders, and ordered pallets. Vehicles that are used in the model are forklifts, these are positioned near inbound and outbound docks for (un)loading and parked when idle at a designated parking spot. Another resource used in the model are cranes, used by the corresponding crane server. For the second model on Piacenza 2, two extra resources are present: shuttles and lifts. However, the shuttles are only modelled as server, not as an additional resource, due to the amount (more than 10) making it unorganized. The paths (where entities travel on) used in the model are timepaths, conveyor belts (including an adjusted conveyor belt item for the monorail), paths, and connectors. Timepaths are paths defined by a set travel time on that specific path, and are used for trucks driving to the inbound- and outbound docks. Conveyor belts are used for inbound lanes, transporting pallets to, from, and in the I-point area. Conveyors are also used to transport pallets in and out of the highbay and buffer storage, next to the outbound (buffer)lanes, and the picking area. The monorail is a specific conveyor belt with its own characteristics in the model. The monorail transports pallets in a circle with pickup- and drop-off points located on the monorail for every connected area. Paths have a distance and speed defined for movements based on real data. Forklifts drive to their destinations in free space, however for picking- and full rework activities they drive on a path. Connectors do not have travel times, and are used for modelling structure and visibility. For example, between the arrival of orders and its separation into pallets. Lastly, workers are also a resource in the warehouse. Workers drive forklifts, do manual reworks, and picking. When idle, workers go to the canteen.

The objects used in the model are: sources, servers, combiners, separators, and sinks. Sources are used to create entities, in this model the entities created are inbound trucks and orders. Servers are used to model processes for every entity that passes the server. The following servers in the model are in place: truck unloading, truck waiting for unloading, I-point registration, small reworks on buffer at I-point, full reworks at rework station, storage and retrieval by cranes, highbay storage, order waiting, order processing (pick and bulk), highbay retrieval, buffer storage, buffer retrieval, storage and retrieval by lifts, storage and retrieval by shuttles, and truck loading. Extra servers to let entities wait under certain conditions are also in place for inbound- and outbound trucks before moving to a dock, for a dual command cycle for crane operations, as well as for orders until they are completed, and before an ordered pallet moves to the highbay retrieval server. Combiners are used at the truck loading process to combine the pallets of an order with a truck. Separators are used to split an inbound truck load into pallets, as well as for orders coming in that are separated into ordered pallets. Sinks are used for ordered pallets once a completed pallet in the highbay or buffer storage is matched and retrieved, orders when an order is completed, and inbound- and outbound trucks after they leave the warehouse. Nodes are used at decision points in the model. For example, to select a crane when a pallet enters

the monorail at the I-point and the bufferlanes track lane occupancy along the outbound staging. An overview of all objects and its internal processes modelled can be found in appendix B, in table ?? where a distinction is made for every scenario.

Inbound docks are assigned randomly to inbound trucks, just as outbound docks for orders and their corresponding trucks. It takes into account if a dock is already busy because no truck can be sent to a dock that is not available. If no docks are available, the truck will wait at the parking lot. Pallets distributed across storage locations according to a simplified location assignment policy described in section 4.2. Cranes are assigned randomly, taking the assignment queue of the crane into account and only item codes that exist in the model are assigned to a pallet based on a random draw that aligns with the cumulative percentage of historic item code availability in the warehouse. Cranes have a dual-command cycle where if available, and inbound movement is combined with an outbound movement. Requirement for this is that a inbound or outbound pallet is already in the assignment queue of that particular crane. Pallets use the same crane on their way to their storage position as on their way out of the highbay, as their position is only reached by a specific crane. Once an order comes in, an ordered pallet gets assigned an item code that corresponds to an existing pallet in the highbay. In case of the buffer storage, the assignment of lifts is done randomly, taking the assignment queue into account. With the lift assignment, a level of the storage is also assigned based on occupation per level. Shuttle assignment on levels is done the same way. Output table for these queues are in place for during run-time checks. Based on designated nodes or servers, any path in the model can have a selection weight based on a certain state of an entity.

State assignments track each entity through the process to enable KPI measurement at each stage as well as a way to model processes in the model. State assignments can be changed during the model. An example of a state assignment is the storage phase of a pallet, and if it is a pick or a bulk pallet. State assignments are also translated into tally statistics during the runtime of the model, tracking certain values per order, pallet, and hourly indicators. These tracked variables are shown during and after runs in output tables and results. A full overview of the state variables tracked per entity and their values can be found in appendix B, table ??.

Property assignments are set values for parameters before initialization of the model, this is used for example for the targeted capacity of the highbay and its occupancy ratio. An overview of the property variables can be found in appendix C, table ??.

5.2. Model setup

The warehouse system modelled in this research is a non-terminating system that operates continuously without a natural end point. For this system, an initial warm-up period is required to bring the simulation from its empty starting state to a steady-state condition, representing real operations before KPI collection begins. The property variable that rounds up the warm-up period is the highbay occupancy, which in the real Piacenza warehouse averages 76,34% in July 2025, combined with a time until the system acts realistically and steady. Starting the simulation with an empty highbay would produce artificially low crane utilization and outbound retrieval times in the early stages of the run. The approach to have a warm-up period with only inbound movements, avoids the need to simulate several months of operations solely to reach realistic occupancy levels, which is computationally costly and introduce unnecessary stochastic variability into the starting state.

Each replication runs for a simulated duration of 2 weeks (14 days) following the warm-up period. Only scenario 1 has a simulated run of a month (31 days). The run length of 2 weeks is selected because of a sample of truck arrival events and outbound delay occurrences, improving the statistical precision of KPI estimates, particularly for rare, but operationally significant events such as buffer overflow and multi-hour outbound delays. It also reduces the influence of day-to-day variability in truck arrival patterns on the results, producing more stable mean estimates across replications. Due to run-time intensity, only the first model has a runtime of 31 days. A total of 5 independent replications are run for each scenario, this number is limited due to timing constraints and intensive runtimes. For scenario 1, 10 replications were executed. A detailed description of the scenario specific warm-up periods is given below.

5.2.1. Scenario 1

A warm-up period of 31 days (744 hours) is applied uniformly across all 10 replications, with the 31-day (744 hours) window used for measurement of the results. This gives a total simulated horizon of 1,488 hours (62 days) per replication. The warm-up length was validated against the highbay occupancy (figure 5.1), averaged across all 10 replications.




Figure 5.1: Warmup stabilisation highbay scenario 1

Occupancy rises steeply from 0% to roughly 76% over the first 740 hours as the model fills from an empty initial state. As a consistency check, the highbay occupancy threshold that drives order generation is independently crossed between 720 and 776 hours across the 10 replications. This is consistent with, and bracketing, the fixed 744-hour warm-up setting for the scenario, despite this threshold-crossing being a stochastic model event rather than the mechanism that determines the warm-up length itself.

A modest downward movement remains visible after warm-up (from a peak of 78% around hour 840 to 71% by the end of the run at hour 1,488). And shows that outbound trucks outpace inbound trucks over the full simulation horizon. This is confirmed by the arrival rates implemented in the model, discussed in appendix C.

5.2.2. Scenario 2

A fixed warm-up period of 23 days (552 hours) is applied uniformly across all replications for scenario 2, with the subsequent 14 days (336 hours) window used for measurement. This gives a total simulated horizon of 888 hours (37 days) per replication. This warm-up length was validated against the highbay occupancy trajectory of the final configuration (7 cranes, 11 outbound docks, 4 shuttles on level 1, 8 buffer storage levels), averaged across all 5 replications. Figure 5.2) shows the warm-up length and stabilization of scenario 2, across all 5 replications.




Figure 5.2: Warmup stabilisation highbay scenario 2

The figure shows that occupancy rises from 0% to approximately 76% over the first 520 hours as the model fills from an empty initial state. Hereafter the line shifts behaviour. As with scenario 1, the order-generation trigger (highbay occupancy reaching threshold) fires stochastically rather than at a fixed simulation time, ranging from 495 to 521 hours across the 5 replications (mean 503.4h), 31 to 57 hours before the fixed 552-hour warm-up cutoff, giving every replication a comfortable margin beyond its own trigger point before measurement begins. A modest downward movement remains visible after warm-up (from a peak of approximately 76% around hour 520 to approximately 66% by the end of the run at hour 888). This is confirmed by the arrival rates implemented in the model, discussed in appendix C.

As a further check, buffer storage occupancy, the metric most directly relevant to this scenario's capacity-dimensioning results, was examined independently (Figure 5.3). It reaches its steady oscillation band (approximately 23-30%) almost immediately once order generation ramps up, and shows no systematic trend after warm-up cutoff. This independently confirms 552 hours as an adequate warm-up length for the metric this scenario's results depend on most directly.




Figure 5.3: Warmup stabilisation buffer storage scenario 2

5.2.3. Scenario 3

A fixed warm-up period of 23 days (552 hours) is applied uniformly across all 5 replications for the final Scenario 3 configuration, with the subsequent 14 day (366 hours) window used for measurement. This gives a total simulated horizon of 888 hours (37 days) per replication. This warm-up length was validated against both the highbay occupancy and buffer storage occupancy, averaged across all 5 replications (figures 5.4 and 5.5). Highbay occupancy rises from 0% to approximately 74% over the first 552 hours, after which its behaviour changes.




Figure 5.4: Highbay occupancy stabilisation, scenario 3

Highbay occupancy rises from 0% to approximately 74% over the first 552 hours, after which its behaviour changes.




Figure 5.5: Buffer storage occupancy stabilisation, scenario 3

Buffer storage occupancy shows an even sharper transition than the highbay occupancy, with its rate of change dropping sharply once the warm-up cutoff is reached. This confirms rapid stabilization for the metric most directly tied to this scenario's capacity-dimensioning results

As with Scenarios 1 and 2, the order-generation trigger fires stochastically rather than at a fixed simulation time, ranging from hour 499 to 515 across the 5 replications. Comfortably before the fixed 552-hour

cutoff in every replication.

5.3. Model input

5.3.1. Piacenza 1

Table 5.1 presents the model input parameters derived from operational data, internal documentation of the Piacenza warehouse and expert consultation. The parameters are organized into four categories: inbound, outbound, and resources. Data was retrieved from NewCold's operational dashboards in PowerBI connected to the warehouse management system. Data was exported of a peak month in 2025, July. Data is analysed with Python.

Operational parameter values reported in this section and in appendix C have been scaled or generalised in accordance with the confidentiality statement. Distribution families, fitting procedures and relative relationships between parameters remain unchanged. Values originating from proprietary NewCold/Newteq technology or non-public supplier specifications are omitted from the public version. The scale factor is known by the thesis committee.

Table 5.1: Model inputs Piacenza 1

Input	Value	Unit	Source
Inbound			
Inbound truck arrival time	Rate table (Appendix C)	Trucks/hour	Data analysis
Inbound truck to dock driving time	10	Minutes	Expert consultation
Trip distribution	$\text{Math.Max}(1, \text{Math.Min}(\text{Math.Round}(\text{Random.Lognormal}(2.9037, 0.3416)), 67))$ (scaled)	Pallets/trip	Data analysis
Manual unloading processing time	$\text{Math.Min}(\text{Random.Gamma}(1.7170, 37.7638), 265.2)$ (scaled)	Minutes	Data analysis
I-point speed	[not disclosed - proprietary system specification]	Pallets/hour	Proprietary system specification
I-point processing time	[not disclosed - proprietary system specification]	Seconds	Proprietary system specification
I-point rejection rate first visit	0.1984 (scaled)	-	Data analysis
I-point rejection rate second visit	0.3364 (scaled)	-	Data analysis
I-point rejection rate third visit	0.4029 (scaled)	-	Data analysis
I-point rejection rate fourth visit	0.4844 (scaled)	-	Data analysis
I-point rejection rate fifth visit	0.5019 (scaled)	-	Data analysis
I-point rejection rate after rework station	0	-	Expert consultation
Reworking at buffer I-point processing time	$\text{Random.JohnsonSB}(-1.0572, 0.6179, 1.6715, 2.4749)$ (scaled)	Minutes	Data analysis
Reworking at buffer rejection rate	0.02739	-	Data analysis
Rework station processing time	10	Minutes	Expert consultation
Outbound			
Order arrival time	Rate table (Appendix ??)	Orders/hour	Data analysis
Trip distributions (pallets/trip)	$\text{Math.Max}(1, \text{Math.Min}(\text{Math.Round}(\text{Random.JohnsonSU}(0.0303, 0.9613, 33.4013, 6.0503)), 66))$ (scaled)	Pallets/trip	Data analysis
Pick distribution of orders	[not disclosed]	%	Data analysis
Bulk distribution of orders	[not disclosed]	%	Data analysis

Continued on next page

Table 5.1 – continued from previous page

Input	Value	Unit	Source
Bulk pallets retrieval activation until planned pickup time	[not disclosed]	Hours	Expert consultation
Pick pallets retrieval activation	[not disclosed]	Hours	Expert consultation
Picking processing time	[not disclosed]	Minutes	Expert consultation
Outbound truck arrival time	Planned pickup time + arrival delay	Minutes	Data analysis
Planned pickup time from moment order comes in	$\text{Math.Min}(\text{Math.Max}(\text{Random.LogLogistic}(63.1949, 2.6190), 12.23), 357.98)$ (scaled)	Hours	Data analysis
Arrival delay from pickup time	$\text{Math.Min}(\text{Math.Max}(\text{Random.JohnsonSU}(0.0336, 0.7250, -26.1801, 96.6533), -1528.5392), 923.1400)/60$ (scaled)	Minutes	Data analysis
Outbound truck to dock driving time	10	Minutes	Expert consultation
Manual loading processing time per truck	$\text{Math.Min}(\text{Math.Max}(1.3357, \text{Random.Weibull}(0.7194, 39.9204) + (1.1380)), 233.6504)$ (scaled)	Minutes	Data analysis
Resources			
Employee work schedule	[not disclosed]	Hours	Expert consultation
Employee off shift rule	Suspend processing	-	Expert consultation
Crane speed	[not disclosed - proprietary system specification]	Pallets/hour	Proprietary system specification
Crane processing time	[not disclosed - proprietary system specification]	Seconds	Proprietary system specification
Crane capacity inbound	1	Pallets	Expert consultation
Crane capacity outbound	1	Pallets	Expert consultation
Conveyor (accumulating) speed	[not disclosed - proprietary system specification]	Pallets/hour	Proprietary system specification
Conveyor (chain) speed	[not disclosed - proprietary system specification]	Pallets/hour	Proprietary system specification
Conveyor speed model	0.45	m/s	Expert consultation
Monorail speed	2.5	m/s	Expert consultation
Monorail capacity	1	Pallets	Expert consultation

Distributions were fitted in Python using the fitter library, evaluating a candidate set of distributions against each dataset and ranks them using sum of squared errors (SSE), the Akaike Information Criterion (AIC), a Kolmogorov-Smirnov (KS) Goodness-of-fit test. A chi square test was also added to these statistical tests. Time distributions apply to the processing times, planned pickup time from order receipt and arrival delay of outbound trucks. The distribution with the best overall ranking across these criteria is selected. Once the best fitting distributions were identified, each were translated into a Simio-compatible random expression, shown in table 5.1. Simio supports about 14 distributions. Because the fitted distributions are unbounded and can occasionally generate implausible extreme values due to outliers in the datasets, chosen was to apply a maximum and/or minimum percentile of respectively 99 and 1. In appendix C, there is a detailed description of the calculations of these values. Inbound truck and order arrivals have been analysed and an interarrival rate table was created in Simio to incorporate for peaks and deviation of a week in July compared to a standard flat distribution. The table is also found in appendix C.

5.3.2. Piacenza 2

Scenario 2 and Scenario 3 share the same set of model inputs, since they differ only in the routing of picked pallets mostly not in system capacity or process parameters. Most inputs carry over unchanged from table 5.1. Table 5.2 lists the parameters that differ from the Piacenza 1 layout for the Piacenza

2 layout, either because they were refitted to Piacenza 2-specific data, or because they reflect structural elements (lifts, shuttles, automatic loading) that do not exist in the Piacenza 1 layout modeled in scenario 1.

Table 5.2: Model inputs that differ for Piacenza 2

Input	Value	Unit	Source
Inbound			
Inbound truck arrival time	Rate table (appendix C)	Trucks/hour	Data analysis
Outbound			
Order arrival time	Rate table (appendix C)	Orders/hour	Data analysis
Bulk pallets retrieval activation until planned pickup time	[not disclosed]	Hours	Expert consultation
Pick pallets retrieval activation until planned pickup time	[not disclosed]	Hours	Expert consultation
Manual loading processing time per pallet	$\text{Math.Min}(\text{Math.Max}(0.0664, \text{Random.LogLogistic}(1.0502, 0.3147)) + (0.0580)), 11.9475)$ (scaled)	Minutes	Data analysis
Automatic loading processing time	[not disclosed]	Minutes	Expert consultation
Resources			
Shuttle processing time	[not disclosed - proprietary system specification]	Seconds	Proprietary system specification
Shuttle speed	[not disclosed - proprietary system specification]	m/s	Proprietary system specification
Shuttle capacity	1	Pallets	Expert consultation
Lift capacity inbound	1	Pallets	Expert consultation
Lift capacity outbound	1	Pallets	Expert consultation
Lift speed 10 - 15m	[not disclosed - proprietary system specification]	m/s	Proprietary system specification
Lift speed 10 - 15m	[not disclosed - proprietary system specification]	Pallets/hour	Proprietary system specification

All other model inputs including manual unloading time, I-point rejection rates, pick/bulk pallet ratio, and pallet trip size distributions, are unchanged from scenario 1 and are not repeated here. Since the Piacenza 2 facility was not yet operational at the time of this research, no historical interarrival-time dataset comparable to the July 2025 Piacenza data was available. The inbound truck arrival time and order arrival time distributions were therefore derived from NewCold's estimated operations timeline, which provides projected hourly inbound and outbound pallet volumes for a representative peak day at Piacenza 2. An interarrival rate table was created in Simio to incorporate for peaks and deviation of a day in July compared to a standard flat distribution. The table is found in appendix C.

5.4. KPI's

Key performance indicators that are evaluated on the models and scenarios include time-based, throughput and productivity, space and storage, resource, and performance KPI's. Supply-chain, social, safety, IT, and environmental KPI's that are discussed in chapter 2 are not taken into account for the analysis of the outbound performance of an automated, cold-storage warehouse. The focus of the model performance measurement lies on throughput and balanced operations in terms of resources and space. The main KPIs evaluated for the model are shown in table 5.3. Supporting indicators are reported in appendix D to explain observed differences in the key KPIs and to support model validation.

Table 5.3: Key KPIs and their applicability by scenario

KPI	Unit	Piacenza 1	Piacenza 2
Trucks loaded per dock per hour	—	X	X
Highbay occupancy	%	X	X
Truck loading time (arrival–departure)	Hours	X	X
Total time on bufferlane	Hours	X	X
Bufferlane fill time	Hours	X	
Truck punctuality	Hours	X	X
Planned pickup to actual departure	Hours	X	X
Crane retrieval time	Hours	X	X
Crane utilisation	%	X	X
Buffer storage occupancy	%		X
Buffer storage to bufferlane time	Hours		X
Order ready wait time	Hours	X	X
Outbound dock assignment wait time	Hours	X	X
Inbound dock utilisation	%	X	X
Outbound dock utilisation	%	X	X

The amount of trucks loaded per dock per hour reflects both the processing and loading time of a truck and the utilisation of a dock, and is the main KPI of this research. Highbay occupancy is the proportion of available highbay storage capacity occupied at a given time, tracked hourly throughout the simulation, and is sized approximately for the peak expected demand. Truck loading time is defined as the time from the arrival of a truck until its departure, independent of the planned pickup time. The total time on bufferlane is the full duration pallets of an order spend staged on the bufferlane, from the first pallet's arrival to the truck's actual departure. This consists of two components: the bufferlane fill time, the time from the first pallet of an order entering the bufferlane to the last pallet of that order entering the bufferlane, and the order ready wait time. This is defined by the time the last pallet of an order enters the bufferlane, until the actual truck departure.

Truck punctuality is the difference between a truck's actual arrival time and its planned pickup time; a negative value indicates an early arrival. Planned pickup to actual departure measures the full cycle instead, from the planned pickup time until the truck is completely loaded and departs. If a truck arrives far outside its scheduled window, its loading duration does not by itself reflect a punctuality problem, since the truck's arrival, not the loading process, caused the deviation. With reporting punctuality and planned-pickup-to-departure separately allows these two effects to be distinguished. Crane retrieval time is the time from order activation until the crane completes retrieval from the highbay. Crane utilisation is the proportion of time each crane spends actively retrieving or storing pallets. Buffer storage occupancy is the proportion of available buffer storage capacity occupied at a given time, tracked hourly throughout the simulation, and is sized approximately for the expected peak demand. Buffer storage to bufferlane time is the travel time from the moment a pallet is retrieved from buffer storage until it enters the bufferlane. Outbound dock assignment wait time is the time a truck spends in the input buffer of the waiting for dock server, where it waits for an outbound dock to be free. Inbound dock utilisation and outbound dock utilization are the proportion of time each dock spends occupied by a truck.

5.5. Assumptions and simplifications

Given the complexity of the real warehouse system, several assumptions and simplifications are made to keep the model within the research scope and time available for this research. Table 5.4 presents these assumptions including the reasoning behind each, and an indication of which scenarios they apply to.

Table 5.4: Assumptions and simplifications

Assumption or simplification	Reasoning	S1	S2/S3
Once a truck does not arrive on its scheduled time, the order will stay on the buffer-lane until the truck arrives	Within the EU, buffer lanes are hardly ever emptied out at NewCold warehouses	x	
Charging of vehicles and resources are ignored	It is assumed that these applications are charged during times that vehicles do not have to be used or that batteries are manually changed.	x	x
Manual labour capacity for inbound and outbound movements, picking and reworks is considered non-limiting for the activities in a warehouse.	Labour capacity is based on peak performance of the warehouse	x	x
Differences between various pallet sizes are not taken into account.	Modelling complexity, dimensions of the warehouse are not modelled. The maximum pallet size is taken as a reference for design, which is an industrial pallet with size 1000x1200mm.	x	x
Pallet orientation and therefore turntables are not considered.	Outside the scope of research.	x	x
No losses of pallets in the warehouse are taken into account. All pallets at the I-point are assumed to be adjusted correctly once rejected maximally 5 times and infeed again.	The number of pallet losses is very low,	x	x
Temperature requirements are not considered in the model.	The warehouse is already completely designed for temperature control and cooling. Airlocks are listed within mechanisms for oxygen reduction design for fire safety in automated areas of the warehouse.	x	x
Deeper levels of MHE systems further than the systems listed are not considered in this research.	Outside the scope of research.	x	x
No reshuffle movements in the highbay are considered.	Modelling complexity.	x	x
No relocations in the outbound buffer storage are considered.	Modelling complexity.		x
Equipment has no maintenance related downtime modeled.	Modelling complexity.	x	x
Stacking and destacking of pallets is not taken into account	A small percentage of pallets come in stacked or are ordered stacked.	x	x
Pick/bulk ratio is applied uniformly per order.	Picking is no key process in the model and there customer data cannot be published.	x	x
The highbay is modelled generically without taking actual storage positions into account	Modelling complexity	x	x
Bufferlanes are modelled with capacity matching the trip size rather than a fixed physical length.	Modelling complexity	x	
Picking is not scheduled	Modelling complexity	x	x
One trip is one order	Modelling complexity	x	x

The assumptions fall either in the category of physical and spatial simplifications, like: pallet size, orientation, stacking, and highbay modelling. These choices are made to keep the model computationally manageable, preserving throughput and timing behavior relevant for this research. Other simplifications reflect resource and capacity: labour, vehicle charging, and maintenance. Assumptions on the process scope like the pick/bulk ratio, pallet losses are based on their low frequency in the real world. Model specific is the buffer lane length that is adjusted to a trip (or order) size, this should be taken into account for the interpretation of KPIs for the reference scenario. Relocations impact the crane, lift, and shuttle activities mirror a less busy system than in reality. Broader implications of these assumptions

and simplifications are discussed in the limitations section.

5.6. Verification and validation

From the conceptual models (flowcharts and IDEF0 diagrams), the warehouse flow including all the rules and assumptions that are made, is proposed to industry experts. Discussions took place about the scope and extent of detail that was necessary to analyse a warehouse with the focus on throughput and outbound performance. With an externally validated, conceptual model of the warehouse and the new warehouse, the simulation model was developed.

Every model and scenario was developed step by step, where the simulation in the beginning is very simple and small with sub-models of all different processes. Models of inbound-, storage-, retrieval-, and outbound processes were merged once the models were verified and working as designed following the conceptual model. The sub-models gave controllable debugging and the certainty of model correctness at every new element or complexity added. Especially the first scenario, the reference scenario of this research, is validated with real-case data extracted from NewCold's warehouse management system and data dashboards. Next to that, various experts from the company were consistently assessed for the validation of input variables, warehouse flows, operational rules and logics. This results in highly validated input for the model. The source of the model inputs is noted in the corresponding tables in this chapter. Results are reviewed for consistency with real life or expected patterns in inbound and outbound flows, process times, and utilization rates. Model parameters are based on real data from NewCold's Piacenza warehouse, and assumptions are described and stated. Simplifications are made, however real life strategies and logics are taken into account.

During the development of the models, output tables are used where various variables like time and occupancy rates are calculated during runs for pallets, orders, throughput and resources. Based on the values in appendix D, models are validated and iterated on as inconsistencies easily show here. Next to the output tables, visible blockages in the model, as well as the automatically generated results table and the model trace is used for the validation and development of the model. Errors during a run show the modelling flaws and are fixed until the model works correctly as it was designed.

Scenario design and Simulation

6.1. Scenario 1

Scenario 1 is modelled to reflect the real warehouse operations in Piacenza. 10 replications are done for a runtime of 31 days, after a warm-up period of the highbay reaching its average occupancy rate (31 days). The occupancy rate of the highbay is the real occupancy for July 2025 in Piacenza.

6.1.1. Validation

Scenario 1 is validated against observed data from the Piacenza warehouse for a representative peak month in 2025. Table 6.1 compares simulated outputs with observed data as a percentage deviation to preserve confidentiality; absolute figures are available in the version held by the thesis committee.

Simulated averages and standard deviations (SD) are computed across 10 replications, measured over a 31-day window. Real averages for truck loading time, truck punctuality, and planned pickup versus actual departure, are calculated from trip-level records ($n = 1,966$). For truck loading time and planned pickup versus departure, an upper-tail 99th-percentile cap was applied to exclude a small number of extreme outliers, consistent with the filtering approach used elsewhere in the model's input distribution fitting. No cap was applied to truck punctuality, as its full range reflects genuine early/late arrival behavior. Highbay occupancy, truck counts, and trucks loaded per dock per hour are reported as single monthly aggregates from operational records.

Table 6.1: Scenario 1 validation

KPI	Deviation
Truck loading time (actual arrival–departure)	+11.3%
Truck punctuality (arrival delay)	+14.3%
Planned pickup to actual departure	+24.3%
Highbay occupancy	-2.3%
Total inbound trucks	-2.3%
Total outbound trucks	-5.6%
Trucks loaded per dock per hour	-5.0%

Simulated standard deviations reflect variability across the 10 replications and support the claim that this replication count is sufficient: for every KPI, the standard deviation is small relative to the mean. This indicates that the simulated average is a stable estimate rather than sensitive to a particular set of random streams.

Not every KPI in the table carries equal evidential weight. Truck punctuality and truck loading time were

used as fitted input distributions, so close agreement here mainly confirms that these distributions were implemented correctly rather than providing independent evidence that the model behaves realistically. Planned pickup to actual departure is not itself a fitted input, but as a measure it is largely derived from these same two distributions. It inherits similar sensitivity, the real-data distributions of all three KPIs are heavily skewed, meaning a resampled mean is more sensitive to which tail it happens to draw in a given set of replications.

The simulated distributions closely track the real shape, including the right skew in loading time and the heavy tails in punctuality and planned-pickup-vs-departure, confirming that the underlying fitted input distributions (arrival delay and loading time) were implemented correctly, with the derived planned-pickup-vs-departure measure showing the same close agreement as a consequence.

The remaining KPIs: highbay occupancy, total inbound- and outbound truck counts, and trucks loaded per dock per hour are outputs of the model rather than resampled inputs. These provide the strongest evidence for the validity of the model. As shown in table 6.1, these four KPIs deviate from the real averages by 2.3-5.6%, notably lower than the 11.3-24.3% deviation seen in the resampled-input KPIs. Trucks loaded per dock per hour shows the closest agreement (5.0% deviation), followed by highbay occupancy (2.3%) and total inbound trucks (2.3%), with total outbound trucks showing a somewhat larger gap (5.6%). Taken together, this is the strongest and most direct validation evidence available for scenario 1, since these KPIs emerge from the model's internal logic rather than being drawn from the same fitted distributions used to construct its inputs.

6.1.2. Results

Table 6.2 summarizes the full set of the KPIs of scenario 1, used as a baseline against scenario 2 and 3. For KPIs measured per truck or per hour, the mean, standard deviation, minimum and maximum values are computed across the 10 replications, and shows replication variation. For the four utilisation KPIs (cranes, inbound and outbound docks, and outbound dock waiting) the value is already an average per physical resource, the values are describing how utilisation varies across resources of that type, not across replications. KPIs that are compared to real data are not disclosed in this table.

Table 6.2: Scenario 1 KPI summary

KPI	Mean	SD	Min	Max
Trucks loaded per dock per hour	-	-	-	-
Highbay occupancy	-	-	-	-
Truck loading time (arrival–departure)	-	-	-	-
Total time on bufferlane (first pallet–departure)	3.40 h	0.08 h	3.30 h	3.51 h
Bufferlane fill time	0.22 h	0.01 h	0.21 h	0.23 h
Order ready wait time	3.18 h	0.08 h	3.07 h	3.31 h
Truck punctuality	-	-	-	-
Planned pickup to actual departure	-	-	-	-
Crane retrieval time	4.43 h	0.25 h	3.91 h	4.77 h
Crane utilization	45.6%	0.07 pp	45.5%	45.7%
Inbound dock utilization	47.6%	8.27 pp	37.2%	59.2%
Outbound dock utilization loading	18.9%	0.83 pp	17.5%	20.2%
Outbound dock utilization waiting for order ready	3.0%	0.21 pp	2.6%	3.4%

Crane utilisation averages 45.6%, with a standard deviation of only 0.07 percentage points across the cranes and almost no spread between the busiest (45.7%) and least busy (45.5%) crane. This shows that retrieval demand is spread evenly across the highbay's aisles. Figure E.3 shows crane retrieval activity over a representative day. Inbound dock utilisation shows substantial variation (a standard

deviation of 8.27 percentage points, range 37.2-59.2%, and an average of 47.6%). The utilisation decreases the first to last dock, consistent with the dock-search-order pattern identified previously rather than genuine demand differences between docks. Outbound dock utilisation is narrow and low (standard deviation of 0.83 percentage points, range 17.5-20.2%, and an average of 18.9% across the docks). The amount of outbound trucks in the system remain substantial throughout a representative day, at roughly 40-55 trucks per day. This is well above the number of outbound docks, confirming that most outbound trucks are queued before their order is ready rather than waiting at a dock itself, consistent with the low outbound-dock-loading utilisation reported above. In appendix E, figures E.1 and E.2 show inbound and outbound trucks in the warehouse over the same day. The additional time outbound trucks spend waiting for their order to be fully staged before loading can start, averages 3.0% of dock time (standard deviation of 0.21 percentage points), indicating that once a truck occupies an outbound dock, it is loading for the large majority of that time rather than waiting, consistent with the modelled waiting server before the outbound docks.

The bufferlane fill time (0.22 h and a standard deviation of 0.01 h) remains short relative to the order ready wait time (3.18 h and standard deviation of 0.08 h). This shows that once the first pallet of an order reaches the bufferlane, the remaining pallets follow quickly. Most of the delay before a truck can depart still comes from waiting for the truck itself to arrive and begin loading, or from the start of retrieval, rather than from staging the remaining pallets.

To complement these averages with a view of within-day dynamics, figure 6.1 shows pallets in the highbay over a representative day (15th day of runtime), averaged across the 10 replications. Shown are dips through the morning before refilling later in the day. In figure E.4 the corresponding inbound registration throughput driving this refill is shown. Across all KPIs, replication-level standard deviations are small relative to their means, supporting the adequacy of the 10-replication design used throughout scenario 1.




Figure 6.1: Pallets in highbay over a representative day, scenario 1

6.2. Scenario 2

Scenario 2 is modelled to simulate and analyse the impact of a shuttle buffer storage, replacing traditional bufferlanes, for a second warehouse in Piacenza. The first experiment was based off the proposed layout, and from there modifications were made, resulting in an optimal configuration of the warehouse. In total, seven configurations were tested. For all experiments, 5 replications are done for a runtime of 14 days, after a warmup period of the highbay reaching its average occupancy rate of scenario 1, and 24 hours before the model reaches a steady state (total is 23 days).

6.2.1. Model configuration

The model configuration was based on the same KPIs driving the modelling decisions: truck loading time, planned pickup to actual departure, trucks loaded per dock per hour, crane utilization, buffer storage occupancy, and outbound dock wait time. Table 6.3 and 6.4 show all configurations made on the Piacenza 2 model. Below the tables, each configuration is discussed.

Table 6.3: Scenario 2 capacity dimensioning KPIs

KPI	Configuration 1 -		Configuration 2 + 1 crane		Configuration 3 + 1 ob dock		Configuration 4 + 1 shuttle	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Truck loading time	56.36 h	12.78 h	52.98 h	7.81 h	54.96 h	29.24 h	17.51 h	11.42 h
Planned pickup to actual departure	55.57 h	12.79 h	52.24 h	7.83 h	54.21 h	29.29 h	16.76 h	11.42 h
Trucks loaded per dock per hour	0.234	0.012	0.236	0.014	0.211	0.053	0.292	0.039
Crane utilization	97.23%	1.4 pp	84.44%	0.89 pp	83.68%	1.9 pp	84.41%	1.53 pp
Buffer storage occupancy	149.8%	23.3 pp	126.1%	16.5 pp	115.7%	41.9 pp	51.7%	19.7 pp
Outbound dock assignment wait time	45.98 h	10.98 h	40.13 h	4.68 h	38.38 h	19.02 h	10.81 h	8.21 h

Table 6.4: Scenario 2 capacity dimensioning KPIs (continued)

KPI	Configuration 5 + 1 level		Configuration 6 + 1 level		Configuration 7 - 1 ob dock	
	Mean	SD	Mean	SD	Mean	SD
Truck loading time	7.79 h	11.61 h	1.94 h	0.26 h	1.97 h	0.32 h
Planned pickup to actual departure	7.17 h	11.55 h	1.26 h	0.21 h	1.32 h	0.52 h
Trucks loaded per dock per hour	0.327	0.025	0.345	0.013	0.377	0.012
Crane utilization	83.21%	1.65 pp	83.67%	2.38 pp	83.30%	1.44 pp
Buffer storage occupancy	37.7%	24.2 pp	27.2%	0.53 pp	26.9%	0.95 pp
Outbound dock assignment wait time	4.78 h	9.96 h	0.026 h	0.033 h	0.052 h	0.051 h

Configuration 1

The baseline configuration does not reach steady state within the measurement window. Buffer storage occupancy exceeds nominal capacity around hour 650 and continues rising through the end of the run (figure E.5), reaching 280–340% of the proposed capacity across replications. Reported KPIs for this configuration therefore reflect an actively insufficient system rather than a stable operating point, and are included here only to characterize the severity of the crane bottleneck that motivated the subsequent capacity increase.

Configuration 2

The second configuration consists of an extra crane compared to configuration 1. Increasing the cranes reduces crane utilisation substantially, from 97.2% to 84.4% on average, and correspondingly shortens crane retrieval time (from 2.66h to 0.51h). However, this configuration also fails to reach steady state within the measurement window. Buffer storage occupancy again exceeds nominal capacity, around hour 600–700 depending on replication, and continues rising through the end of the run (figure E.6), reaching 180–350% of capacity across replications, a comparable magnitude to the baseline despite the substantially lower crane utilisation. This indicates that cranes, while a genuine constraint, are not the sole bottleneck limiting this layout: relieving crane pressure alone is insufficient to bring the system into balance, and a further capacity constraint elsewhere in the buffer storage pathway is still being exceeded by demand. Reported KPIs for this configuration are, as with the baseline, included only to characterize this continued instability rather than as a stable operating point.

Configuration 3

The third configuration consists of an extra outbound dock compared to configuration 2. Increasing outbound docks, changes essentially nothing about the underlying highbay bottleneck. Crane utilisation is similar as in the last configuration (83.7% vs. 84.4%) and outbound dock assignment wait time (38.4h vs. 40.1h) too. The additional dock does not address the constraint limiting throughput, and this configuration again fails to reach steady state. Buffer storage occupancy reaches 139–286% of capacity across replications by the end of the run E.7. The extra dock does change variability. Loading time's standard deviation nearly quadruples, from 7.81h to 29.24h, with individual replications ranging from 18.0h to 93.1h despite running the identical configuration. This sign that with no spare capacity anywhere in the pipeline, small differences in demand happen to arrive compound into widely varied outcomes rather than averaging out. Also, level 1 shuttle utilisation rises sharply in this configuration, from 77.7% to 93.8%, while all other levels remain in the 38–53% range. This confirms that the shuttle-level imbalance identified in the previous step is not only present but intensifying as other constraints are relieved, and motivating shuttle capacity as the next parameter to test.

Configuration 4

The fourth configuration consists of another shuttle on the first level of the buffer storage. Adding a this shuttle to the first level produces the first configuration to achieve genuine steady-state behavior, though only partially (figure E.8). Level 1 shuttle utilisation falls sharply, from 93.8% to 56.5%, and the reported averages improve substantially across every KPI. Loading time decreases from 54.96h to 17.51h, and buffer storage occupancy decreases from 115.7% to 51.7%. Examining the full trajectory confirms this is a real improvement, not only a lower average masking the same instability. All five replications settle into a genuine plateau around 15–35% occupancy from roughly hour 500 to 650. However, this stability is not sustained for the full run. In every replication, buffer storage occupancy begins climbing again around hour 650–680, diverging by the end of the run. Two replications even exceed 170% of capacity, while the remaining three end between 72% and 95%.

Examining shuttle utilisation by level at the replication level, clarifies what actually drives this divergence. It is not that one specific level is structurally under-resourced, but that whichever level happens to receive a disproportionate share of demand in a given replication's random draw becomes that replication's bottleneck, since three shuttles per level leaves no slack to absorb such a concentration.

Configuration 5

Configuration 5 consists of an extra level in the buffer storage. Increasing buffer storage produces a substantially more robust result than any previous configuration (E.9). Four of the five replications now achieve genuine steady-state behavior for the entire 888-hour run. Buffer storage occupancy

settles around 20-35% and remains there throughout. Replication 5 is the clear exception, destabilising from around hour 600 onward and reaching 192% of capacity by the end of the run. Because this replication diverge so sharply from the other four, the statistics in table 6.4 should be read with this in mind. Outbound dock assignment wait time, for instance, averages 4.78h with a standard deviation of 9.96h, but this reflects near-zero waits (0.01-1.24h) in four replications against 22.57h in replication 5 alone, not a genuinely representative spread.

The resource-level picture also improves significantly at the aggregate level. Shuttle utilisation is now much more even across all levels (52.5-66.6%), with no single level dominating. Lift utilisation shows a similar convergence: the pronounced split between inbound and outbound lifts seen in every prior configuration (roughly a 1.5-1.6x gap) has narrowed to near parity (35.6-39.0%). This shows a more balanced inflow and outflow of the buffer storage.

Configuration 6

The sixth configuration consists of an additional level in the buffer storage compared to configuration 6. Increasing buffer storage resolves the instability seen in every previous configuration. All five replications now achieve genuine steady-state behavior for the entire 888-hour run, with buffer storage occupancy settling tightly around 20-40% and no replication showing the late-run divergence seen at in other configurations (figure E.10). This is reflected directly in the per-replication details, shown in table E.8. Unlike configurations 4 and 5, where one or two replications diverged sharply from the rest, every KPI here is tightly clustered across all five replications, with loading time ranging only 1.63-2.27h and buffer storage occupancy only 26.4-27.9%. Outbound dock assignment wait time falls to near zero (0.026h on average, compared to 4.78h at 7 levels and 45.98h at the baseline), and loading time settles at 1.94h, closely matching this scenario's target benchmark of comparable performance to the validated Scenario 1 baseline.

The resource-level picture is consistent with this outcome. Lift utilisation has fully converged, with all lifts within a tight 35.9-39.3% band, compared to the roughly 1.5-1.6x gap between the busiest and least-busy lifts present in every prior configuration. This imbalance resolved on its own alongside the shuttle-level fix, without requiring its own dedicated intervention. Shuttle utilisation remains reasonably balanced across all levels (47.9-70.2%). Although level 1 shows a higher standard deviation across replications (8.2 percentage points) than any other level, a residual signal is not destabilising the system.

Configuration 7

The seventh, and final configuration decreases an outbound dock from configuration 6. With the buffer storage and shuttle-level instability fully resolved at configuration 6, the final step tests whether the outbound dock added at configuration 3, is actually necessary. Reducing outbound docks, with all other parameters held at configuration 6's settings, reproduces configuration 6's results almost exactly: loading time (1.97h vs. 1.94h), buffer storage occupancy (26.9% vs. 27.2%), and crane utilization (83.30% vs. 83.67%). These KPIs are all within normal replication variance of one another, and every replication remains fully stable for the entire run (table E.9). All replications here also show steady-state behavior E.11.

This confirms directly what configuration 3 first suggested: the extra outbound dock provides no measurable benefit, since removing it produces no significant change in system performance. Individual crane utilisation remains tightly clustered (82.61-83.79%), and lift utilisation has fully converged to near parity across all lifts (35.60-38.65%). This matches configuration 6 and confirming the lift imbalance identified in earlier configurations remains resolved regardless of outbound dock count. Shuttle utilisation across all levels remains reasonably balanced (46.95-69.83%), with level 1 again showing the highest standard deviation across replications (8.05 percentage points) of any level.

Configuration 7 is therefore adopted as the final scenario 2 configuration: it matches configuration 6's fully resolved, stable performance while using one fewer outbound dock, with total loading time of 1.97h closely comparable to scenario 1's own validated benchmark.

6.2.2. Results

Table 6.5 summarises the full set of KPIs for the final scenario 2 configuration, used as the baseline against which scenario 3 is compared. The differences between KPIs of scenario 1 are shown, but

must be taken into perspective as the model operates on different arrival patterns and processes.

Table 6.5: Scenario 2 KPI summary final configuration

KPI	Mean	SD	Min	Max	vs. S1
Trucks loaded per dock per hour	0.377	0.012	0.366	0.391	+65.4%
Highbay occupancy	70.16%	1.14 pp	68.81%	71.32%	-5.9%
Truck loading time (arrival–departure)	1.97 h	0.32 h	1.56 h	2.38 h	-23.0%
Total time on bufferlane*	0.74 h	0.09 h	0.63 h	0.85 h	-78.2%
Truck punctuality	-0.65 h	0.23 h	-0.96 h	-0.32 h	-9.7%
Planned pickup to actual departure	1.32 h	0.52 h	0.80 h	1.88 h	-28.3%
Crane retrieval time*	0.46 h	0.04 h	0.42 h	0.51 h	-89.6%
Crane utilization	83.30%	1.44 pp	82.61%	83.79%	+82.7%
Buffer storage occupancy	26.92%	0.95 pp	25.65%	27.92%	n/a
Buffer storage retrieval time	0.49 h	0.04 h	0.45 h	0.56 h	n/a
Buffer storage to bufferlane time	0.11 h	0.004 h	0.11 h	0.12 h	n/a
Order ready wait time*	0.089 h	0.0020 h	0.087 h	0.091 h	-97.2%
Outbound dock assignment wait time	0.052 h	0.051 h	0.006 h	0.131 h	n/a
Inbound dock utilisation	88.15%	–	87.57%	88.66%	+85.2%
Outbound dock utilisation loading	19.26%	–	5.00%	36.26%	+1.9%

Four KPIs marked with an asterisk are not directly comparable to scenario 1 despite sharing a name: total time on bufferlane, crane retrieval time, and order ready wait time all measure a materially shorter segment of the process in scenario 2 than in scenario 1. The large apparent deviations for these three KPIs (-78% to -97%) therefore reflect this structural difference in what each metric captures, not a genuine reduction in overall process time. Buffer storage occupancy, buffer storage retrieval time, buffer storage to bufferlane time, and outbound dock assignment wait time have no equivalent at all, since scenario 1's layout has no buffer storage stage. Among the KPIs that are genuinely comparable, scenario 2 shows improved truck-facing performance relative to scenario 1: loading time falls by 23.0%, and planned pickup to actual departure falls by 28.3%. Crane utilisation rises by 82.7% and inbound dock utilisation by 85.2% relative to scenario 1. This indicates that scenario 2's faster truck turnaround is achieved with higher resource utilisation.

Crane utilisation averages 83.30%, with individual cranes tightly clustered (82.61-83.79%, figure E.9), consistent with the even retrieval demand seen in scenario 1. Inbound dock utilisation is similarly tight across the docks (87.57-88.66%, averaging 88.15%) but at a much higher level than crane utilisation, consistent with the limitation discussed in 7 that inbound dock capacity was never varied during the capacity-dimensioning process and may warrant its own investigation.

Outbound dock utilisation, by contrast, shows substantial spread across the individual loading docks (5.00-36.26%, averaging 19.26%) despite the aggregate outbound-dock-assignment wait time being negligible (0.052h). This uneven distribution across docks, without any corresponding queueing delay, suggests a dock-assignment pattern rather than a capacity constraint. The assignment pattern is based on shared bufferlanes between outbound docks. Shuttle and lift utilisation show the converged, well-balanced pattern established at this configuration, with no individual level or lift standing out as a constraint.

To complement these run-level averages with a view of within-day dynamics, figures 6.2 and 6.5 show inbound trucks in system, outbound trucks in system, crane retrievals per hour, and pallets in buffer storage over a representative day within the measurement window (day 5), averaged across the 5

replications. Pallets in buffer storage (figure 6.5) show a clear within-day pattern, dipping to their lowest point overnight (hours 8-13) before climbing to a peak in the evening (hours 21-23). This is mirroring the pattern found for highbay pallets in scenario 1 (figure 6.6).




Figure 6.2: Inbound trucks in system over a representative day, scenario 2




Figure 6.3: Outbound trucks in system over a representative day, scenario 2




Figure 6.4: Crane retrievals per hour over a representative day, scenario 2




Figure 6.5: Pallets in buffer storage over a representative day, scenario 2




Figure 6.6: Pallets in highbay over a representative day, scenario 2

6.3. Scenario 3

Scenario 3 replaces the retrieval routing used in scenario 2 with a direct-to-buffer-storage path that skips the return trip to the highbay. Rather than repeating scenario 2's full bottom-up capacity search, the configuration process for scenario 3 starts from scenario 2's final layout and tests whether this same layout remains sufficient under scenario 3's routing, adjusting capacity only where it does not. For all experiments, 5 replications are done for a runtime of 14 days, after a warm-up period of the highbay reaching its average occupancy rate of scenario 1, and 24 hours before the model reaches a steady state (total is 23 days).

6.3.1. Model configuration

Table 6.6 summarizes the KPIs for each configuration tested in the configuration process. The following subsections discuss each configuration and the reasoning behind each subsequent adjustment. Buffer storage occupancy, highbay occupancy, and throughput were found to be affected by a warm-up period misconfiguration and were corrected by recomputing them directly from raw simulation output. Crane, shuttle, and lift utilisation were affected by the same issue but have no equivalent raw data source, so they could not be corrected in the same way. These figures are therefore omitted for configurations 1-3 and reported only for configuration 4, which was independently rerun and confirmed.

Table 6.6: Scenario 3 capacity dimensioning KPIs

KPI	Configuration 1 -		Configuration 2 + 1 level		Configuration 3 - 1 crane		Configuration 4 -1 crane	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Truck loading time	18.54 h	19.56 h	1.80 h	0.08 h	1.86 h	0.10 h	1.90 h	0.23 h
Planned pickup to actual departure	17.86 h	19.61 h	1.07 h	0.11 h	1.10 h	0.15 h	1.12 h	0.25 h
Trucks loaded per dock per hour	0.331	0.058	0.382	0.012	0.379	0.013	0.383	0.006
Crane utilisation	–	–	–	–	–	–	94.53%	0.28 pp
Buffer storage occupancy	67.9%	49.6 pp	28.4%	0.53 pp	28.1%	1.08 pp	28.2%	0.51 pp
Outbound dock assignment wait time	14.48 h	17.90 h	0.025 h	0.012 h	0.040 h	0.021 h	0.067 h	0.084 h

Configuration 1

The first configuration is the final scenario 2 configuration. Truck loading time is significantly higher than the loading time in scenario 2 (a mean of respectively 18.54h versus 1.97h), and has a high standard deviation of 19.56 hours. Not only seen in truck loading time, but also for planned pickup to actual departure, buffer storage occupancy, and outbound dock wait time, point to a system with no steady-state behaviour and high variations throughout the simulation time.

The same layout that reached full steady state under scenario 2's routing does not do so here. Three of five replications remain stable throughout. But two replications (4 and 5) destabilise from around hour 600, and a third (replication 2) destabilises later, around hour 800, ending near 125% of capacity by the end of the run (Figure E.12).

The mechanism behind this instability is less clear-cut than in scenario 2. Buffer storage retrieval time is consistently and substantially elevated in the unstable replications (1.84-1.91h versus 0.55–0.63h in the stable replications), confirming the downstream consequence of the instability even though its root cause is not confirmed. The system shows clear signs of operating without spare capacity, with buffer storage retrieval time substantially elevated in the unstable replications. Crane utilisation is not reported for this configuration, but the routing change relative to scenario 2, is expected to reduce crane demand. The buffer storage constraint identified here is therefore assumed to lie downstream of the cranes, motivating buffer storage capacity as the next parameter to adjust rather than crane count.

Configuration 2

The second configuration has another level in the buffer storage. Increasing buffer storage fully resolves the instability identified in configuration 1. All five replications now achieve genuine steady-state behaviour for the entire 864-hour run, with buffer storage occupancy tightly clustered around 27.7-28.9% (table E.12). No replication show the late-run divergence seen at configuration 1 (figure E.13). Loading time is correspondingly tight and low across all replications (1.69-1.90h, compared to the 1.56-41.0h split seen in configuration 1), and outbound dock assignment wait time falls to near zero (0.005-0.037h, compared to 0.15-34.9h previously). These KPIs are both consistent with a system operating with comfortable spare capacity rather than one intermittently overwhelmed.

Having established a configuration that reaches full steady state, the next step tests whether crane capacity can be reduced without reintroducing instability. This step follows from the same expectation raised in configuration 1. Since scenario 3's routing removes the highbay return trip for picked pallets, total crane demand per pallet is inherently lower than in scenario 2, suggesting the current crane fleet may have more capacity than this configuration actually requires.

Configuration 3

The third configuration has one crane less than configuration 2. Reducing crane capacity, with buffer storage held the same, tests whether scenario 3's expected lower crane demand allows for a smaller crane fleet without reintroducing instability. It does: all five replications remain fully stable for the entire measurement window, with buffer storage occupancy tightly clustered around 26.9-29.8% (table E.13), closely matching configuration 2's 27.7-28.9% at 7 cranes (figure E.14). Loading time (1.76-1.99h) and outbound dock assignment wait time (0.016-0.064h) remain low and consistent across all replications, with no sign of the instability seen in configuration 1.

Crane retrieval time increases from 0.14h at configuration 2 to 0.22h at configuration 3, reflecting the reduced spare capacity per crane, though this remains a small fraction of total loading time and has no visible effect on system-level stability. With 6 cranes confirmed sufficient, the next step tests a further reduction in cranes to establish how far this scaling can be pushed before instability returns.

Configuration 4

The fourth, and final, configuration consists of one less crane. Reducing crane capacity further, tests the limit of how far scenario 3's reduced crane demand allows the fleet to be scaled down. The system remains stable: all five replications complete the full measurement window without destabilising, buffer storage occupancy stays tightly clustered around 27.4-28.7% (table E.14). This matches configurations 2 and 3, and loading time remains low (1.65-2.24h) across all replications (figure E.15). Crane utilisation, however, is substantially higher than at any previous configuration: 94.53% on average,

and evenly distributed across the individual cranes (94.07-94.82%), indicating the system is operating close to crane capacity saturation, though without concentrating stress on any single crane and heavy effects on the system.

This near-saturated crane utilisation is seen in several other indicators. Crane retrieval time nearly quadruples relative to configuration 3, from 0.22h to 0.82h, a substantially steeper increase than the step produced (0.14h to 0.22h). Outbound dock assignment wait time also increases, from 0.040h to 0.067h on average, and its standard deviation (0.084h) now exceeds its own mean, indicating a wider and less predictable spread across replications than other configurations. Lift and shuttle utilisations stay stable, indicating that the near-saturated crane condition at this configuration does not reintroduce the instabilities in the broader system.

This configuration therefore demonstrates that this amount of cranes is technically sufficient to maintain steady-state operation under scenario 3's routing, but with substantially less margin than other crane amounts: the system operates close to crane capacity saturation, not merely with a visibly reduced buffer. Which amount of cranes represents the more appropriate recommendation depends on how much weight is given to this high operating point against the resource savings of the smaller fleet, a trade-off discussed further in 7, rather than a question the stability result alone resolves.

6.3.2. Results

Table 6.7 summarises the full set of KPIs for the final scenario 3 configuration. This configuration is selected over the configuration 3 on the basis that the underlying demand data represents peak conditions, average real-world demand would leave more margin than either configuration shows here. The differences between KPIs of scenario 2 are shown and discussed.

Table 6.7: Scenario 3 KPI summary

KPI	Mean	SD	Min	Max	vs. S2
Trucks loaded per dock per hour	0.383	0.006	0.376	0.390	+1.6%
Highbay occupancy	69.67%	0.85 pp	68.99%	71.09%	-0.7%
Truck loading time (arrival–departure)	1.90 h	0.23 h	1.65 h	2.24 h	-3.6%
Total time on bufferlane	0.69 h	0.04 h	0.64 h	0.74 h	-6.8%
Truck punctuality	-0.78 h	0.13 h	-0.92 h	-0.60 h	-20.0%
Planned pickup to actual departure	1.12 h	0.25 h	0.86 h	1.55 h	-15.2%
Crane retrieval time	0.82 h	0.14 h	0.65 h	1.00 h	+78.3%
Crane utilisation	94.53%	0.28 pp	94.07%	94.82%	+13.5%
Buffer storage occupancy	28.20%	0.51 pp	27.38%	28.66%	+4.8%
Buffer storage retrieval time	0.48 h	0.01 h	0.47 h	0.48 h	-2.0%
Buffer storage to bufferlane time	0.12 h	0.01 h	0.11 h	0.13 h	+9.1%
Order ready wait time	0.091 h	0.0013 h	0.090 h	0.093 h	+2.2%
Outbound dock assignment wait time	0.067 h	0.084 h	0.003 h	0.208 h	+28.8%
Inbound dock utilisation	86.90%	–	85.98%	87.86%	-1.4%
Outbound dock utilisation loading	19.38%	–	5.45%	32.60%	+0.6%

Comparing the final configuration of scenario 2 with scenario 3, highlights the trade-offs introduced by scenario 3's routing. Outbound performance is similar or modestly better in scenario 3. Loading time falls by 3.6%, and planned pickup to actual departure improves by 15.2%. However, this is not achieved with less crane demand overall, despite scenario 3 skipping the highbay return trip that scenario 2 requires, crane utilisation is 13.5% higher (94.53% versus 83.30%), and crane retrieval time is 78.3% longer (0.82h versus 0.46h). This apparent contradiction resolves once utilisation is considered along-

side workload: scenario 3 requires less total crane work per pallet, but the smaller crane fleet needed to exploit this results in each individual crane operating closer to saturation. Each retrieval waits longer in queue even though the aggregate demand on the crane fleet is lower. Outbound dock assignment wait time shows a similar pattern, rising 28.8% in absolute percentage terms (0.067h versus 0.052h) despite both figures being operationally negligible in practice.

Buffer storage occupancy is 4.8% higher in scenario 3 (28.20% versus 26.92%), and buffer storage to bufferlane time is 9.1% higher. This is consistent with buffer storage having a more central role in scenario 3's routing than in scenario 2's, where picked pallets stay in the storage for a longer period of time. Highbay occupancy and inbound dock utilisation remain close between the two scenarios (within 1.5% of each other), confirming that the two scenarios diverge specifically in downstream retrieval and buffer-storage behaviour rather than in inbound area processes. Overall, scenario 3 trades a smaller, more heavily utilised crane fleet for comparable or slightly improved outbound outcomes. A genuine resource-efficiency gain in terms of crane count, but one that leaves markedly less operational slack than scenario 2's configuration.

Crane utilisation averages 94.53%, with individual cranes tightly clustered (94.07-94.82%, table 6.7), indicating the near-saturated workload remains evenly distributed rather than concentrating on any single crane. Inbound dock utilisation remains tight across the docks (85.98-87.86%, averaging 86.90%), consistent with every prior scenario and configuration, and is now close to crane utilisation. Outbound dock utilisation is lower and more evenly spread than in scenario 2. This spread continues to occur without any corresponding queueing delay (outbound dock assignment wait time averages only 0.067h).

Figures 6.7 and 6.11 show inbound trucks in the warehouse, outbound trucks in system, crane retrievals per hour, pallets in buffer storage, and pallets in the highbay over a representative day within the measurement window ((8th day), averaged across the 5 replications). Pallets in buffer storage (figure 6.10) show a steady within-day climb from approximately 1,200 pallets at the start of the day to a peak of roughly 1,430 by hour 23, rather than the rhythm found in scenario 2. This is consistent with scenario 3's routing feeding buffer storage more continuously via the direct path rather than through a highbay affected flow. Crane retrievals per hour (figure 6.9) remain active throughout the day with no extended idle period, consistent with cranes now operating at 94.53% utilisation. Outbound trucks in the warehouse (figure 6.8) stay comparatively low throughout, consistent with the low outbound dock utilisation and near-zero dock assignment wait time reported in table 6.7.




Figure 6.7: Inbound trucks in system over a representative day, scenario 3




Figure 6.8: Outbound trucks in system over a representative day, scenario 3




Figure 6.9: Crane retrievals per hour over a representative day, scenario 3




Figure 6.10: Pallets in buffer storage over a representative day, scenario 3




Figure 6.11: Pallets in highbay over a representative day, scenario 3

6.4. Scenario comparison

Table 6.8 compares the final configuration of each scenario across the KPIs measured throughout this research. Two notes apply to this comparison: scenario 1 is validated against real Piacenza 1 operational data, while scenarios 2 and 3 rely on demand projections for the not-yet-built Piacenza 2 facility. This means that the comparisons between scenario 2 and 3 versus scenario 1, compare a validated baseline against internally consistent but unvalidated projections rather than two equally-grounded scenarios. Four KPIs marked with an asterisk measure a structurally different segment of the process in scenario 1 than in scenarios 2 and 3, since scenario 1's layout has no buffer storage stage, so their large deviations partly reflect this difference rather than a purely operational improvement. The results are discussed below alongside each scenario's expected outcomes, before a concluding

summary of the overall findings.

Table 6.8: Scenario comparison, final configurations

KPI	Scenario 1	Scenario 2	Scenario 3	Δ S2 vs. S1	Δ S3 vs. S2	Δ S3 vs. S1
Highbay occupancy	74.57%	70.16%	69.67%	-5.9%	-0.7%	-6.6%
Truck loading time (arrival–departure)	2.56 h	1.97 h	1.90 h	-23.0%	-3.6%	-25.8%
Planned pickup to actual departure	1.84 h	1.32 h	1.12 h	-28.3%	-15.2%	-39.1%
Trucks loaded per dock per hour	0.228	0.377	0.383	+65.4%	+1.6%	+68.0%
Inbound dock utilisation	47.6%	88.15%	86.90%	+85.2%	-1.4%	+82.6%
Outbound dock utilisation	18.9%	19.26%	19.38%	+1.9%	+0.6%	+2.5%
Crane utilisation*	45.6%	83.30%	94.53%	+82.7%	+13.5%	+107.3%
Total time on bufferlane*	3.40 h	0.74 h	0.69 h	-78.2%	-6.8%	-79.7%
Crane retrieval time*	4.43 h	0.46 h	0.82 h	-89.6%	+78.3%	-81.5%
Order ready wait time*	3.19 h	0.089 h	0.091 h	-97.2%	+2.2%	-97.1%

Two main patterns emerge from the table as a whole, before diving into the individual hypotheses discussed below. First, crane utilisation rises steadily across the full progression from scenario 1 to scenario 3 (+107.3% overall), even as crane retrieval time itself falls. This reflects that both buffer-storage scenarios were dimensioned with smaller crane fleets than scenario 1, given the crane workload reduction the buffer storage design allows. Utilisation rises because the fleet was actively reduced to match that lower workload, not because the design itself increases crane pressure. This should also be read in light of how this research's capacity-dimensioning process was conducted: configurations were searched down toward the minimum crane fleet that remained stable, closer to a true operating limit than a design decision made in practice would typically allow. A real-life decision would more likely retain deliberate operating margin. The resulting crane utilisation figures should therefore be read as a demonstration of how far the design can be pushed, not necessarily as what a practical implementation would choose to run at.

Second, the buffer storage design's central benefit, decoupling retrieval from truck arrival timing, shows up almost entirely in the transition from scenario 1 to scenario 2. With order ready wait time and total time on bufferlane both falling by more than 75% at that step, while the further transition from scenario 2 to scenario 3 changes these same KPIs comparatively little. The routing change tested in scenario 3 therefore appears to be a refinement of an already-realised benefit, not an independent source of the same kind of improvement. Compared to scenario 2, scenario 3 falls even further in loading time, and trucks loaded per dock per hour increases even more. Loading time falls a further 3.6% and trucks loaded per dock per hour rises 1.6%, a modest but genuine improvement on top of the benefit already secured by scenario 2's buffer storage design.

6.4.1. Expected outcomes

Scenario 2

In table 6.9, the expected outcomes of scenario 2 that were discussed in chapter 4 are assessed and compared to the actual simulation results.

Table 6.9: Scenario 2 expected outcomes

Aspect	Actual result	Verdict
Bufferlane occupancy	-78.2%	Confirmed
Highbay occupancy	-5.9%	Confirmed
Loading time	-23.0%	Refuted
Truck arrival-to-departure	-23.0%	Confirmed
Order activation to retrieval	-89.6%	Confirmed

Loading time and truck arrival-to-departure time both refer to total loading time as defined in this re-

search (actual arrival to departure), predicted smaller in scenario 2 compared to scenario 1. This is confirmed by a decrease of -23%. Order activation to retrieval time is predicted lower. The actual result confirms the matter, a large reduction of -89.6%. Bufferlane occupancy and highbay occupancy are both confirmed as hypothesised: pallets spend significantly less time occupying the bufferlane before departure once retrieval is triggered by actual truck arrival rather than planned pickup time, and highbay occupancy falls as pallets are now divided between the highbay and buffer storage rather than held in the highbay alone. The main goal of releasing pressure on bufferlanes is achieved by this.

Scenario 3

In table 6.10, the expected outcomes of scenario 3 that were discussed in chapter 4 are assessed and compared to the actual simulation results.

Table 6.10: Scenario 3 expected outcomes

Aspect	Actual result	Verdict
Crane utilisation	+13.5%	Refuted
Order activation to retrieval	+78.3%	Refuted
Buffer storage occupancy	+4.8%	Confirmed
Highbay occupancy	-0.7%	Confirmed
Loading time	-3.6%	Confirmed
Trucks loaded per dock per hour	+1.6%	Confirmed

Crane utilisation rises rather than falls as hypothesised, since the crane fleet was reduced next to the routing change to match scenario 3's lower total crane workload per pallet. This leaves the smaller fleet more heavily utilised despite requiring less total work. Buffer storage occupancy, highbay occupancy, and trucks loaded per dock per hour are all confirmed as hypothesised, though the highbay occupancy change is negligible in magnitude. Order activation to retrieval time, pointing at the queues of cranes, is predicted as lower than scenario 2. The actual result shows a substantial increase (+78.3%), refuting both predictions. This can be explained by the crane utilisation. Loading time is predicted as slightly lower. The actual result (-3.6%) confirms this.

6.5. Sensitivity analysis

A one-factor-at-a-time (OFAT) sensitivity analysis was conducted on three input parameters: pick share, bulk retrieval activation time, and the distribution scale of manual loading time. Each were varied by $\pm 25\%$ from its nominal value, using the final scenario 3 configuration. The adjusted values are not shown in the public version.

A higher pick share is expected to decrease crane retrieval times, since pick pallets have an earlier retrieval time from bulk pallets. This will spread out crane retrievals more per order. Bulk retrieval activation time was expected to influence buffer storage occupancy directly, a shorter lead time should reduce the time pallets spend in buffer storage before departure. Therefore the occupancy of the buffer storage decreases, and the highbay occupancy increases as pallets stay there for a longer period of time. Loading time scale was expected to impact total loading time. A lower scale, and therefore loading time per pallet, decreases total loading time. Although this is depending on the speed of retrieval of an order from the buffer storage.

6.5.1. Results

Table 6.11 shows the percentage deviation from the baseline for each KPI under each tested condition. 'Bulk retrieval high' is deleted from this table. The results of this run were identical to the nominal baseline across every KPI and all 5 replications to several decimal places. This is a level of agreement inconsistent with two independently executed stochastic simulations, even under a genuine null effect, which would still be expected to show replication-level noise. Bulk retrieval activation time's sensitivity is therefore characterised only in the low (-25%) direction in this analysis.

Table 6.11: Sensitivity analysis scenario 3

KPI	Pick share (±25%)		Bulk retrieval (−25%)	Loading time scale (±25%)	
	Low	High	Low	Low	High
Total loading time	-6.9%	+1.5%	+1.3%	+28.3%	-11.1%
Planned pickup to actual departure	-7.9%	+7.9%	-0.4%	+42.0%	-19.6%
Trucks loaded per dock per hour	-1.7%	+1.7%	+0.7%	-0.3%	-1.2%
Highbay occupancy	+0.4%	-0.9%	-0.3%	-0.5%	-0.9%
Buffer storage occupancy	-0.4%	+1.1%	+1.1%	+1.2%	-2.8%
Crane retrieval time	-9.1%	+15.7%	-0.8%	-10.4%	-0.4%
Buffer storage retrieval time	+3.8%	-5.1%	-2.5%	+2.6%	-4.9%
Buffer storage to bufferlane time	-3.3%	+2.6%	+2.0%	+143.1%	-44.8%
Order ready wait time	+1.0%	-0.3%	-0.4%	+44.3%	-28.4%
Crane utilisation	-1.0%	+0.1%	+0.3%	-1.2%	-1.3%

Crane utilisation remains essentially constant across every tested condition (93.3-94.8%, compared to 94.53% at nominal), confirming that none of the three tested parameters meaningfully affects crane workload, since all three govern pick-, loading-, or retrieval-timing behaviour rather than crane demand directly. This also confirms that the system operates close to crane saturation consistently across the tested parameter space, not only at the nominal configuration.

The pick share was varied by $\pm 25\%$. This produces moderate, broadly symmetric effects across most KPIs, with deviations generally within $\pm 10\%$ and no single KPI standing out. The pick share's hypothesis is refuted, because crane retrieval time was expected to decrease under a higher pick share. The data shows the opposite: crane retrieval time rises 15.7% under Pick high and falls 9.1% under Pick low. A higher pick share is associated with longer, not shorter, crane retrieval times. The effects of a low or high pick share are similar in margin on the KPIs.

The bulk retrieval activation time was varied by $\pm 25\%$. Its measured effects are moderate, within $\pm 10\%$ across most KPIs. Its hypothesis is inconclusive rather than refuted. A shorter lead time (Bulk low) was expected to reduce buffer storage occupancy and correspondingly increase highbay occupancy. The observed deviations are both negligible and in the opposite direction to what was expected (buffer storage occupancy +1.1%, highbay occupancy -0.3%), but their magnitude falls within normal replication variance. The missing high-lead-time comparison point further limits what can be concluded about this parameter's true effect on buffer storage sizing.

The loading time scale was varied by $\pm 25\%$. The results in both directions show different sizes where buffer storage to bufferlane time is 3 times the size for low loading time scale than high loading time scale. This results in higher impact of changing to a lower loading time scale, than a high loading time scale. The scale produces a substantially larger effect than either of the other two parameters: total loading time rises 28.3% under the low-scale condition but falls only 11.1% under the high-scale condition. This is the inverse of the direction implied, since a lower scale value corresponds to shorter individual loading operations. This inversion is most pronounced in two downstream metrics: buffer storage to bufferlane time rises 143.1% under low scale and falls 44.8% under high scale, and order ready wait time shows a similar pattern (+44.3%/-28.4%). Loading time scale's hypothesis is refuted: total loading time was expected to fall under a lower scale value, reflecting shorter individual per-pallet loading times. This is consistent with what was already anticipated in the hypothesis, namely that the effect depends on the speed of order retrieval from buffer storage. This could be explained by an incorrect implementation of the parameters in opposite direction. Even more, a faster individual loading rate increases the rate at which docks turn over and demand new pallets from the buffer storage, so time saved on loading could be outweighed by additional queueing time upstream.

The loading time scale is therefore identified as the most sensitive of the three tested parameters, both in magnitude and in the unpredictability of its effect direction, though part of this larger effect is attributable to its larger absolute step size, not only to the system's underlying sensitivity to it. This distinguishes it from pick share and bulk retrieval activation time only in scale and clarity of mechanism, not in reliability: none of the three parameters can be said to behave in a fully understood or fully

validated way based on the present analysis, and loading time scale simply shows the largest and most consequential departure from expectation.

6.6. Economic analysis

With a new technology implementation in a warehouse, there are also economic consequences. An AS/RS system with a 2D shuttle has multiple cost drivers, including racking, cranes, lifts, shuttles, charging systems, and more. Based on internal supplier data, an estimation comparing the traditional bufferlane design and the buffer storage design can be made. The supplier data used in this analysis was compiled from internal NewCold conversations and past project experience. This comparison is made between the optimal scenario 3 configuration, a buffer storage design by NewCold, and a traditional-bufferlane design for Piacenza 2.

CAPEX (capital expenditure) refers to the one-time upfront investment required to acquire and install the physical equipment for a given warehouse design, in this case everything associated with: cranes, lifts, shuttles, docks, and associated other buffer storage equipment. The additional equipment related to buffer storage is scaled by the number of lifts. Operating expenditure (OPEX) covers the operational costs of the facility (e.g., maintenance, labour, and energy). This analysis compares CAPEX only; OPEX differences between the two designs are discussed qualitatively below.

Table 6.12 shows the equipment across three design configurations for Piacenza 2. The costs are left out in the public version. The traditional design (corrected) column reflects NewCold's standard bufferlane-based design, adjusted to include the additional inbound dock this research identifies as necessary to handle expected inbound arrivals. The buffer storage design column reflects the supplier's own proposed shuttle-based configuration for Piacenza 2's given demand and throughput, prior to this correction. The buffer storage design (corrected) column is the final configuration of the simulated warehouse model developed in this research, with the same inbound dock correction applied for consistency.

Table 6.12: Equipment and cost comparison across design configurations

	Traditional design (corrected)	Buffer storage design	Buffer storage design (corrected)
Cranes	-	-2	-2
Inbound docks	-	-1	-
Outbound docks	-	+4	+4
Pickfloor	1.5	1	1
Lift	—	[not disclosed]	+1 from buffer storage design
Shuttles	—	[not disclosed]	+11 from buffer storage design
Additional buffer related equipment	—	[not disclosed]	- 1
vs. Traditional (%)	—	-13.84%	-6.14%
vs. Buffer storage (%)	—	—	+8.94%

Based on internal NewCold cost data regarding the Piacenza 2 design, the shuttle-based configuration shows a CAPEX approximately lower than the corrected conventional design with bufferlanes. This uses the corrected inbound-dock layout as the baseline rather than NewCold's original standard configuration; comparing against the uncorrected baseline would understate the shuttle design's relative cost advantage, since the original configuration itself under-provisions the inbound dock capacity this research identifies as required.

It is worth noting that this thesis's own configuration process specifies a larger shuttle fleet than the supplier's own proposed shuttle-based design. This may reflect the more rigorous, steady-state-tested approach used in this study relative to the supplier's estimate. The only equipment count that decreases relative to the supplier's proposed design is lifts, this is outweighed by additional shuttles. Although the total costs of the traditional and buffer storage designs are broadly comparable once corrected, this comparison carries enough uncertainty that it should not be read as a recommendation.

Operating cost differences (maintenance, labour) are not quantified but would plausibly reinforce, rather than offset, the shuttle design's cost advantage. A more maintenance-intensive system is expected for the buffer storage design, given its greater number of electronic components and moving vehicles (shuttles and lifts) compared to the comparatively simple conveyor-based traditional bufferlane system. Decoupling pallet retrieval from planned truck arrival time substantially reduces the system's dependence on truck-arrival forecasting accuracy. This is a reduction that translates into less operational effort spent managing exceptions caused by early, late, or mistimed truck arrivals. From a managerial perspective, this represents a genuine reduction in day-to-day operational handling. However, a design with less spare crane capacity offers less resilience to disruption, than a design with more comfortable operating margin. Finally, the buffer storage design's greater reliance on automated, electronically controlled equipment (shuttles and lifts, in place of largely mechanical conveyor-based bufferlanes) implies a different maintenance profile and a larger number of individual components with the potential to fail. This is relative to the traditional design's comparatively simple failure modes. NewCold should weigh the buffer storage design's lower upfront cost against these longer-term operational and risk considerations before finalizing a recommendation for Piacenza 2.

7

Discussion

7.1. Key findings

The results support the central question of this research: a shuttle-based buffer can function as a decoupling system between highbay retrieval and the outbound dispatch area. Although this comes at the cost of substantially higher crane utilisation once the crane fleet is dimensioned down to match the buffer design's reduced workload. It cannot be directly concluded that a shuttle-based buffer is the best alternative in every context, as this trade-off, and its acceptability, varies with the resource cost a warehouse is willing to accept. Several findings deserve particular attention and are discussed below.

First, the pick-pallet routing change in scenario 3, where picked pallets are sent directly into the shuttle buffer rather than back through the highbay, reduces the total crane work required per pallet. However, this does not translate into lower crane utilisation once the corresponding fleet reduction is accounted for. This risks more impact of errors or maintenance to the cranes, but also point at efficiency. Scenario 3's final configuration operates at 94.53% crane utilisation, compared to 83.30% in scenario 2's final configuration (Table 6.8), despite scenario 3 requiring substantially less total crane work per pallet by skipping the highbay return trip. This is because the crane fleet was reduced alongside the routing change, in line with the reduced total workload. The benefit of the routing change therefore surfaces as an opportunity to operate with fewer cranes, not as lower utilisation of the cranes that remain. Crane retrieval time reflects this trade-off too, rising from 0.46h (scenario 2) to 0.82h (scenario 3), consistent with a smaller crane fleet operating closer to saturation. This remains a small fraction of total loading time (1.90h) and has no visible effect on system-level stability at this configuration, pointing at an efficient system. Outbound throughput is essentially unaffected by the routing adjustment (0.377 vs. 0.383 trucks loaded per dock per hour), and total loading time is modestly better in Scenario 3 (1.90h vs. 1.97h). This is a valuable finding independent of whether it changes throughput, since it demonstrates that the buffer-routing benefit in this research shows up as a resource-efficiency trade-off: fewer cranes are needed for comparable output, at the cost of one additional buffer storage level, rather than as reduced pressure on the cranes that are actually deployed. Scenario 3 distributes the resource capacity decreased from cranes to increase buffer storage capacity, rather than reducing total resource requirements.

Secondly, the order ready wait time, which is the delay absorbed by fixed-lane staging to account for arrival uncertainty and dock capacity, fell from 3.19 hours in scenario 1 to approximately 0.09 hours in scenarios 2 and 3 (table 6.8). This is explained by the fact that pallets in Scenarios 2 and 3 move to the bufferlane only once the truck has actually arrived, rather than hours ahead of a planned time that may not hold. This is a direct, quantified confirmation of the buffer design's core purpose. Decoupling retrieval from bufferlane staging substantially reduces the outbound system's dependence on the accuracy of truck arrival forecasts. This is also explained by planned pickup to actual departure, which falls by 28.3% in scenario 2 and by a further 15.2% in scenario 3 (39.1% relative to Scenario 1 overall). Actual truck departure lies substantially closer to the planned pickup time once the buffer storage design is implemented.

Against these two results, total loading time improved from scenario 1 (2.56h) to scenario 2 (1.97h) and scenario 3 (1.90h) in a way that is not fully explained by the results discussed above. The loading time is numerically lower in both buffer storage scenarios, but this reduction is smaller than might be expected given the near-elimination of order ready wait time. Total loading time is measured from actual truck arrival to actual truck departure, and is therefore also affected by the processing time of the loading activity itself, a component the buffer design does not directly address and depends on manual labour and resources.

The instability observed in the intermediate configurations during the scenario 2 and scenario 3 capacity-dimensioning process is a useful negative result in its own right. It shows that a shuttle-based buffer storage does not absorb outbound uncertainty regardless of how it is sized. Sizing decisions shift load onto the crane fleet in a way that must be explicitly dimensioned for, and the final scenario 3 configuration's high crane utilisation (94.53%) demonstrates that this risk persists even in the recommended configuration, not only in the deliberately undersized intermediate steps tested during dimensioning. This is consistent with (Wang et al., 2015) and (Carlo & Vis, 2012), who find that shuttle-system performance depends strongly on coordination between system components and control policies rather than the shuttle mechanism in isolation. This research extends that principle to a dispatch-buffering context specifically, where the critical coordination point is upstream crane capacity rather than internal shuttle-lift scheduling.

7.2. Implications

7.2.1. Academic contribution

Chapter 2 identified a knowledge gap along four related aspects. First, warehouse subcomponents as design, storage assignment, and order picking, are generally studied in isolation rather than as an integrated system (J. Gu et al., 2007; Rouwenhorst et al., 2000). AS/RS research often analyses these systems as self-contained long-term storage, without explicit consideration of interaction with outbound staging or dock operations (Roodbergen & Vis, 2009a; Wang et al., 2015). Second, cold-storage requirements and dispatch-oriented applications are rarely addressed in this literature, which positions AS/RS predominantly as a technology for storage and retrieval in high-bay or reserve areas (Kriehn et al., 2018; Roodbergen & Vis, 2009a; Wu et al., 2020) rather than as a solution for outbound staging with uncertain truck arrival times. Third, there is a lack of empirical case studies that apply real-life operational data to warehouse design and optimisation (J. Gu et al., 2007). Fourth, the technology choice itself is rarely framed as a managerial trade-off between operational flexibility and investment, with existing research evaluating shuttle-based AS/RS almost exclusively on operational performance (Yener & Yazgan, 2019). This research addresses these aspects, applying a validated, real-data discrete-event simulation to a shuttle-based system used specifically as an outbound dispatch buffer in a cold-storage context, and pairing the operational simulation results with a broader economic comparison that makes this managerial trade-off explicit rather than implicit.

Several findings from this research can agree and complement to claims in the literature reviewed in chapter 2. Wang et al. (2015) and Carlo and Vis (2012) find that shuttle-system performance is strongly influenced by coordination between system components. Specifically coordination between parallel shuttle operations and sequential lift movements internal to the shuttle system itself are influencing performance. This research's results confirm that this internal coordination dimension remains relevant in a dispatch-buffering role. The capacity-dimensioning process in chapter 6 shows that buffer storage instability could only be resolved by adjusting the number of shuttle levels. This research's central finding, however, is that crane utilisation, not the buffer storage mechanism itself, is the resource most sensitive to the routing and capacity choices tested. This is a second, external coordination dependency not characterised in studies that treat the shuttle system as self-contained. Buffer-internal coordination and crane-buffer coordination are therefore both necessary, but neither is sufficient on its own to explain this research's results. An outbound buffer's performance depends on both dimensions simultaneously, extending the internal-coordination finding of Wang et al. (2015) and Carlo and Vis (2012) to include this external coordination element.

Battarra et al. (2024) propose that shuttle-based AS/RS performance cannot be optimised by considering any one layout dimension (racks, material handling vehicles, or buffer areas) in isolation from the others. This research's results are consistent with that principle but also refine it: of the three inter-

nal dimensions tested in this research's capacity-dimensioning process, shuttle-level capacity proved the binding constraint, while lift capacity did not, remaining at low utilisation throughout every tested configuration. This diverges from Lerher et al. (n.d.), who identify lifts as the typical bottleneck determining a shuttle-based system's throughput ceiling. In this research, that role is instead taken by the upstream crane resource supplying the buffer from the highbay, a resource entirely outside the shuttle system's own three layout dimensions. This suggests that the principle of joint optimisation of Battarra et al. (2024), may need to extend beyond the shuttle system's internal boundary in a dispatch-buffering context specifically, to include the external resource feeding the buffer.

This research's finding that scenario 3's routing change reduces required crane capacity while maintaining comparable outbound performance directly confirms Ekren et al. (2015), who find that appropriately combined warehouse design choices can improve system performance while simultaneously reducing resource requirements, rather than only improving a single performance metric in isolation. It also confirms the bottleneck-shifting behaviour reported by Kriehn et al. (2018), who find that storage policy changes shift where a shuttle-based system's bottleneck lies rather than removing it. In this research, the routing change tested in scenario 3 does not remove the crane-buffer trade-off but shifts its expression, from crane retrieval time in scenario 2 to crane utilisation approaching saturation in scenario 3's smaller, more heavily-loaded fleet.

Methodologically, this research also demonstrates that combining a validated real-data baseline (scenario 1) with structurally distinct design scenarios (scenarios 2 and 3) built on projected future demand is a viable way to evaluate a not-yet-built facility's design choices. It should be made explicit about which comparisons are validated against real operations and which are internally consistent projections only. This directly addresses the disconnection between academic research and real-life case data noted by J. Gu et al. (2007), and offers a concrete methodological template for evaluating design choices at facilities where full operational validation is not yet possible.

Taken together, this research contributes to the literature in four specific ways: it applies shuttle-based storage and retrieval system methodology to an outbound dispatch-buffering role in a cold-storage context. This directly addresses the gap identified in chapter 2, it identifies upstream crane capacity, rather than an internal shuttle-system component, as the resource most clearly affected by dispatch-buffer design choices. And extends existing bottleneck and coordination findings beyond the shuttle system's own boundaries, it demonstrates a resource-reallocation pattern consistent with, and extending, existing findings on joint resource optimisation in shuttle-based systems. It also provides a methodological template, for evaluating a not-yet-built facility's design choices against a validated real-data baseline, applying a case-study and managerial perspective.

7.2.2. Practical implications

For NewCold's Piacenza 2 design, the results support adopting scenario 3's routing logic, in which picked pallets move directly into the shuttle buffer rather than back through the highbay. This routing change achieves comparable or slightly improved outbound performance (loading time, planned pickup to actual departure, throughput) relative to scenario 2, while requiring fewer cranes to do so, making it the more resource-efficient of the two buffer-based designs tested in this research.

This efficiency gain comes with a real risk that NewCold should account for explicitly. Any implementation should be dimensioned with enough crane capacity to stay clear of the saturation risk demonstrated in this research's undersized configurations, and can occur with crane failures or maintenance too. A buffer storage system redistributes rather than removes retrieval load. The final scenario 3 configuration itself operates close to crane saturation (94.53% utilisation), indicating this risk is not merely theoretical but present even in the recommended design. This figure reflects peak-demand conditions specifically, so average operational conditions are expected to be less intensive than this. The peak figure remains the appropriate basis for a sizing decision, however, since a design sized only for average demand would be under-provisioned during the periods that matter most for stability. Sizing the buffer without sizing the crane fleet to match would risk the instability seen in this research's initial configurations. Given this, NewCold should treat crane capacity as a primary design parameter for Piacenza 2, sized against peak rather than average demand. The economic analysis in chapter 6 shows the buffer storage design's CAPEX to be lower than the corrected traditional design. This analysis should be elaborated on further, and operational consequences must be translated into monetary terms, to

give a complete recommendation.

Beyond the specific design recommendation, the simulation model and validated dataset developed in this research are themselves a practical asset for NewCold independent of the Piacenza 2 decision. The model was built and validated against real Piacenza 1 operational data, and its structure, model inputs, and KPI framework are not specific to the parameters tested here. NewCold can reuse this model to evaluate further design variants beyond the three scenarios tested in this research, to test sensitivity to updated design specifications, or to support design decisions at other automated cold-storage facilities with a similar dispatch-buffering need. This increases the validation basis for future design decisions and reduces the marginal cost of future design evaluation, since the validation effort already completed in this research does not need to be repeated for related design questions.

More broadly, the results suggest that automated cold-storage warehouses considering a shuttle-based dispatch buffer should evaluate it primarily as a tool for freeing crane capacity and reducing dependence on truck-arrival forecasting accuracy. This research's results do not support a proportional reduction in total truck dwell time as a significant additional benefit: total loading time improved only modestly despite order ready wait time falling by over 95%, consistent with manual per-pallet loading time remaining fixed across scenarios and therefore setting a floor on how much total dwell time can be reduced. This has two implications for the wider industry. First, the crane-buffer trade-off identified in this research is a design consideration that is easy to overlook when a shuttle buffer is evaluated only on outbound metrics. Warehouse operators evaluating similar designs should explicitly track crane utilisation alongside throughput and metrics as arrival-to-departure times and loading time, not treat it as a secondary check performed only if problems arise. Second, this research's approach, combining a real-data-validated baseline with simulated design variants for a not-yet-built facility, offers other automated warehouse operators a way to evaluate buffer designs before construction, without requiring an existing design in order to validate the model that evaluates it.

7.3. Limitations

Within this research, several limitations regarding the simulation model and analysis are to be made. To begin with, scenario 1 is validated against real operational data, while scenarios 2 and 3 rely on demand estimates derived from the company's projections for the not-yet-built Piacenza 2 facility. These demand estimates were translated into interarrival rate tables to mirror the hourly behaviour of the warehouse. The estimates consist of a single peak-day provided by NewCold, rather than a fuller characterisation of expected demand variability. A more detailed estimation of expected interarrival times for inbound trucks and orders, ideally incorporating multiple representative days, would strengthen the credibility of scenarios 2 and 3. Throughput figures for scenarios 2 and 3 should therefore be read as internally consistent with each other, not as fully validated predictions of real Piacenza 2 performance, and are not directly comparable to the scenario 1 baseline. The estimations come with uncertainties about customers and demand, and therefore lack precise data for a detailed design. However, lots of takeaways from Piacenza 1 and other warehouses can be taken into account. Secondly, not all KPIs from the scenario 1 results had a validation comparison with real data. Crane and dock utilisation rates in particular were not compared to real data. Next to that, scenario 1's warm-up period was also determined differently from scenarios 2 and 3: it was set at the point where order generation began, without the additional stabilisation margin of approximately one day added beyond this threshold used in scenario 2 and 3. This introduces a minor inconsistency in warm-up methodology across scenarios that was not resolved within this research.

The arrival rate of inbound trucks also caused congested routes to the I-point in all scenarios. Inbound flow was not the main focus of this research, but this behaviour raises questions about how unloading processing times were calculated. Inbound dock utilisation in particular was out of this research's scope but did not behave as expected (way above proposed NewCold design for Piacenza 2), and is flagged here rather than investigated further. Related to this, unloading time at the inbound docks was also uncertain due to a low sample size of the calculations. No better estimation could be established within the scope of this research. Also, the outbound dock utilisation observed across all three scenarios suggests the outbound docks tested in this research may be larger than the outbound flow requires under the demand modelled.

A warm-up period misconfigured in hours rather than days, discovered later in the analysis, affected several resource-utilisation statistics (crane, shuttle, and lift utilisation) across the scenario 3 capacity-dimensioning processes. Unlike buffer storage occupancy, highbay occupancy, and throughput, which could be recomputed directly from raw simulation output and are correctly stated throughout this thesis, these utilisation statistics have no equivalent raw data source and could not be corrected. Time constraints allowed only the final configuration of the scenario to be rerun and independently verified.

Time constraints limited the scope of the experimental design in several ways. Most scenarios were run with only 5 replications per configuration. This limits the statistical confidence that can be placed on the reported means and reduces the ability to distinguish genuine effects from run-to-run variation, particularly for KPIs that already showed a wide spread across replications. The measurement window itself was also constrained: each replication covered a 14-day period following warm-up. Except for scenario 1, which used a 31-day runtime to allow direct comparison against a full month of real operational data. A 14-day window captures only a limited number of low-frequency but operationally significant events, such as extreme truck delays or buffer near-capacity moments, so rarer tail behaviour may be under- or over-represented in the reported KPIs. Computational complexity also played a role, with some experiments taking up to 20 hours to run. The consistency of replication behaviour, reflected in the narrow ranges of most reported KPIs, limits the practical impact of this limitation.

Task assignment within the model was also kept deliberately simple. Highbay crane assignment is random across available cranes, with no policy directing retrievals toward, for example, the least-loaded or nearest crane. Buffer storage assignment places all pallets of a given order on a single row, spread across the available levels. But the assignment does not optimise placement for retrieval efficiency in another way. Lift assignment for moving pallets from buffer storage to the bufferlanes is similarly random, with no policy directing which lift serves which pallet based on travel distance, current load, or bufferlane destination. This can introduce congestions on the conveyors to the bufferlanes, and can be limited by a policy. All three assignment policies could plausibly be improved with a more deliberate strategy, and doing so may reveal additional performance improvements not captured in this research.

Several further operational aspects were simplified or excluded from the scope of this research. Picking was modelled as a standard fixed process, pallets requiring picking proceed directly to the pickfloor, without accounting for employee shift schedules or picking capacity constraints that would apply in practice. Reshuffling, whereby a stored pallet is relocated within storage for efficiency purposes, is not modelled for cranes. In practice, reshuffle movements account for roughly 3.5% of total warehouse movements, representing additional crane workload not reflected in this research's crane-utilisation figures. Charging of shuttles and lifts is not modelled, and could constrain buffer storage performance in practice if vehicles are periodically unavailable for charging during operational hours. The model also assumes an unlimited number of employees and forklifts are available wherever needed; in practice, labour and equipment availability represent a real operational constraint not captured here. This research evaluated system performance in operational terms only, the investment and operational cost of a shuttle-based buffer relative to conventional buffer lanes was not incorporated into the simulation model itself.

A manual one-factor-at-a-time (OFAT) sensitivity analysis was conducted instead, varying pick share, bulk retrieval activation time, and loading time scale individually. However, this analysis was not extensive, covering three parameters at a single variation each. And one condition (bulk retrieval activation time, high) produced results identical to the nominal baseline, leaving that parameter's sensitivity characterised only in one direction.

7.4. Future research


Multiple future research recommendations follow directly from the limitations above. Firstly, in scenarios 2 and 3, it is recommended to optimise the sizing of the inbound and outbound docks, and the highbay and buffer storage, rather than only the resources acting on them. For example, the aisle size or level size can be adjusted in the buffer storage. Next to that, in both buffer storage scenarios, the highbay retrieval activation times are set with a maximum deviation of 2 hours. Future research could test wider activation time spreads and their effect on crane and buffer storage performance. Another interesting addition to the scenarios laid out in this research is the exploration of additional flows for the shuttle-

based dispatch buffer, such as inbound overflows and incomplete pick pallets, which could possibly reduce stress on warehouse operations further.

The crane, shuttle, and lift utilisation figures for scenario 3's intermediate configurations, affected by the warm-up misconfiguration described in the limitations above, should be re-verified with a corrected warm-up period. Only the final configuration was independently confirmed within this research's available time. Rerunning would allow the full capacity-dimensioning journey, not only its endpoint, to be reported with confidence. The sensitivity analysis conducted in this research should also be extended further. The approach used proved workable and could be extended to a wider set of parameters, a wider range of variation per parameter, and more replications per condition. The one condition that produced inconsistent results in this research (bulk retrieval activation time, high) should specifically be rerun and diagnosed. Doing so would provide a more complete and reliable ranking of which operational uncertainties (truck arrival timing, retrieval activation timing, pick share, and others) the shuttle buffer design is most sensitive to.

The task assignment policies kept deliberately simple in this research, specifically random highbay crane assignment, semi-random buffer storage placement, and random lift assignment to the bufferlane, could be replaced with deliberate assignment strategies to establish whether additional performance headroom exists beyond what this research's simplified policies allow.

Beyond these, incorporating investment and operating cost directly into the simulation-based comparison, rather than as the separate general analysis presented in this research, is recommended, along with validating the Piacenza 2 demand model once further operational details are available; both would strengthen the practical applicability of the findings. Running experiments for a longer simulated time and with more replications will also improve result reliability. With these suggestions, this research is a good starting point for further research into the performance effects of implementing a shuttle-based buffer in the dispatch area of warehouses.



Conclusion

This research investigated whether a shuttle-based dispatch buffer can be designed and controlled to minimise crane congestion and outbound delays in an automated cold-storage warehouse, using discrete-event simulation of NewCold's Piacenza facility. The main research question is answered by the following five sub-questions.

SQ1: What causes outbound delays in the dispatch area of automated cold storage warehouses, and which design alternatives exist to address these causes?

The literature review, expert consultation, and desk research identified truck arrival uncertainty, variable order compositions, and activity peaks for cranes as the dominant sources of outbound variability. Existing fixed buffer lanes are structurally unable to absorb this uncertainty, resulting in manual real-time exception handling and delays in daily operations. Literature on automation and automated storage- and retrieval systems (AS/RS) is extensive, but focuses mainly on ambient warehouses and treats individual warehouse processes largely in isolation from one another. Among the AS/RS design alternatives reviewed, shuttle-based storage and retrieval systems (SBS/RS) were identified as the most suitable technology for this outbound buffering role, given their fast-retrieval, small-footprint, and flexible-release characteristics relative to alternatives. However, the application of these AS/RS systems, particularly shuttle-based ones, is rarely studied in the context of a warehouse's outbound staging area specifically. Shuttle-based AS/RS systems are most suitable for the automated and cold storage environment.

SQ2: What are the process flows and system characteristics of the Piacenza warehouse that determine outbound performance?

The flowchart and IDEF0 diagrams characterised the Piacenza warehouse's inbound, storage, and outbound flows, and quantified key operational parameters. The characteristics and processes defined here with expert consultation and desk research, combined with the facility's actual layouts, formed the operational baseline for the simulation model. Uncertainties and assumptions within the general warehouse flow and its processes were also identified through desk research and expert consultation. The same approach as for the existing Piacenza warehouse was applied to design and structure the scenarios for the new Piacenza warehouse. Expected results for scenario 2 and 3 were included as well as benchmark against which results discussed later were compared.

SQ3: Which system parameters, constraints, and key performance indicators (KPIs) are relevant for modelling and evaluating the dispatch buffer area of an automated, cold-storage warehouse?

For both Piacenza warehouses, validation and performance metrics were researched. Most KPIs overlap between the two warehouses (and, respectively, scenario 1, and scenarios 2 and 3), with the exception of buffer-storage-related metrics, which apply only to scenario 2 and 3. The shared KPI framework centred on truck loading time, punctuality, bufferlane fill time, throughput, occupancy, and other crucial transport and activity times. Resource utilisation rates for cranes and docks were included across all scenarios, while shuttle and lift utilisation were added specifically for the buffer storage resources in

scenarios 2 and 3. All KPIs were drawn from literature, internal data analysis, and expert consultation, to identify parameters and constraints relevant to the model's inputs and process logic.

SQ4: How can the context of a shuttle-based dispatch buffer be modelled?

A discrete-event simulation model was developed in Simio across three scenarios, isolating the effect of the buffer design through step-by-step configuration changes. Scenario 1 is a validated baseline reflecting Piacenza 1's current traditional-bufferlane layout, scenario 2 introducing a shuttle-based buffer for ordered pallets, and scenario 3 extending this to route picked pallets directly into the buffer rather than back through the highbay. The model was validated against real Piacenza 1 operational data, reproducing trucks loaded per dock per hour almost exactly, alongside plausible truck-count and occupancy figures. Not all KPIs (notably crane and dock utilisation) had a direct real-data comparison available. All final configurations are analysed in terms of steady-state achievement, and the individual process-and element behaviour over a representative day.

SQ5: How does the shuttle based dispatch buffer perform compared to traditional buffer lanes under different operational conditions?

Comparing the shuttle-based buffer against the traditional bufferlane layout showed that the buffer's core mechanism, decoupling retrieval and its resulting uncertainties from bufferlane staging, works as intended. Order ready wait time fell from 3.19 hours in the validated baseline to approximately 0.09 hours with the buffer storage design, confirming substantially reduced dependence on truck-arrival accuracy. Routing picked pallets directly into the shuttle buffer, rather than returning them to the highbay first, reduced the total crane work required per pallet and allowed the crane fleet to be reduced for comparable truck-facing performance. However, this came at the cost of the remaining cranes operating at higher utilisation (94.53%, close to saturation, versus 83.30%) and the required extra level in the buffer storage, compared to scenario 2. Total truck loading time improved relative to the traditional-bufferlane baseline (2.56h, versus 1.97h and 1.90h for the two buffer-based designs), though by less than the near-elimination of order ready wait time might suggest. An under-resourced buffer configuration was also shown to destabilise the system rather than absorb demand. Which confirmed that a buffer storage does not passively absorb outbound uncertainty regardless of how it is sized. A manual sensitivity analysis of loading time, retrieval activation timing, and pick share confirmed the final configuration's robustness to pick share specifically. Loading time scale was identified as a more unpredictable driver of system behaviour. A directional economic comparison further indicated that the buffer storage design carries a lower capital expenditure than the corrected traditional bufferlane design. This suggests the operational trade-offs identified above are not impacted by a higher upfront cost, although this comparison was not integrated into the simulation model itself and should be read as indicative. Other consequences of the buffer storage design, including its higher expected maintenance intensity and lower operating costs relative to the traditional bufferlane system, must be taken into account in evaluating the business case.

In answer to the main research question:

"How can a shuttle-based dispatch buffer be designed and controlled to minimise crane congestion and outbound delays, improving outbound performance, in an automated cold-storage warehouse?"

A shuttle-based dispatch buffer can meaningfully reduce dependence on accurate truck arrival estimates and reduce the total crane workload required per pallet, provided picked pallets are routed directly into the buffer rather than back through the highbay. Under the demand levels and layout tested, this benefit shows up primarily as a resource-reallocation opportunity, fewer cranes for comparable outbound performance, rather than as a proportional reduction in total truck loading time, which instead depends on fixed manual labour. This result comes with a trade-off in operating margin that a warehouse operator must weigh explicitly: because the system was dimensioned for peak demand, it is sized with limited slack, leaving it more vulnerable to disruption from errors or maintenance than a design with greater built-in margin. Despite this trade-off, and despite requiring more shuttles than the alternative configurations tested, direct picked-pallet routing remains the more resource-efficient of the configurations evaluated in this research.

For NewCold, this research provides a validated basis for the Piacenza 2 design decision, together with a reusable simulation model and analysed model inputs that can support further design evaluation

beyond the specific scenarios tested here. For the wider warehouse industry handling with pallets, this research demonstrates that a shuttle-based dispatch buffer's primary value lies in decoupling outbound operations from truck-arrival uncertainty and in reallocating resource pressure away from cranes. Academically, this research addresses a gap in the SBS/RS literature by applying shuttle-based storage and retrieval methodology to an outbound dispatch-buffering role, validated against real operational data, and identifies external resource capacity as a coordination dependency not previously acknowledged in studies that treat the shuttle system as self-contained.

References

- Abdul Rahman, N. S. F., Karim, N. H., Md Hanafiah, R., Abdul Hamid, S., & Mohammed, A. (2023). Decision analysis of warehouse productivity performance indicators to enhance logistics operational efficiency. *International Journal of Productivity and Performance Management*, 72(4), 962–985. <https://doi.org/10.1108/IJPPM-06-2021-0373>
- Akkerman, R., Farahani, P., & Grunow, M. (2010). Quality, safety and sustainability in food distribution: a review of quantitative operations management approaches and challenges. *OR Spectrum* 2010 32:4, 32(4), 863–904. <https://doi.org/10.1007/s00291-010-0223-2>
- Albert, P. W., Rönqvist, M., & Lehoux, N. (2023). Trends and new practical applications for warehouse allocation and layout design: a literature review. *SN Applied Sciences*, 5(12). <https://doi.org/10.1007/s42452-023-05608-0>
- Amjath, M., Kerbache, L., Smith, J. M. G., & Elomri, A. (2023). Optimisation of Buffer Allocations in Manufacturing Systems: A Study on Intra and Outbound Logistics Systems Using Finite Queueing Networks. *Applied Sciences* 2023, Vol. 13, Page 9525, 13(17), 9525. <https://doi.org/10.3390/app13179525>
- Ascheuer, N., Grötschel, M., & Abdel-Aziz Abdel-Hamid, A. (1999). Order picking in an automatic warehouse: Solving online asymmetric TSPs. *Mathematical Methods of Operations Research*, 49(3), 501–515. <https://doi.org/10.1007/s001860050064>
- Aung, M. M., & Chang, Y. S. (2023). *Cold Chain Management*. Springer International Publishing. <https://doi.org/10.1007/978-3-031-09567-2>
- Azadeh, K., Koster, R., & Roy, D. (2019). EUR Research Information Portal Robotized and Automated Warehouse Systems: Review and Recent Developments. <https://doi.org/10.1287/trsc.2018.0873>
- Baglio, M., Creazza, A., & Dallari, F. (2023). The “Perfect” Warehouse: How Third-Party Logistics Providers Evaluate Warehouse Features and Their Performance. *Applied Sciences* 2023, Vol. 13, Page 6862, 13(12), 6862. <https://doi.org/10.3390/app13126862>
- Bajec, P., Tuljak-Suban, D., & Bajor, I. (2020). A warehouse social and environmental performance metrics framework[Socialni in okoljski kriteriji za merjenje uspešnosti v skladišču]. *Promet - Traffic and Transportation*, 32(4), 513–526. <https://doi.org/10.7307/ptt.v32i4.3390>
- Ballou, R. H. (2007). The evolution and future of logistics and supply chain management. *European Business Review*, 19(4), 332–348. <https://doi.org/10.1108/09555340710760152>
- Baruffaldi, G., Accorsi, R., Manzini, R., & Ferrari, E. (2020). Warehousing process performance improvement: a tailored framework for 3PL. *Business Process Management Journal*, 26(6), 1619–1641. <https://doi.org/10.1108/BPMJ-03-2019-0120>
- Battarra, I., Accorsi, R., Lupi, G., & Manzini, R. (2024). A design framework for shuttle-based automated storage systems. *IFAC-PapersOnLine*, 58(19), 1216–1221. <https://doi.org/10.1016/J.IFACOL.2024.09.086>
- Battini, D., Glock, C. H., Grosse, E. H., Persona, A., & Sgarbossa, F. (2016). Human energy expenditure in order picking storage assignment: A bi-objective method. *Computers & Industrial Engineering*, 94, 147–157. <https://doi.org/10.1016/J.CIE.2016.01.020>
- Blatnický, M., Dižo, J., Sága, M., Gerlici, J., & Kuba, E. (2020). Design of a Mechanical Part of an Automated Platform for Oblique Manipulation. *Applied Sciences* 2020, Vol. 10, Page 8467, 10(23), 8467. <https://doi.org/10.3390/app10238467>
- Bozek, P., Karavaev, Y. L., Ardentov, A. A., & Yefremov, K. S. (2020). Neural network control of a wheeled mobile robot based on optimal trajectories. *International Journal of Advanced Robotic Systems*, 17(2). <https://doi.org/10.1177/1729881420916077>
- BRCGS Food Safety Global Standard. (n.d.). <https://www.brcgs.com/our-standards/food-safety/>
- Carlo, H. J., & Vis, I. F. (2012). Sequencing dynamic storage systems with multiple lifts and shuttles. *International Journal of Production Economics*, 140(2), 844–853. <https://doi.org/10.1016/j.ijpe.2012.06.035>


- Chen, N., Liu, Q., Stević, Ž., Andrejić, M., & Pajić, V. (2024). An integrated cost based approach for warehouse performance evaluation: A new multiphase model. *Alexandria Engineering Journal*, 101(5), 62–77. <https://doi.org/10.1016/j.aej.2024.05.063>
- Chiaraviglio, A., Grimaldi, S., Zenezini, G., & Rafele, C. (2025). Overall Warehouse Effectiveness (OWE): A New Integrated Performance Indicator for Warehouse Operations. *Logistics* 2025, Vol. 9, Page 7, 9(1), 7. <https://doi.org/10.3390/logistics9010007>
- Cooksey, D. L. (1986). Comparison of conventional and automated refrigerated storage systems. *International Journal of Refrigeration*, 9(2), 99–104. [https://doi.org/10.1016/0140-7007\(86\)90041-1](https://doi.org/10.1016/0140-7007(86)90041-1)
- D'Antonio, G., & Chiabert, P. (2019). Analytical models for cycle time and throughput evaluation of multi-shuttle deep-lane AVS/RS. *International Journal of Advanced Manufacturing Technology*, 104(5-8), 1919–1936. <https://doi.org/10.1007/S00170-019-03985-8/FIGURES/4>
- Datex. (n.d.). How Technology Solves Labor Challenges in Cold Storage Warehouses. <https://datexcorp.com/how-technology-solves-labor-challenges-in-cold-storage-warehouses/>
- de Koster, R. B., Le-Duc, T., & Roodbergen, K. J. (2007). Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, 182(2), 481–501. <https://doi.org/10.1016/j.ejor.2006.07.009>
- Directive 2006/42/EC - machinery directive | Safety and health at work EU-OSHA. (n.d.). <https://osha.europa.eu/en/legislation/directives/directive-2006-42-ec-of-the-european-parliament-and-of-the-council>
- Discrete Event Simulation Software (DES) | Simio. (n.d.). <https://www.simio.com/discrete-event-simulation>
- Dukic, G., & Opetuk, T. (2012). Warehouse layouts. *Springer* 4, 9781447122746, 55–69. <https://doi.org/10.1007/978-1-4471-2274-6\3/TABLES/1>
- Eder, M. (2020a). Analytical model to estimate the performance of shuttle-based storage and retrieval systems with class-based storage policy. *The International Journal of Advanced Manufacturing Technology* 2020 107:5, 107(5), 2091–2106. <https://doi.org/10.1007/s00170-020-04990-y>
- Eder, M. (2020b). An approach for performance evaluation of SBS/RS with shuttle vehicles serving multiple tiers of multiple-deep storage rack. *International Journal of Advanced Manufacturing Technology*, 110(11-12), 3241–3256. <https://doi.org/10.1007/s00170-020-06033-y>
- Ekren, B. Y., Sari, Z., & Lerher, T. (2015). Warehouse Design under Class-Based Storage Policy of Shuttle-Based Storage and Retrieval System. *IFAC-PapersOnLine*, 48(3), 1152–1154. <https://doi.org/10.1016/j.ifacol.2015.06.239>
- Faveto, A., Traini, E., Bruno, G., & Chiabert, P. (2023). Review-based method for evaluating key performance indicators: an application on warehouse system. *The International Journal of Advanced Manufacturing Technology* 2023 130:1, 130(1), 297–310. <https://doi.org/10.1007/s00170-023-12684-4>
- Frazelle, E. (2016). *World-Class Warehousing and Material Handling, 2nd Edition* (Vol. 1). McGraw-Hill Education. <https://learning-oreilly-com.tudelft.idm.oclc.org/library/view/world-class-warehousing-and/9780071842839/cover.html>
- Gagliardi, J.-P., Renaud, J., Ruiz, A., Gagliardi, J.-P., Renaud, J., & Ruiz, A. (2012). A simulation modeling framework for multiple-aisle automated storage and retrieval systems. *Journal of Intelligent Manufacturing* 2012 25:1, 25(1), 193–207. <https://doi.org/10.1007/s10845-012-0686-x>
- Gould, W. P. (2019). Cold Storage. *Quarantine Treatments for Pests of Food Plants*, 119–132. <https://doi.org/10.1201/9780429303654-8>
- Gu, J., Goetschalckx, M., & McGinnis, L. F. (2007). Research on warehouse operation: A comprehensive review. *European Journal of Operational Research*, 177(1), 1–21. <https://doi.org/10.1016/j.ejor.2006.02.025>
- Gu, J., Goetschalckx, M., & McGinnis, L. F. (2010). Research on warehouse design and performance evaluation: A comprehensive review. *European Journal of Operational Research*, 203(3), 539–549. <https://doi.org/10.1016/j.ejor.2009.07.031>
- Gu, W., Tang, N., Wang, L., Guo, Z., Cao, Y., & Yuan, M. (2025). Coordinated Scheduling for Zero-Wait RGV/ASR Warehousing Systems with Finite Buffers. *Machines* 2025, Vol. 13, Page 546, 13(7), 546. <https://doi.org/10.3390/machines13070546>

- Guo, X., Yu, Y., & De Koster, R. B. (2016). Impact of required storage space on storage policy performance in a unit-load warehouse. *International Journal of Production Research*, 54(8), 2405–2418. <https://doi.org/10.1080/00207543.2015.1083624>;PAGE:STRING:ARTICLE/CHAPTER
- Ha, Y., & Chae, J. (2018). Free balancing for a shuttle-based storage and retrieval system. *Simulation Modelling Practice and Theory*, 82, 12–31. <https://doi.org/10.1016/j.simpat.2017.12.006>
- Han, J. W., Zuo, M., Zhu, W. Y., Zuo, J. H., Lü, E. L., & Yang, X. T. (2021). A comprehensive review of cold chain logistics for fresh agricultural products: Current status, challenges, and future trends. *Trends in Food Science and Technology*, 109, 536–551. <https://doi.org/10.1016/j.tifs.2021.01.066>
- Huertas, J. I., Díaz Ramírez, J., & Trigos Salazar, F. (2007). Layout evaluation of large capacity warehouses. *Facilities*, 25(7-8), 259–270. <https://doi.org/10.1108/02632770710753307>
- IDEFØ – Function Modeling Method – IDEF. (n.d.). https://www.idef.com/idefo-function_modeling_method/
- ISO 14001:2026 - ISO 14001 - Environmental management systems. (n.d.). <https://www.iso.org/standard/14001>
- ISO 45001:2018 - Occupational health and safety management systems. (n.d.). <https://www.iso.org/standard/63787.html>
- José, J., Reyes, R., Solano-Charris, E. L., & Montoya-Torres, J. R. (2019). The storage location assignment problem: A literature review. *International Journal of Industrial Engineering Computations*, 10, 199–224. <https://doi.org/10.5267/j.ijiec.2018.8.001>
- Kalafat, İ., Hekimoğlu, M., Yücekaya, A. D., Kirkil, G., Ediger, V., & Yıldırım, Ş. (2025). An Integrated Framework for Internal Replenishment Processes of Warehouses Using Approximate Dynamic Programming. *Applied Sciences* 2025, Vol. 15, Page 7767, 15(14), 7767. <https://doi.org/10.3390/app15147767>
- Kriehn, T., Schloz, F., Wehking, K. H., & Fittinghoff, M. (2018). Impact of class-based storage, sequencing of retrieval requests and warehouse reorganisation on throughput of shuttle-based storage and retrieval systems. *FME Transactions*, 46(3), 320–329. <https://doi.org/10.5937/fmet1803320K>
- Küçükyaşar, M., Y. Ekren, B., & Lerher, T. (2021). Cost and performance comparison for tier-captive and tier-to-tier SBS/RS warehouse configurations. *International Transactions in Operational Research*, 28(4), 1847–1863. <https://doi.org/10.1111/itor.12864>
- Kumar, S., & Mahapatra, R. P. (2021). Design of multi-warehouse inventory model for an optimal replenishment policy using a Rain Optimization Algorithm. *Knowledge-Based Systems*, 231, 107406. <https://doi.org/10.1016/J.KNOSYS.2021.107406>
- Law, A., & Kelton, W. (2000). *Simulation modeling & analysis*. McGraw Hill.
- Legowo, N., & Wijaya, W. (2022). Measurement of Warehouse Management System Performance using IT-BSC Method in the FMCG Industry. *Journal of System and Management Sciences*, 12(5), 230–251. <https://doi.org/10.33168/JSMS.2022.0514>
- Lerher, T. (2015, December). *Throughput and Energy Related Performance Calculations for Shuttle Based Storage and Retrieval Systems – Nova Science Publishers*. Nova Science Publishers. <https://novapublishers.com/shop/throughput-and-energy-related-performance-calculations-for-shuttle-based-storage-and-retrieval-systems/>
- Lerher, T., Ekren, B. Y., Sari, Z., & Rosi, B. (n.d.). Method for evaluating the throughput performance of shuttle based storage and retrieval systems. <https://doi.org/10.17559/TV-20141022121007>
- Lerher, T., Ficko, M., & Palčič, I. (2021). Throughput performance analysis of Automated Vehicle Storage and Retrieval Systems with multiple-tier shuttle vehicles. *Applied Mathematical Modelling*, 91, 1004–1022. <https://doi.org/10.1016/j.apm.2020.10.032>
- Li, G., Ye, B., Ye, J., Sheng, J., Li, M., Liu, J., Zhao, Y., Yu, G., & Xu, J. (2024). Data Collection and System Integration of Distributed Message Queues for Stacker Cranes. *2024 3rd International Conference on Artificial Intelligence and Computer Information Technology, AICIT 2024*. <https://doi.org/10.1109/AICIT62434.2024.10730029>
- Li, Y., Yeo, S., Keng Swee, Li, H., & Chew, E. P. (2026). Multi-objective simulation-optimization for a distribution center resource planning. *IIE Transactions*, 58(3), 324–340. <https://doi.org/10.1080/24725854.2025.2475307>
- Li, Y., & Li, Z. (2022). Shuttle-Based Storage and Retrieval System: A Literature Review. *Sustainability* 2022, Vol. 14, Page 14347, 14(21), 14347. <https://doi.org/10.3390/su142114347>

- Liu, Q., Tang, Y., & Yu, M. (2024). A new discrete-time queueing model to optimize cargo dispatch for a warehouse. *Operational Research 2024* 25:1, 25(1), 7–. <https://doi.org/10.1007/s12351-024-00888-9>
- Liu, Z., Wang, Y., Jin, M., Wu, H., & Dong, W. (2021). Energy consumption model for shuttle-based Storage and Retrieval Systems. *Journal of Cleaner Production*, 282, 124480. <https://doi.org/10.1016/J.JCLEPRO.2020.124480>
- Logistics. (n.d.). <https://www.ifs-certification.com/en/ifs-portfolio/standards/logistics-standard>
- Lu, C., Zeng, B., & Liu, S. (2020). A Study on the Block Relocation Problem: Lower Bound Derivations and Strong Formulations. *IEEE Transactions on Automation Science and Engineering*, 17(4), 1829–1853. <https://doi.org/10.1109/TASE.2020.2979868>
- Lu, C. S., & Yang, C. C. (2010). Logistics service capabilities and firm performance of international distribution center operators. *The Service Industries Journal*, 30(2), 281–298. <https://doi.org/10.1080/02642060802123392>
- Marchet, G., Melacini, M., Perotti, S., & Tappia, E. (2013). Development of a framework for the design of autonomous vehicle storage and retrieval systems. *International Journal of Production Research*, 51(14), 4365–4387. <https://doi.org/10.1080/00207543.2013.778430>
- Marchi, B., Zaroni, S., & Jaber, M. Y. (2020). Energy implications of lot sizing decisions in refrigerated warehouses. *Energies*, 13(7). <https://doi.org/10.3390/en13071739>
- Marolt, J., Sgarbossa, F., Jimenez, J. A., Sharotry, A., & Lerher, T. (2025). Analytical model and swapping policy assessment of a vertical lift module–buffer integrated storage system. *International Journal of Production Research*, 63(6), 2150–2169. <https://doi.org/10.1080/00207543.2024.2396513>
- Meneghetti, A., & Monti, L. (2015). Greening the food supply chain: An optimisation model for sustainable design of refrigerated automated warehouses. *International Journal of Production Research*, 53(21), 6567–6587. <https://doi.org/10.1080/00207543.2014.985449>
- MHI & Deloitte. (2023). *THE RESPONSIBLE SUPPLY CHAIN 2023 MHI Annual Industry Report* (tech. rep.).
- Moghadam, M. F., Khalili damghani, K., Ghezavati, V., & Komijan, A. R. (2025). A Bi-level optimization method for integrated warehouse design, storage location assignment, and order picking. *MethodsX*, 15(6), 103595. <https://doi.org/10.1016/j.mex.2025.103595>
- Mohammed, A., Rahman, M., Demir, G., Shah, B., & Alaya, A. (2025). When warehouses dream: a systems-informed exploration of Industry 5.0 principles in warehousing performance management. *International Journal of Systems Science: Operations and Logistics*, 12(1), 2559257. <https://doi.org/10.1080/23302674.2025.2559257>
- Mohamud, I. H., Kafi, A. A., Shahron, S. A., Zainuddin, N. B., & Musa, S. (2023). The Role of Warehouse Layout and Operations in Warehouse Efficiency: A Literature Review. *Journal Europeen des Systemes Automatises*, 56(1), 61–68. <https://doi.org/10.18280/jesa.560109>
- Nasir, D., Rakesh Venkitasubramony, ., Suresh, ., Jakhar, K., Venkitasubramony, R., Kumar Jakhar, S., & Nasir, D. (2024). Ergonomics in warehouse design and operations: a systematic literature review. *Operational Research 2024* 25:1, 25(1), 10–. <https://doi.org/10.1007/s12351-024-00892-z>
- Nof, S. Y. (2007). Handbook of industrial robotics. *Handbook of Industrial Robotics*, 1–1327. <https://doi.org/10.1002/9780470172506>
- Popovic, D., Vidovic, M., & Bjelic, N. (2012). Application of genetic algorithms for sequencing of AS/RS with a triple-shuttle module in class-based storage. *Flexible Services and Manufacturing Journal* 2012 26:3, 26(3), 432–453. <https://doi.org/10.1007/s10696-012-9139-2>
- Ribino, P., Cossentino, M., Lodato, C., & Lopes, S. (2017). Agent-based simulation study for improving logistic warehouse performance. *Journal of Simulation* 2017, 12(1), 1–20. <https://doi.org/10.1057/s41273-017-0055-z>
- Ritter, B., Barker, C., & Karl, K. (2023, January). Expanding Global Cold Chains: Effective Adaptation, or Dangerous Contribution to Climate Change? – State of the Planet. <https://news.columbia.edu/2023/01/27/expanding-global-cold-chains-effective-adaptation-or-dangerous-contribution-to-climate-change/>
- Roodbergen, K. J., Sharp, G. P., & Vis, I. F. (2007). EUR Research Information Portal Designing the Layout Structure of Manual Order-Picking Areas in Warehouses. *IIE Transactions*, 40, 1032–1045. <https://doi.org/10.1080/07408170802167639>

- Roodbergen, K. J., & Vis, I. F. (2009a). A survey of literature on automated storage and retrieval systems. *European Journal of Operational Research*, 194(2), 343–362. <https://doi.org/10.1016/j.ejor.2008.01.038>
- Roodbergen, K. J., & Vis, I. F. (2009b). A survey of literature on automated storage and retrieval systems. *European Journal of Operational Research*, 194(2), 343–362. <https://doi.org/10.1016/j.ejor.2008.01.038>
- Roosbeh Nia, A., Haleh, H., & Saghaei, A. (2017). Dual command cycle dynamic sequencing method to consider GHG efficiency in unit-load multiple-rack automated storage and retrieval systems. *Computers & Industrial Engineering*, 111, 89–108. <https://doi.org/10.1016/j.cie.2017.07.007>
- Rouwenhorst, B., Reuter, B., Stockrahm, V., van Houtum, G. J., Mantel, R. J., & Zijm, W. H. M. (2000). Warehouse design and control: Framework and literature review. *European Journal of Operational Research*, 122(3), 515–533. [https://doi.org/10.1016/S0377-2217\(99\)00020-X](https://doi.org/10.1016/S0377-2217(99)00020-X)
- Schuberg Philis. (2023, December). 'Snelle groei NewCold door hoge mate van automatisering' | NT. <https://www.nt.nl/brandstories/2023/12/19/snelle-groei-newcold-door-hoge-mate-van-automatisering/?gdpr=accept>
- Shahroudnejad, A., Mousavi, P., Perepelytsia, O., Sahir, Staszak, D., Taylor, M. E., & Bawel, B. (2024). A novel framework for automated warehouse layout generation. *Frontiers in Artificial Intelligence*, 7, 1465186. <https://doi.org/10.3389/FRAI.2024.1465186/TEXT>
- Skawińska, E., & Zalewski, R. I. (2022). Economic Impact of Temperature Control during Food Transportation—A COVID-19 Perspective. *Foods* 2022, Vol. 11, Page 467, 11(3), 467. <https://doi.org/10.3390/foods11030467>
- Statista Research Department. (2025, December). Cold chain logistics - statistics & facts | Statista. <https://www.statista.com/topics/4321/cold-chain-logistics/#topicOverview>
- Staudt, F. H., Alpan, G., Di-Mascolo, M., & Rodríguez, C. M. T. (2015). Warehouse performance measurement: A literature review. *International Journal of Production Research*, 53(18), 5524–5544. <https://doi.org/10.1080/00207543.2015.1030466>
- Tako, A. A., & Robinson, S. (2012). The application of discrete event simulation and system dynamics in the logistics and supply chain context. *Decision Support Systems*, 52(4), 802–815. <https://doi.org/10.1016/J.DSS.2011.11.015>
- Taoukis, P. S., Gogou, E., Tsironi, T., Giannoglou, M., Dermesonlouoglou, E., & Katsaros, G. (2016). Food Cold Chain Management and Optimization. In *Food engineering series* (pp. 285–309). Springer. https://doi.org/10.1007/978-3-319-24040-4_16
- Tokat, S., Karagul, K., Sahin, Y., & Aydemir, E. (2022). Fuzzy c-means clustering-based key performance indicator design for warehouse loading operations. *Journal of King Saud University - Computer and Information Sciences*, 34(8), 6377–6384. <https://doi.org/10.1016/j.jksuci.2021.08.003>
- Trab, S., Bajic, E., Zouinkhi, A., Abdelkrim, M. N., Chekir, H., & Ltaief, R. H. (2015). Product Allocation Planning with Safety Compatibility Constraints in IoT-based Warehouse. *Procedia Computer Science*, 73, 290–297. <https://doi.org/10.1016/J.PROCS.2015.12.033>
- Tutam, M., & White, J. A. (2019). A multi-dock, unit-load warehouse design. *IIE Transactions*, 51(3), 232–247. <https://doi.org/10.1080/24725854.2018.1488307;WGROU:STRING:PUBLICATIO>
- Wang, Y., Mou, S., & Wu, Y. (2015). Task scheduling for multi-tier shuttle warehousing systems. *International Journal of Production Research*, 53(19), 5884–5895. <https://doi.org/10.1080/00207543.2015.1012604>
- Wu, Y., Zhou, C., Ma, W., & Kong, X. T. (2020). Modelling and design for a shuttle-based storage and retrieval system. *International Journal of Production Research*, 58(16), 4808–4828. <https://doi.org/10.1080/00207543.2019.1665202>
- Xu, Z., Xu, L., Shao, H., Lehmann, T., & Matta, A. (2025). Operation time and rack stability optimisation in tier-to-tier shuttle-based storage and retrieval systems with multiple retrieval locations. *Computers in Industry*, 173, 104385. <https://doi.org/10.1016/j.compind.2025.104385>
- Yener, F., & Yazgan, H. R. (2019). Optimal warehouse design: Literature review and case study application. *Computers and Industrial Engineering*, 129, 1–13. <https://doi.org/10.1016/j.cie.2019.01.006>

- Ying, L., Xun, W., Hongqiang, F., & Jingtian, Z. (2026). Improved A* algorithm design for container relocation problem in four-way shuttle vehicle-based digital warehousing systems. *Digital Engineering*, 9(4), 100086. <https://doi.org/10.1016/j.dte.2025.100086>
- Yu, W., & Egbelu, P. J. (2008). Scheduling of inbound and outbound trucks in cross docking systems with temporary storage. *European Journal of Operational Research*, 184(1), 377–396. <https://doi.org/10.1016/j.ejor.2006.10.047>
- Zhang, F., Yang, X., Dang, Y., Wu, C., & Chen, B. (2025). Construction and Practice of Intelligent Warehouse Management Decision Support System Based on Artificial Intelligence. *Procedia Computer Science*, 259(2), 711–720. <https://doi.org/10.1016/j.procs.2025.04.022>
- Zhao, Y., Zhang, X., & Xu, X. (2019). Application and research progress of cold storage technology in cold chain transportation and distribution. *Journal of Thermal Analysis and Calorimetry 2019* 139:2, 139(2), 1419–1434. <https://doi.org/10.1007/s10973-019-08400-8>
- Zhen, L., & Li, H. (2022). A literature review of smart warehouse operations management. *Frontiers of Engineering Management*, 9(1), 31–55. <https://doi.org/10.1007/s42524-021-0178-9>
- Zou, B., Xu, X., Gong, Y., & De Koster, R. (2016). Modeling parallel movement of lifts and vehicles in tier-captive vehicle-based warehousing systems. *European Journal of Operational Research*, 254(1), 51–67. <https://doi.org/10.1016/J.EJOR.2016.03.039>



IDEF0 Model

Complementary to the flowcharts presented in chapter 4, IDEF0 diagrams are made to illustrate warehouse flows and processes for all three scenarios.

A.1. Scenario 1

The IDEF0 model consists of interconnecting diagrams with a hierarchy. The highest level is the whole warehouse process as one activity (A-0 level), further levels are separate activities with subactivities if applicable. Every activity in the model has inputs and outputs, and can have multiple controls (upper arrow) and mechanisms (lower arrow). The A-0 level is shown in figure A.1, whereafter an explanation of the elements is given.




Figure A.1: A-0 diagram for an automated, cold-storage warehouse.

The main activity is the distribution of orders from an automated, cold-storage warehouse. The input of this activity is trucks that are loaded with pallets at the inbound docks, and the output of the activity are trucks loaded with orders at the outbound docks. The controls, are the following: OSA (operating service agreement), outbound planning, truck arrival times, dock scheduling system, location assignment policy, order management system, food safety regulations, and health and safety rules. The mechanisms are human operators, MHE (material handling equipment), racking, WMS (warehouse management system), MFCS (material flow control system), and airlocks.

Controls and mechanisms

OSA is the agreement between customers and the warehouse operator. The OSA can influence the prioritization of activities and handling. Outbound planning is the result of orders being processed into type of products, bulk or pick pallets, and amounts. Depending on the order characteristics, the planning is created for following activities as retrieval from the highbay, picking, and staging. The outbound planning also controls other activities unrelated to the outbound processes due to the prioritization of outbound processes in warehouse systems that also deal with inbound pallets. Truck arrival times, inbound and outbound, activate processes happening in the warehouse. The dock scheduling system deals with the registration of incoming trucks to inbound and outbound docks, connecting orders and trips, as well as pallets to be unloaded, to a dock. A location assignment policy consists of the standard rules that are used for decision making regarding the allocation of pallets to storage location in the highbay. The order management system registers customer orders, consisting of product details and pickup scheduling.

Food safety policies as frozen temperatures, fire risk reduction, and the limitation of human interaction are followed by ISP140001, IFS and BRC. ISO 14001 is the Environmental Management System NewCold holds, a standard that provides a framework for identifying environmental impacts, reduce them, and improve them ("ISO 14001:2026 - ISO 14001 - Environmental management systems", n.d.). Relevant for cold-storage warehouses is the energy consumption and waste management. The British Retail Consortium (BRC) Global Standard for Storage and Distribution is a specific storage and distribution standard ("BRCGS Food Safety Global Standard", n.d.) that NewCold complies with. It covers food safety management systems, temperature control, facility standards, and vehicle and equipment standards. A similar standard, is the IFS Logistics International Featured Standards covering the safety and quality of logistics services of food products ("Logistics", n.d.). The IFS uses a scoring system where each requirement is scored numerically, partial compliance is recognised with this. Health and safety requirements take human activities and machines into account. The machinery in the warehouse should comply to the CE machinery Directive (2006/42/EC) and ISO 45001 is held for general health and safety. ISO 45001 is another management system that NewCold holds, but for Occupational Health and Safety ("ISO 45001:2018 - Occupational health and safety management systems", n.d.). This is the international standard that requires organizations to identify and control health and safety risks, with continuous improvement. For machinery, safety standards are included in the CE machinery Directive (2006/42/EC) ("Directive 2006/42/EC - machinery directive | Safety and health at work EU-OSHA", n.d.).

Human operators are the warehouse employees that are in charge of activities in the lowbay such as unloading trucks, reworking pallets, picking pallets and loading trucks. MHE systems that are taken into account are: conveyor systems, monorail, label applicator, I-point, ATU and ATL, lifts, EPT/forklifts, cranes, and buffer lanes. A conveyor system consists of all types of conveyors including turntables and T-carts, used for transporting pallets throughout the whole warehouse. A monorail is used to connect the inbound area with the highbay and outbound area. A label applicator is used depending on the OSA and is placed on the monorail between the highbay and the outbound area. The I-point registers a pallet once it arrives and checks for irregularities. Lifts are often used for a non-ground level picking area. EPTs (Electric Pallet Trucks) or forklifts are used for the movement of pallets in the lowbay area for unloading, reworking, picking and loading pallets. These vehicles are always driven by a human operator. Cranes transport pallets within the highbay, storing and retrieving pallets, and relocating pallets for efficient operations. Buffer lanes are conveyor belts where trips are staged before truck loading. Racking is the base of storage locations with several levels, used in the highbay and on the pickfloor. Airlocks secure oxygen reduction in the highbay, just as temperature efficiency. The Warehouse Management System manages the whole operational planning in a warehouse and gives orders to the MFCS, the Material Flow Control System. This software system executes the WMS orders in terms of instructions for routes, MHE assignment and sequencing. The MFCS is the system that decides which pallets go to which buffer lane and when and which crane retrieves a pallet for an order instruction from the WMS. Next to these two control systems, there are others (shown in A.2): mechanical (physical hardware like a conveyor belt), electrical (cables and sensors on mechanical systems), PLC (Programmable Logic Controller, low-level controls), than the WCS (Warehouse Control System, in this case the MFCS), the WMS and the ERP (Enterprise Resource Planning).




Figure A.2: Control systems of MHE

The A-0 level is divided into three sub activities: processing inbound pallets, storing pallets, and sending out orders, shown in figure A.3.




Figure A.3: A-1 diagram for an automated, cold-storage warehouse.

Processing inbound pallets

A decomposition of the A1 activity in the A-1 diagram, is shown in figure A.4, where the unloading and processing of inbound pallets happens including registration at the I-point.




Figure A.4: A-2 diagram of processing inbound pallets in an automated, cold-storage warehouse.

Inbound truck arrivals have limited effects on warehouse performance. Inbound truck arrivals are not planned or booked in advance. Trucks that arrive get a random inbound dock assigned. Trucks are unloaded manually by a maximum of two employees that drive a forklift. Pallets are placed on the inbound conveyors and routed to the I-point. Once a pallet is accepted, the WMS assigns a storage location in the highbay according to a rule-based location assignment policy that takes item characteristics, location availability and workload balancing into account. The detailed decision rules and their order of precedence are not disclosed in this public version. The MFCS controls the movements of the pallet to the highbay and within the highbay. Once a pallet gets rejected based on size, damage or weight, it will move to a parking spot for reworking. When the pallet is broken or cannot be easily fixed, the pallet will be transported by an employee and a forklift to the rework station where the employee will repair the pallet. Hereafter it will be infed again at the I-point by the employee and the forklift.

Storage of pallets

A pallet is stored in an optimal location decided by the WMS and controlled by the MFCS.

Sending out orders

A decomposition of the A3 activity in the A-1 diagram is shown in figure A.5 where the outbound process of orders is displayed in six subactivities: processing orders, retrieval, picking, transportation to the bufferlanes, staging on bufferlanes, and truck loading.




Figure A.5: A-2 diagram for sending out orders in an automated, cold-storage warehouse.

A customer will let the carrier know which order they want to receive and at what moment, in the order system. The carrier will plan its operations accordingly as one trip (full truck) can include multiple orders and drop-off locations, and book a timeslot at the warehouse for pickup. There is also a difference between the automation of the order process of customers and transport companies per warehouse. For the warehouse in Piacenza, the site crew needs to proactively contact customers on their arrival times as most customers do not have automated order systems. The order is specified by bulk and/or pick pallets. For a picked pallet, the retrieval time is higher than a bulk pallet, as the pallet needs to move to the highbay again, and retrieved from there. Requiring an earlier activation of retrieval. Ordered pallets (bulk or already picked) will move to the bufferlanes where the pallets are staged until pickup. Lastly, the pallet will be loaded into the truck at the outbound dock by a an employee that drives a forklift.

The picking activity (A33) is decomposed, shown in figure A.6 with transport included.




Figure A.6: A-3 diagram for sending out orders in an automated, cold-storage warehouse.

When the pallet arrives at the pick floor, an employee with a forklift will transport the pallet to a free area and perform the picking where the goods are combined from another pallet and/or separated from other goods on the pallet. After picking the pallet is brought to the inbound conveyor by the employee and forklift, and the pallet will transport to the highbay again. Locations in the highbay with limited travel times to exit are reserved for these short-term storage of picked pallets.

A.2. Scenario 2

Compared to scenario 1, some updates are made in the warehouse flow for scenario 2. An extra mechanism is added to the process: the 2D shuttle, lifts and its racking. In figure A.7 the A-2 level for the shuttle based buffer implementation is shown.




Figure A.7: A-2 diagram for sending out orders in an automated, cold-storage warehouse with shuttle based buffer storage.

For pick pallets, the retrieval will activate from 30 earlier, go back to the highbay and be activated again, acting like a bulk pallet. The pallet will move to the buffer storage, and the order will only be retrieved once the corresponding truck arrives. This will possibly minimize dock occupancy time, and increase the amount of trucks that can be loaded per dock. However, the time a truck waits at a dock should not increase compared to the traditional outbound layout of the warehouse. The staging of a pallet in the outbound buffer storage has three subactivities: transporting pallet to outbound buffer storage, staging pallet, and retrieving pallet. The process is shown in figure A.8.




Figure A.8: A-3 diagram for staging pallet in outbound buffer storage in an automated, cold-storage warehouse.

A.3. Scenario 3

Compared to scenario 2, the A-2 level changes with the adaption for the changed routing for picked pallets. The diagram is shown in figure A.9.





Figure A.9: A-2 diagram for sending out orders in outbound buffer storage in an automated, cold-storage warehouse.

B

Model structure

Model processes, and model states discussed in chapter 4 and 5 and describe implementation and organisation of the model in Simio. The appendix is not disclosed in this public version.



Model parameters

Model inputs for Piacanza 1 and 2, discussed in chapter 5 are described and reasoned. Model properties that are discussed in chapter 5 for every scenario are also shown. The appendix is not disclosed in this public version.

D

Model validation

92

The model is setted up with various variables measuring flow, throughput, start- and end times of processes and activities, and counter variables. These variables contribute to constant model validation and help explaining results next to debugging. The validation variables are found in table D.1 below.

Variable	Description	Calculation (Simio)	Location in model (process)
Pallets			
Time pallet to be retrieved by crane	The time a pallet is requested to be retrieved from the outbound crane of the corresponding aisle a pallet is stored in, there could be a queue and the ranking of this queue is based on the closest planned pickup time.	Pallet.TimePalletAtRetrievalHighbay	Cranes entered
Time pallet retrieved by crane	The time a pallet is retrieved by the outbound crane and can continue its outbound journey.	Pallet.timePalletRetrievedHighbay	Cranes exited
Time pallet retrieval activated Highbay	The time an order activates and the ordered pallet matches with the corresponding pallet and activates the retrieval request.	Pallet.TimePalletRetrievalHighbayActivated	Highbay retrieval processing
Time pallet on buffer-lane	The time a pallet arrives on its assigned buffer-lane.	Pallet.TimePalletOnBufferlane	Combiner docks member entered / Bufferlanes entered
duration of retrieval Highbay	The time between a retrieval request and the time that the pallet is retrieved from the highbay.	Pallet.TimePalletRetrievedHighbay-Pallet.TimePalletAtRetrievalhighbay	

Continued on next page

Table D.1 – continued from previous page

Variable	Description	Calculation (Simio)	Location in model (process)
Time retrieval start Highbay	The time a crane arrives at the pallet and the pallet is transported by the crane.	Pallet.StartCraneRetrieval	Cranes before processing
Time retrieval activated to completed Highbay	The time between a retrieval request from the highbay and the time that the pallet is retrieved from the highbay.	Pallet.TimePalletRetrievedHighbay-Pallet.TimePalletRetrievalHighbayActivated	
Time retrieval activated to on bufferlane	The time between a retrieval request from the highbay and the time a pallet arrives on the bufferlane.	Pallet.TimePalletOnBufferlane-Pallet.TimePalletRetrievalHighbayActivated	
Time pallet retrieval activated buffer storage	The time a pallet is requested to be retrieved from the buffer storage, also the time the corresponding truck arrives at the outbound dock.	Pallet.TimePalletRetrievalBufferActivated	Buffer retrieval processing
Time pallet retrieved by shuttle	The time a pallet is retrieved by a shuttle on the level of storage in the buffer storage.	Pallet.TimePalletRetrievedByShuttle	Shuttles after processing
Time pallet retrieved by lift	The time a pallet is retrieved by a lift on the level of storage to the ground floor, and arriving at the tcar for transportation to the bufferlanes.	Pallet.TimePalletRetrievedByLift	Lifts after processing
time retrieval activated to completed buffer storage	The time between a retrieval request from the buffer storage to the completed retrieval by shuttle and lift.	Pallet.TimePalletRetrievedByLift-Pallet.TimePalletRetrievalBufferActivated	
Order activated to retrieval activated highbay	The time between an order is activated and the retrieval request of corresponding pallets is activated from the highbay.	Pallet.TimePalletRetrievalActivated-OrderActivationTime[Pallet.OrderID]	
Order activated to retrieval activated buffer storage	The time between an order is activated and the retrieval request of corresponding pallets is activated from the buffer storage.	Pallet.timePalletRetrievalBufferActivated-OrderActivationTime[Pallet.OrderID]	
time between shuttle and lift	The time between retrieval by lift and retrieval by shuttle in the buffer storage.	Pallet.TimePalletRetrievedByLift-Pallet.TimePalletRetrievedByShuttle	
time between pallet retrieved buffer storage and bufferlane	The time between a pallet retrieved from the buffer storage and its arrival on a bufferlane.	Pallet.TimePalletOnBufferlane-Pallet.TimePalletRetrievedByLift	
Orders			
planned pickup time	The planned pickup time of an order.	Trucks.PlannedPickupTime	Order arrival
first pallet on bufferlane	The time the first pallet of an order enters the bufferlane.	FirstPalletOnBufferlaneTime[Trucks.TruckOrderID]	Bufferlanes entered
actual arrival time	The actual arrival of a truck.	Trucks.ActualArrivalTime	Order arrival

Continued on next page

Table D.1 – continued from previous page

Variable	Description	Calculation (Simio)	Location in model (process)
actual departure time	The departure time of a loaded truck.	TimeNow	Combiner docks after processing
last pallet on bufferlane	The time the last pallet of an order enters the bufferlane.	LastPalletOnBufferlaneTime[Trucks.TruckOrderID]	Bufferlanes entered
total time on bufferlane	The time until an order is complete at a bufferlane or dock.	TimeNow-FirstPalletOnBufferlaneTime[Trucks.TruckOrderID]	
truck punctuality	The time between the actual arrival and the planned pickup time of an outbound truck. A negative value means that the truck is early, and a positive value means that the truck is late.	Trucks.ActualArrivalTime-Trucks.PlannedPickupTime	
pallets in order	The amount of pallets in an order that will be loaded into the truck.	PalletsPerOrder[Trucks.TruckOrderID]	Separator orders into ordered pallets before processing
time order activated	The time an order activates and is within the time of retrieval before planned arrival.	OrderActivationTime[Trucks.TruckOrderID]	Order processing bulk pallets before processing
order ready waiting time on bufferlane for truck	The time an order is complete at the bufferlane until the truck arrives at the dock.	TimeNow-LastPalletOnBufferlaneTime[Trucks.TruckOrderID]	
time until all pallets ready	The time between the first and last pallet arriving on the bufferlane, completing an order.	LastPalletOnBufferlaneTime[Trucks.TruckOrderID]-FirstPalletOnBufferlaneTime[Trucks.TruckOrderID]	
orderid	The orderid of a pallet or truck.	Trucks.TruckOrderID	Order arrival
time between planned pickup and actual departure	The time between the planned pickup time and the actual departure of a truck.	TimeNow-Trucks.PlannedPickupTime	
total loading time	The time between the actual arrival and the actual departure of a truck.	TimeNow-Trucks.ActualArrivalTime	
total time order activated	The time between order activation and truck departure.	TimeNow-OrderActivationTime[Trucks.TruckOrderID]	
time between planned pickup and order ready	The time between the planned pickup time and the completeness of the order at the bufferlane. A positive value means that the order was ready before the planned pickup time, a negative value means that the truck needed to wait until the order completed.	Trucks.PlannedPickupTime-LastPalletOnBufferlaneTime[Trucks.TruckOrderID]	

Continued on next page

Table D.1 – continued from previous page

Variable	Description	Calculation (Simio)	Location in model (process)
time between actual arrival and order ready	The time between the actual arrival time of a truck and the completeness of the order at the bufferlane. A positive value means that the order was ready before the planned pickup time, a negative value means that the truck needed to wait until the order completed.	$\text{Trucks}.\text{ActualArrivalTime} - \text{LastPalletOnBufferlaneTime}[\text{Trucks}.\text{TruckOrderID}]$	
order arrival time	The time an order is placed and scheduled.	$\text{DateAndTimeofArrivalOrders}[\text{Trucks}.\text{TrucksOrderID}]$	
Model states			
Time		TimeNow	
Pallets in warehouse		PalletsinSystem	Seperator docks inbound member exited / Combiner outbound docks member entered
Pallets in highbay		PalletsinHighbay	Highbay storage entered / Highbay retrieval exited
inbound trucks arrived		InboundTrucksArrived	Inbound truck arrival
outbound trucks arrived		OutboundTrucksArrived	Outbound truck arrival
Orders in system		OrderCounter	Order arrival
Ordered pallets in system		TotalOrderedPallets	Seperator orders into ordered pallets before member exiting
Successful retrievals from Highbay		SuccessfulretrievalsHighbay	Cranes after processing
Pallets passed IPoint per hour		PalletsThisHourIpoint	Ipoint after processing
Pallets in buffer storage		PalletsinBufferStorage	Buffer storage entered / Buffer retrieval before processing
highbay occupancy	Pallets in highbay divided by highbay capacity	$\text{PalletsInHighbay} / \text{HighbayCapacity}$	
buffer storage occupancy	Pallets in buffer storage divided by buffer storage capacity	$\text{PalletsinBufferStorage} / \text{BufferStorageCapacity}$	
trucks loaded per hour		TrucksloadedHourly	Combiner outbound docks after processing
trucks loaded per dock per hour		$\text{Trucksloadehourly} / \text{DocksAmount}$	

Continued on next page

Table D.1 – continued from previous page

Variable	Description	Calculation (Simio)	Location in model (process)
Outbound trucks departed		OutboundTrucksDeparted	Sink outbound trucks
Inbound trucks departed		InboundTrucksDeparted	Sink inbound trucks

Table D.1: Model validation variables

E

Additional Results

Resource utilisation in this appendix is reported in aggregated form by resource type. Individual unit identifiers and total equipment counts are not disclosed in this public version.

E.1. Scenario 1

Figures E.1 and E.2 show inbound and outbound trucks in the warehouse over the same representative day used in section 6.1.2..




Figure E.1: Inbound trucks in system over a representative day, scenario 1




Figure E.2: Outbound trucks in system over a representative day, scenario 1

Figures E.3 and E.4 show crane retrievals per hour and inbound registration throughput over the same 15th day of the experiment.




Figure E.3: Crane retrievals per hour over a representative day, scenario 1




Figure E.4: Inbound registration throughput over a representative day, scenario 1

E.2. Scenario 2

In figure E.1, all KPI details for the experiments are shown.

Table E.1: Scenario 2, full KPI detail by experiment

KPI	Configuration 1 -		Configuration 2 +1 crane		Configuration 3 +1 dock		Configuration 4 +1 shuttle	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Crane utilisation	97.23%	1.4 pp	84.44%	0.89 pp	83.68%	1.9 pp	84.41%	1.53 pp
Total loading time	56.36 h	12.78 h	52.98 h	7.81 h	54.96 h	29.24 h	17.51 h	11.42 h
Planned pickup to actual departure	55.57 h	12.79 h	52.24 h	7.83 h	54.21 h	29.29 h	16.76 h	11.42 h
Highbay occupancy	69.13%	1.74 pp	69.64%	1.02 pp	70.62%	1.70 pp	69.95%	0.84 pp
Buffer storage occupancy	149.8%	23.3 pp	126.1%	16.5 pp	115.7%	41.9 pp	51.7%	19.7 pp
Truck punctuality	-0.79 h	0.17 h	-0.74 h	0.18 h	-0.75 h	0.11 h	-0.76 h	0.20 h
Order ready wait time	0.089 h	0.0015 h	0.084 h	0.0023 h	0.082 h	0.004 h	0.088 h	0.0024 h
Crane retrieval time	2.66 h	2.21 h	0.51 h	0.07 h	0.41 h	0.09 h	0.48 h	0.12 h
Buffer storage retrieval time	2.53 h	0.38 h	2.23 h	0.53 h	1.62 h	0.88 h	1.17 h	0.31 h
Buffer storage to bufferlane time	0.090 h	0.016 h	0.103 h	0.019 h	0.108 h	0.006 h	0.098 h	0.008 h
Trucks loaded per dock per hour	0.234	0.012	0.236	0.014	0.211	0.053	0.292	0.039
Outbound dock assignment wait time	45.98 h	10.98 h	40.13 h	4.68 h	38.38 h	19.02 h	10.81 h	8.21 h

Table E.2: Scenario 2, full KPI detail by experiment (continued)

KPI	Configuration 5 +1 level		Configuration 6 + 1 level		Configuration 7 -1 dock	
	Mean	SD	Mean	SD	Mean	SD
Crane utilisation	83.21%	1.65 pp	83.67%	2.38 pp	83.30%	1.44 pp
Total loading time	7.79 h	11.61 h	1.94 h	0.26 h	1.97 h	0.32 h
Planned pickup to actual departure	7.17 h	11.55 h	1.26 h	0.21 h	1.32 h	0.52 h
Highbay occupancy	70.28%	1.29 pp	70.13%	0.93 pp	70.16%	1.14 pp
Buffer storage occupancy	37.7%	24.2 pp	27.2%	0.53 pp	26.9%	0.95 pp
Truck punctuality	-0.62 h	0.12 h	-0.68 h	0.14 h	-0.65 h	0.23 h
Order ready wait time	0.089 h	0.0021 h	0.091 h	0.0018 h	0.089 h	0.0020 h
Crane retrieval time	0.44 h	0.05 h	0.45 h	0.06 h	0.46 h	0.04 h
Buffer storage retrieval time	0.83 h	0.68 h	0.47 h	0.02 h	0.49 h	0.04 h
Buffer storage to bufferlane time	0.107 h	0.009 h	0.112 h	0.008 h	0.112 h	0.004 h
Trucks loaded per dock per hour	0.327	0.025	0.345	0.013	0.377	0.012
Outbound dock assignment wait time	4.78 h	9.96 h	0.026 h	0.033 h	0.052 h	0.051 h

E.2.1. Configuration 1

In figure E.5, shown is that the buffer storage occupancy is far from stable. Showing a severe bottleneck in this model namely that the buffer storage is constantly filling up with pallets and outbound trucks are loaded slowly.




Figure E.5: Buffer storage occupancy stabilisation, scenario 2, configuration 1

The raw per-replication values for the KPIs of this configuration are shown in figure E.3

Table E.3: Raw per replication KPI values, scenario 2, configuration 1

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	96.36	97.06	99.59	96.49	96.66	97.23	1.40
Total loading time	h	63.33	46.39	63.90	39.29	68.89	56.36	12.78
Planned pickup to actual departure	h	62.68	45.44	63.29	38.54	67.90	55.57	12.79
Trucks loaded per dock per hour	trucks/dock/h	0.221	0.244	0.237	0.245	0.221	0.234	0.012
Highbay occupancy	%	70.99	69.38	66.35	70.00	68.91	69.13	1.74
Buffer storage occupancy	%	157.29	133.34	166.78	118.18	173.41	149.80	23.30
Truck punctuality	h	-0.65	-0.95	-0.61	-0.75	-0.98	-0.79	0.17
Order ready wait time	h	0.089	0.089	0.091	0.090	0.087	0.089	0.0015
Crane retrieval time	h	1.15	2.43	6.53	1.46	1.74	2.66	2.21
Buffer storage retrieval time	h	2.89	2.03	2.52	2.29	2.91	2.53	0.38
Buffer storage to bufferlane time	h	0.069	0.085	0.093	0.090	0.112	0.090	0.016
Outbound dock assignment wait time	h	54.46	37.52	46.58	32.66	58.66	45.98	10.98

E.2.2. Configuration 2

In figure E.6, shown is that the buffer storage occupancy is far from stable. Showing a severe bottleneck in this model namely that the buffer storage is constantly filling up with pallets and outbound trucks are loaded slowly.




Figure E.6: Buffer storage occupancy stabilization, scenario 2, configuration 2

The raw per-replication values for the KPIs of this configuration are displayed in table E.4.

Table E.4: Raw per replication KPI values, scenario 2, configuration 2

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	85.16	83.33	84.78	85.26	83.65	84.44	0.89
Total loading time	h	50.22	48.03	54.91	46.06	65.65	52.98	7.81
Planned pickup to actual departure	h	49.25	47.38	54.42	45.27	64.87	52.24	7.83
Trucks loaded per dock per hour	trucks/dock/h	0.243	0.236	0.239	0.250	0.212	0.236	0.014
Highbay occupancy	%	69.18	69.51	70.17	71.03	68.32	69.64	1.02
Buffer storage occupancy	%	139.59	102.47	143.68	120.05	124.65	126.09	16.49
Truck punctuality	h	-0.97	-0.65	-0.49	-0.80	-0.78	-0.74	0.18
Order ready wait time	h	0.086	0.081	0.085	0.086	0.083	0.084	0.0023
Crane retrieval time	h	0.498	0.570	0.553	0.393	0.538	0.510	0.071
Buffer storage retrieval time	h	2.47	2.05	2.84	2.33	1.44	2.23	0.53
Buffer storage to bufferlane time	h	0.101	0.071	0.116	0.107	0.118	0.103	0.019
Outbound dock assignment wait time	h	40.57	34.99	46.71	36.37	42.03	40.13	4.68

E.2.3. Configuration 3

In figure E.7, still showing that the buffer storage occupancy is far from stable.




Figure E.7: Buffer storage occupancy stabilisation, scenario 2, configuration 3

The raw per-replication values for the KPIs of this configuration are shown in figure E.5.

Table E.5: Raw per replication KPI values, scenario 2, configuration 3

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	85.22	81.08	86.67	83.70	81.73	83.68	2.36
Total loading time	h	55.10	37.11	93.12	17.95	71.55	54.96	29.24
Planned pickup to actual departure	h	54.47	36.41	92.42	17.03	70.73	54.21	29.29
Trucks loaded per dock per hour	trucks/dock/h	0.222	0.211	0.136	0.284	0.200	0.211	0.053
Highbay occupancy	%	68.49	72.36	69.29	70.86	72.08	70.62	1.70
Buffer storage occupancy	%	145.92	88.16	156.63	56.99	130.80	115.70	41.90
Truck punctuality	h	-0.63	-0.69	-0.70	-0.91	-0.82	-0.75	0.11
Order ready wait time	h	0.085	0.081	0.075	0.085	0.082	0.082	0.004
Crane retrieval time	h	0.569	0.362	0.346	0.421	0.378	0.415	0.090
Buffer storage retrieval time	h	2.97	1.39	0.60	1.32	1.82	1.62	0.88
Buffer storage to bufferlane time	h	0.115	0.113	0.107	0.103	0.103	0.108	0.006
Outbound dock assignment wait time	h	45.90	25.74	56.56	11.66	52.06	38.38	19.02

E.2.4. Configuration 4




Figure E.8: Buffer storage occupancy stabilisation, scenario 2, configuration 4

Table E.6: Raw per replication KPI values, scenario 2, configuration 4

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	86.57	83.29	82.78	85.25	84.18	84.41	1.53
Total loading time	h	9.12	24.09	11.06	34.54	8.75	17.51	11.42
Planned pickup to actual departure	h	8.36	23.59	10.08	33.64	8.12	16.76	11.42
Trucks loaded per dock per hour	trucks/dock/h	0.325	0.266	0.315	0.235	0.317	0.292	0.039
Highbay occupancy	%	69.13	71.09	70.56	69.50	69.44	69.95	0.84
Buffer storage occupancy	%	39.87	70.06	38.41	75.89	34.17	51.68	19.66
Truck punctuality	h	-0.77	-0.50	-0.98	-0.90	-0.64	-0.76	0.20
Order ready wait time	h	0.090	0.090	0.089	0.086	0.085	0.088	0.0024
Crane retrieval time	h	0.691	0.416	0.463	0.389	0.433	0.478	0.122
Buffer storage retrieval time	h	1.04	1.55	1.45	0.99	0.85	1.17	0.31
Buffer storage to bufferlane time	h	0.099	0.085	0.104	0.099	0.104	0.098	0.008
Outbound dock assignment wait time	h	4.78	16.66	6.73	22.25	3.63	10.81	8.21

E.2.5. Configuration 5




Figure E.9: Buffer storage occupancy, configuration 5, scenario 2

Table E.7: Raw per replication KPI values, scenario 2 configuration 5

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	83.20	80.68	84.57	82.79	84.80	83.21	1.65
Total loading time	h	2.12	2.03	4.51	1.83	28.46	7.79	11.61
Planned pickup to actual departure	h	1.41	1.35	3.95	1.38	27.73	7.17	11.55
Trucks loaded per dock per hour	trucks/dock/h	0.346	0.327	0.336	0.341	0.283	0.327	0.025
Highbay occupancy	%	70.78	71.90	69.20	70.78	68.73	70.28	1.29
Buffer storage occupancy	%	26.33	24.82	30.72	25.69	80.85	37.68	24.24
Truck punctuality	h	-0.71	-0.67	-0.56	-0.44	-0.73	-0.62	0.12
Order ready wait time	h	0.087	0.090	0.087	0.092	0.089	0.089	0.0021
Crane retrieval time	h	0.449	0.405	0.513	0.384	0.457	0.442	0.050
Buffer storage retrieval time	h	0.505	0.508	0.598	0.514	2.045	0.834	0.678
Buffer storage to bufferlane time	h	0.118	0.105	0.103	0.114	0.096	0.107	0.009
Outbound dock assignment wait time	h	0.050	0.011	1.243	0.014	22.575	4.78	9.96

E.2.6. Configuration 6




Figure E.10: Buffer storage occupancy, configuration 6, scenario 2

Table E.8: Raw per replication KPI values, scenario 2 configuration 6

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	83.61	79.69	85.02	85.84	84.19	83.67	2.38
Total loading time	h	2.27	1.63	2.02	2.07	1.71	1.94	0.26
Planned pickup to actual departure	h	1.48	1.17	1.28	1.41	0.95	1.26	0.21
Trucks loaded per dock per hour	trucks/dock/h	0.351	0.323	0.345	0.359	0.345	0.345	0.013
Highbay occupancy	%	70.86	71.35	69.64	69.12	69.66	70.13	0.93
Buffer storage occupancy	%	27.27	26.42	27.28	27.92	27.22	27.22	0.53
Truck punctuality	h	-0.79	-0.46	-0.74	-0.66	-0.76	-0.68	0.14
Order ready wait time	h	0.087	0.092	0.091	0.092	0.092	0.091	0.0018
Crane retrieval time	h	0.480	0.367	0.518	0.473	0.436	0.455	0.057
Buffer storage retrieval time	h	0.479	0.461	0.469	0.509	0.457	0.475	0.021
Buffer storage to bufferlane time	h	0.105	0.106	0.112	0.124	0.112	0.112	0.008
Outbound dock assignment wait time	h	0.023	0.002	0.014	0.084	0.007	0.026	0.033

E.2.7. Configuration 7




Figure E.11: Buffer storage occupancy, configuration 7, scenario 2

Table E.9: Raw per replication KPI values, scenario 2 configuration 7

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	84.15	81.34	83.29	82.60	85.11	83.30	1.44
Total loading time	h	2.38	1.75	1.56	1.96	2.20	1.97	0.32
Planned pickup to actual departure	h	1.81	0.80	0.84	1.25	1.88	1.32	0.52
Trucks loaded per dock per hour	trucks/dock/h	0.391	0.366	0.372	0.366	0.391	0.377	0.012
Highbay occupancy	%	69.11	71.09	71.32	70.47	68.81	70.16	1.14
Buffer storage occupancy	%	27.92	25.65	26.72	26.54	27.80	26.92	0.95
Truck punctuality	h	-0.57	-0.96	-0.71	-0.71	-0.32	-0.65	0.23
Order ready wait time	h	0.087	0.087	0.091	0.091	0.090	0.089	0.0020
Crane retrieval time	h	0.491	0.435	0.420	0.422	0.513	0.456	0.043
Buffer storage retrieval time	h	0.469	0.477	0.446	0.503	0.556	0.490	0.042
Buffer storage to bufferlane time	h	0.109	0.107	0.119	0.113	0.114	0.112	0.004
Outbound dock assignment wait time	h	0.073	0.013	0.006	0.040	0.131	0.052	0.051

E.3. Scenario 3

Table E.10: Scenario 3, full KPI detail by experiment

KPI	Configuration 1 -		Configuration 2 + 1 level		Configuration 3 -1 crane		Configuration 4 -1 crane	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Crane utilisation	—	—	—	—	—	—	94.53%	0.28 pp
Total loading time	18.54 h	19.56 h	1.80 h	0.08 h	1.86 h	0.10 h	1.90 h	0.23 h
Planned pickup to actual departure	17.86 h	19.61 h	1.07 h	0.11 h	1.10 h	0.15 h	1.12 h	0.25 h
Highbay occupancy	70.00%	0.43 pp	70.30%	1.19 pp	69.93%	1.68 pp	69.67%	0.85 pp
Buffer storage occupancy	67.9%	49.6 pp	28.4%	0.48 pp	28.1%	1.06 pp	28.2%	0.53 pp
Truck punctuality	-0.68 h	0.17 h	-0.73 h	0.14 h	-0.76 h	0.10 h	-0.78 h	0.13 h
Order ready wait time	0.089 h	0.0012 h	0.091 h	0.0014 h	0.092 h	0.0017 h	0.091 h	0.0013 h
Crane retrieval time	0.15 h	0.01 h	0.14 h	0.01 h	0.22 h	0.02 h	0.82 h	0.14 h
Buffer storage retrieval time	1.10 h	0.71 h	0.49 h	0.02 h	0.47 h	0.03 h	0.48 h	0.02 h
Buffer storage to bufferlane time	0.108 h	0.014 h	0.119 h	0.010 h	0.113 h	0.003 h	0.119 h	0.008 h
Trucks loaded per dock per hour	0.331	0.058	0.382	0.012	0.379	0.013	0.383	0.006
Outbound dock assignment wait time	14.48 h	17.90 h	0.025 h	0.012 h	0.040 h	0.021 h	0.067 h	0.084 h

E.3.1. Configuration 1




Figure E.12: Buffer storage occupancy stabilisation, configuration 1, scenario 3

Table E.11: Raw per replication KPI values, configuration 1, scenario 3

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Total loading time	h	2.89	7.72	2.43	40.96	38.69	18.54	19.56
Planned pickup to actual departure	h	2.24	6.89	1.74	40.15	38.28	17.86	19.61
Trucks loaded per dock per hour	trucks/dock/h	0.386	0.338	0.388	0.273	0.270	0.331	0.058
Highbay occupancy	%	69.31	69.15	71.09	68.99	69.83*	70.00	0.43
Buffer storage occupancy	%	28.00	40.64	27.32	127.16	116.29	67.88	49.59
Truck punctuality	h	-0.65	-0.83	-0.68	-0.82	-0.41	-0.68	0.17
Order ready wait time	h	0.090	0.088	0.087	0.087	0.090	0.089	0.0012
Crane retrieval time	h	0.159	0.141	0.146	0.150	0.138	0.147	0.008
Buffer storage retrieval time	h	0.569	0.626	0.547	1.841	1.914	1.099	0.712
Buffer storage to bufferlane time	h	0.121	0.110	0.122	0.093	0.095	0.108	0.014
Outbound dock assignment wait time	h	0.153	4.105	0.163	34.908	33.079	14.482	17.896

E.3.2. Configuration 2




Figure E.13: Buffer storage occupancy stabilisation, configuration 2, scenario 3

Table E.12: Raw per replication KPI values, configuration 2, scenario 3

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Total loading time	h	1.83	1.90	1.78	1.69	1.81	1.80	0.08
Planned pickup to actual departure	h	1.14	0.95	1.22	0.94	1.10	1.07	0.11
Trucks loaded per dock per hour	trucks/dock/h	0.389	0.385	0.394	0.363	0.381	0.382	0.012
Highbay occupancy	%	69.56	70.05	69.59	72.45	69.83	70.30	1.19
Buffer storage occupancy	%	28.42	28.04	28.94	27.71	28.85	28.39	0.48
Truck punctuality	h	-0.69	-0.95	-0.56	-0.74	-0.71	-0.73	0.14
Order ready wait time	h	0.092	0.089	0.091	0.092	0.089	0.091	0.0014
Crane retrieval time	h	0.150	0.143	0.144	0.129	0.146	0.142	0.008
Buffer storage retrieval time	h	0.510	0.479	0.504	0.454	0.489	0.487	0.022
Buffer storage to bufferlane time	h	0.133	0.117	0.118	0.106	0.121	0.119	0.010
Outbound dock assignment wait time	h	0.034	0.024	0.037	0.005	0.024	0.025	0.012

E.3.3. Configuration 3





Figure E.14: Buffer storage occupancy stabilisation, configuration 3, scenario 3

Table E.13: Raw per replication KPI values, configuration 3, scenario 3

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Total loading time	h	1.76	1.96	1.79	1.99	1.80	1.86	0.10
Planned pickup to actual departure	h	1.13	1.12	0.91	1.33	1.01	1.10	0.15
Trucks loaded per dock per hour	trucks/dock/h	0.372	0.378	0.362	0.395	0.386	0.379	0.013
Highbay occupancy	%	70.72	70.81	70.93	66.89	70.28	69.93	1.68
Buffer storage occupancy	%	28.28	27.56	26.86	29.79	28.16	28.13	1.06
Truck punctuality	h	-0.64	-0.84	-0.88	-0.66	-0.79	-0.76	0.10
Order ready wait time	h	0.093	0.092	0.089	0.091	0.094	0.092	0.0017
Crane retrieval time	h	0.217	0.205	0.222	0.238	0.243	0.225	0.016
Buffer storage retrieval time	h	0.458	0.477	0.457	0.531	0.461	0.477	0.032
Buffer storage to bufferlane time	h	0.116	0.109	0.117	0.112	0.111	0.113	0.003
Outbound dock assignment wait time	h	0.016	0.057	0.039	0.064	0.024	0.040	0.021

E.3.4. Configuration 4

**Figure E.15:** Buffer storage occupancy stabilisation, configuration 4, scenario 3)**Table E.14:** Raw per replication KPI values, scenario 3 configuration 4

KPI	Unit	Rep1	Rep2	Rep3	Rep4	Rep5	Mean	SD
Crane utilisation	%	—	—	—	—	—	94.53	0.28
Total loading time	h	2.10	2.24	1.65	1.80	1.74	1.90	0.23
Planned pickup to actual departure	h	1.19	1.55	0.86	0.89	1.14	1.12	0.25
Trucks loaded per dock per hour	trucks/dock/h	0.158	0.162	0.161	0.168	0.164	0.163	0.004
Highbay occupancy	%	69.31	69.15	71.09	68.99	69.83	69.67	0.85
Buffer storage occupancy	%	27.38	28.66	28.05	28.44	28.48	28.20	0.53
Truck punctuality	h	-0.91	-0.69	-0.78	-0.92	-0.60	-0.78	0.13
Order ready wait time	h	0.090	0.091	0.093	0.091	0.092	0.091	0.0013
Crane retrieval time	h	0.867	0.669	0.649	0.997	0.938	0.824	0.144
Buffer storage retrieval time	h	0.466	0.474	0.476	0.484	0.476	0.475	0.007
Buffer storage to bufferlane time	h	0.114	0.109	0.124	0.129	0.116	0.119	0.008
Outbound dock assignment wait time	h	0.074	0.208	0.003	0.041	0.011	0.067	0.084