From Forklifts to AMRs: A Simulation-Based Approach to Support Outbound Loading Decisions

João Tomás Pereira Cardoso

Master's Dissertation

Supervisor at FEUP: Prof. Gonçalo Figueira

U.PORTO

FEUP FACULDADE DE ENGENHARIA UNIVERSIDADE DO PORTO

Master in Industrial Engineering and Management

2026-06-25

Resumo

A adoção de Robôs Móveis Autónomos (AMRs) para transporte interno é uma possibilidade cada vez mais considerada em operações de armazém. No entanto, as implicações de substituir empilhadores operados por humanos por AMRs em operações de carregamento outbound não são imediatas, dependendo da interação entre a dimensão da frota, a atribuição de tarefas, o comportamento de carregamento das baterias, a disponibilidade das docas e a composição das cargas.

Esta dissertação desenvolve um modelo de apoio à decisão baseado em simulação para avaliar esta transição na operação de carregamento outbound de um parceiro industrial, comparando uma configuração existente baseada em empilhadores (AS-IS) com uma configuração proposta baseada em AMRs (TO-BE). Foi implementado em FlexSim um modelo de simulação de eventos discretos do processo de expedição outbound, utilizado como ferramenta de comparação de cenários e não como uma réplica preditiva de um dia operacional específico. Como as chegadas dos camiões e os tempos operacionais são estocásticos, os cenários foram comparados através de repetições independentes, utilizando valores esperados e intervalos de confiança. Ambas as configurações foram avaliadas através de um desenho experimental estruturado, composto por cinco blocos experimentais e um cenário adicional de teste de esforço. Os blocos abrangem uma comparação base da dimensão da frota, uma política de AMR auxiliar em que um segundo AMR pode apoiar uma carga em curso, padrões de chegada estacionários e não estacionários, composições de carga mais exigentes e um cenário de sensibilidade a elevada procura. O cenário de teste de esforço avalia o comportamento do sistema quando todas as ordens de carga são libertadas no tempo zero da simulação, removendo a cadência de chegada como fator limitativo.

Os resultados indicam que a substituição de empilhadores por AMRs é operacionalmente viável, mas exige uma frota aproximadamente entre uma vez e meia e duas vezes e meia maior do que a frota de empilhadores que substitui, aumentando a dimensão necessária da frota em cenários de maior procura e maior complexidade de carregamento. A política de AMR auxiliar transforma capacidade adicional da frota em redução do tempo de carregamento por camião, permitindo que a configuração baseada em AMRs se aproxime do desempenho de referência baseado em empilhadores, embora o seu benefício estabilize quando o sistema dispõe de capacidade suficiente para evitar filas persistentes. A gestão de bateria não limitou o desempenho nas configurações competitivas testadas, considerando a lógica operacional implementada. A principal contribuição é um modelo reutilizável de apoio à decisão baseado em simulação para avaliar configurações de carregamento outbound sob condições controladas e representativas.

Abstract

The adoption of Autonomous Mobile Robots (AMRs) for internal transport is an increasingly common consideration in warehouse operations. However, the performance implications of replacing human-operated forklifts with AMRs in outbound loading are not straightforward, as they depend on the interaction between fleet size, task assignment, charging behaviour, dock availability and load composition.

This dissertation develops a simulation-based decision-support model to evaluate this transition in the outbound loading operation of an industrial partner, comparing an existing forklift-based configuration (AS-IS) with a proposed AMR-based configuration (TO-BE). A discrete-event model of the outbound expedition process was implemented in FlexSim and used as a scenario-comparison tool rather than as a predictive replica of a specific operating day. Because arrivals and operating times are stochastic, scenarios were compared across independent replications using expected values and confidence intervals. Both configurations were evaluated under a structured experimental design composed of five experimental blocks and one additional stress-test scenario. The blocks cover a baseline fleet-size comparison, an auxiliary AMR policy in which a second AMR may assist an ongoing load, stationary and non-stationary arrival patterns, more demanding load compositions, and a high-demand sensitivity scenario. The stress-test scenario evaluates system behaviour when all load orders are released at simulation time zero, removing the arrival cadence as a limiting factor.

The results indicate that substituting forklifts with AMRs is operationally viable, but requires a fleet roughly one and a half to two and a half times larger than the forklift fleet it replaces, with the required size increasing under higher demand and greater loading complexity. The auxiliary AMR policy converts additional fleet capacity into reduced per-truck loading time, allowing the AMR-based configuration to approach the forklift-based benchmark, although its benefit saturates once the system has sufficient capacity to avoid sustained queues. Battery management did not constrain performance in the tested competitive configurations under the implemented operational logic. The main contribution is a reusable simulation-based decision-support model for assessing outbound loading configurations under controlled and representative conditions.

Agradecimentos

I would like to express my gratitude to the Faculty of Engineering of the University of Porto (FEUP) for the education and the opportunities provided throughout these years, and to INESC TEC for hosting this work and providing the conditions that made it possible.

I am deeply grateful to my academic supervisor, Professor Gonçalo Figueira, for his guidance, availability, and valuable insight throughout this dissertation.

I would also like to thank Eng. Romão Santos and Eng. Henrique Piqueiro for their support and contributions along the way, and my company supervisor, Professor António Galrão Ramos, for his guidance and for laying the foundations that made this project possible.

To my parents, I am thankful for the constant support and for the opportunity to pursue this path; none of this would have been possible without them. I am also grateful to my brother, who has always been someone I can rely on, a constant presence to share ideas with and to turn to in moments of doubt, and whose support and perspective I deeply value.

Finally, I would like to thank my friends, who have always encouraged me and pushed me to challenge myself and to grow.

"All models are wrong, but some are useful"

George E. P. Box

Contents

Acronyms and Symbols

AMRs	Autonomous Mobile Robots
DES	Discrete Event Simulation
KPI	Key Performance Indicator
RFID	Radio-Frequency Identification
RMFS	Robotic Mobile Fulfillment System

List of Figures

List of Tables

Chapter 1

Introduction

Warehouse intralogistics is a central determinant of performance in modern supply chains, where the efficiency of internal material handling directly influences throughput, cost, service levels and operational responsiveness. Within this domain, outbound operations, including the preparation and loading of goods onto outbound vehicles, are particularly relevant because they concentrate the interaction between internal transport resources, limited loading infrastructure and externally arriving trucks. When these elements are not properly coordinated, local improvements in material handling may not translate into better end-to-end performance, since dock availability, truck waiting time or loading constraints may become the limiting factors.

At the same time, warehousing operations are under increasing pressure to improve throughput, reduce dependence on manual handling and adopt more flexible automation technologies. Labour availability reinforces this pressure: a 2024 survey of 1,000 supply chain and logistics decision-makers reported that 76% of respondents were experiencing notable workforce shortages, with 56% identifying warehouse operations as one of the functions most affected by resource shortages (Descartes Systems Group, 2024). In parallel, mobile robotics has become increasingly relevant in logistics applications. According to the International Federation of Robotics, professional service robot sales reached almost 200,000 units in 2024, a 9% year-on-year increase, and transportation and logistics was the largest application class, with 102,900 units sold, corresponding to a 14% year-on-year increase. This category mainly includes mobile robots for the transport and handling of goods (International Federation of Robotics, 2025). These developments help explain why many operators are considering the replacement or augmentation of conventional forklift fleets with Autonomous Mobile Robots (AMRs) for repetitive internal transport activities.

However, the performance implications of such a transition are not straightforward in an outbound loading context. Although AMRs offer potential benefits in terms of automation, traceability, standardised execution and integration with digital management systems, their effective contribution depends on several interdependent operational decisions. Fleet size, task assignment, charging behaviour, dock availability and load composition jointly determine whether an AMR-based operation can achieve performance comparable to a forklift-based one. As highlighted in the AMR planning and control literature, the performance of robotized intralogistics systems depends not only on the technology adopted, but also on the configuration and operating rules under which it is deployed (Azadeh et al., 2019; Fragapane et al., 2021).

This dissertation was developed in collaboration with an industrial partner operating an outbound expedition facility, where finished goods are loaded onto trucks for dispatch. In the current operation, truck loading is performed by human-operated forklifts, and loading execution relies substantially on the experience and judgement of the operators. The industrial partner identified the replacement of forklifts with AMRs as a candidate operational change, motivated by the prospect of greater automation and more controlled resource utilisation in the expedition area. This raised a practical question that motivates the present work: under what conditions, and with what fleet size, could an AMR-based configuration achieve performance comparable to the existing forklift-based operation? Evaluating this question directly on the real system would be costly and disruptive, which makes a simulation-based approach particularly suitable.

Existing literature provides relevant insights into AMR fleet sizing, dispatching and charging behaviour, as well as into dock scheduling, truck loading and loading feasibility constraints. Nevertheless, these streams are often developed separately. AMR studies are frequently centred on robotic mobile fulfilment or order-picking systems, where the workload structure and operating objectives differ from outbound pallet loading. Conversely, dock scheduling and loading studies often treat the internal transport layer as an external input, without explicitly modelling how transport resource availability affects the sequence and timing of loads reaching the dock area. This limits the direct applicability of existing results to a forklift-to-AMR transition in an outbound loading operation where transport resources, dock capacity, loading requirements and battery behaviour interact dynamically. This dissertation addresses this gap by evaluating the forklift-to-AMR transition within a single integrated simulation model of the outbound loading operation.

1.1 Problem Description and Objectives

The problem addressed in this dissertation is the comparative evaluation of two configurations of the same outbound loading operation: an AS-IS configuration based on human-operated forklifts, and a TO-BE configuration based on AMRs. The central difficulty is that the relative performance of the two configurations depends on the interaction of several operational factors that cannot be evaluated independently, such as fleet size, the different travel and handling characteristics of forklifts and AMRs, task-assignment policy, battery and charging behaviour, dock availability, and the composition of the loads to be handled. The question of how many AMRs would be required to match the forklift-based operation therefore has no single answer; it depends on the operating conditions under which the comparison is made.

The main objective of this work is to develop a simulation-based decision-support model that enables this comparison to be carried out under controlled and representative conditions. More specifically, the dissertation aims to:

- Represent the outbound loading process for both the AS-IS and TO-BE configurations within a single simulation model;
- Characterise the performance gap between the forklift-based and AMR-based configurations across different workload levels;
- Evaluate the effect of an auxiliary AMR assignment policy on the competitiveness of the AMR-based configuration;
- Assess the robustness of the comparison under more demanding operating conditions, including temporally concentrated arrivals and load compositions with a higher share of complex pallet layouts.

The aim is not to optimise the loading process or to develop new routing or scheduling algorithms, but to provide a structured basis for comparing resource configurations and supporting the associated operational decision.

1.2 Methodological Approach

The work adopts a discrete-event simulation (DES) approach, implemented in FlexSim, to represent the outbound loading process and to compare the two configurations. The model is explicitly conceived as a decision-support tool for evaluating alternative scenarios rather than as a predictive replica of a specific operational day: its validity rests on the internal consistency of the comparison between equivalently parameterised configurations, not on the exact reproduction of historical throughput.

The simulation study followed the main phases of a simulation project, including operational characterisation, data collection and treatment, workload construction, conceptual modelling, implementation in FlexSim, verification of the process logic, plausibility validation of the model behaviour, and replicated scenario experimentation. Both configurations are represented within a single model, in which the active configuration and its parameters are selected through the simulation experimenter. This ensures that differences in performance can be attributed to the loading resource configuration rather than to differences in the experimental context.

The comparison is structured through an experimental design organised into sequential blocks, each isolating a specific dimension of analysis: a baseline fleet-size comparison, the evaluation of the auxiliary AMR policy, a test of robustness under non-stationary arrivals, a sensitivity analysis to more complex load compositions, and an assessment under substantially higher demand, complemented by a stress-test scenario.

Because truck arrivals and operational times are subject to variability, the outbound process is inherently stochastic, and a single deterministic calculation cannot capture the resulting performance distribution. Each configuration is therefore evaluated over multiple independent replications, allowing performance to be assessed in terms of expected values and the associated uncertainty, through a set of system-level, resource-level, and energy-related measures.

1.3 Dissertation Structure

The remainder of this dissertation is structured to guide the reader from the conceptual background of the problem to the simulation results and final conclusions. Chapter 2 reviews the research streams that support the study, including simulation for operations decision support, AMR-based material handling systems, fleet configuration and operating rules, and outbound dock and loading operations. Building on this background, Chapter 3 presents the case study by describing the outbound expedition process, the AS-IS and TO-BE configurations, the representative load typologies, and the performance indicators used to frame the analysis. Chapter 4 then details the simulation methodology, covering the model architecture, operational parameters, input data, arrival process, output measurement, verification, and experimental design. Chapter 5 presents and discusses the results obtained from the five experimental blocks and the stress-test scenario, before synthesising the main findings and their operational implications. Finally, Chapter 6 summarises the conclusions of the dissertation, highlights the contribution of the work, and outlines directions for future research.

Chapter 2

Literature Review

This chapter reviews the research streams that support the simulation study developed in this dissertation. It begins by framing the outbound logistics context and the role of simulation as an analytical tool for evaluating operational alternatives under variability (Section 2.1), before introducing the methodological concepts relevant to modelling warehouse operations with discrete-event and resource-specific behavioural logic (Section 2.2). Sections 2.3 and 2.4 then address robotized intralogistics systems, covering AMR technology, fleet configuration, operating rules, and charging behaviour as interdependent decisions whose performance implications require system-level evaluation. Section 2.5 turns to outbound dock operations, truck loading, and the physical constraints governing how loads are placed inside vehicles. The chapter closes in Section 2.6 by synthesising the limitations identified across these streams and positioning the research contribution. The review prioritizes peer-reviewed journal articles and well-established conference proceedings for the domain-specific streams, selected for their methodological relevance and proximity to the problem studied.

2.1 Warehouse outbound operations

Outbound logistics in warehouses comprises a set of tightly coupled activities that determine truck turnaround time and shipping throughput. Typical processes include order release from upstream storage areas, buffering and staging near dock doors, internal transport to the shipping area, and the physical loading of trucks. These processes interact through shared resources - dock doors, staging capacity, and handling equipment - and through time-dependent constraints such as truck arrival and departure schedules, carrier time windows, and limited buffer space. Local improvements in a single activity do not necessarily translate into improved end-to-end performance: a faster internal transport resource, for instance, may simply increase congestion at staging if dock availability becomes the binding constraint. Evaluating outbound performance therefore requires a system-level perspective that captures the interactions among these operational layers simultaneously.

Warehousing operations are increasingly subject to pressure for higher throughput, greater flexibility, and reduced reliance on manual labour, motivating the adoption of robotized intralogistics systems and systems based on AMRs (Azadeh et al., 2019; Fragapane et al., 2021). In response, many distribution centres are exploring the replacement or augmentation of conventional forklift fleets with AMRs, particularly for repetitive internal transport tasks where resource availability, task execution and system throughput can be monitored and controlled more systematically. However, the performance implications of such a transition in an outbound loading context are not straightforward. The effective throughput of an AMR-assisted operation depends not only on fleet size, but also on how robots are assigned to tasks, how their charging needs are managed, and how they interact with dock availability and loading requirements. These interdependencies mean that the transition from an existing forklift-based AS-IS operation to a proposed AMR-based TO-BE operation cannot be evaluated by reasoning about individual components separately.

Understanding this transition requires a decision support framework capable of representing the full outbound loading cycle - internal transport, dock allocation, staging, and truck loading - under both operating modes, and of generating performance metrics that reflect interactions across the operation rather than local efficiency. A DES perspective is particularly suitable for this purpose, since it can represent stochastic arrivals, shared resources, competing policies, and sequential constraints simultaneously, while also allowing resource-specific behaviours to be incorporated where relevant (Law, 2015). Section 2.2 therefore reviews the main simulation paradigms used in operations decision support and discusses how verification, validation, and scenario-based experimentation support rigorous simulation studies.

2.2 Simulation methods for operations decision support

Building on the system-level perspective introduced in Section 2.1, this section reviews the simulation concepts required to support an operations decision framework for outbound warehouse systems. The focus is not on simulation as a generic modelling taxonomy, but on the methodological elements that are relevant to this dissertation: DES as the primary modelling backbone, agent-like behavioural logic when resources follow state-dependent rules, and the verification, validation, and experimentation procedures required to make simulation results credible. In this context, simulation supports the testing of alternative configurations and control rules under realistic variability, while maintaining traceability between modelling assumptions, system mechanisms, and performance outcomes.

A DES perspective represents the system as a set of entities and resources whose state changes at discrete points in simulated time when events occur (Law, 2015). Typical events include arrivals, service starts and completions, resource seize and release, and transfers between processing stages. This event-driven representation is particularly suited to outbound operations naturally described as process flows and queues, such as staging buffers, dock door occupation, loading tasks, and internal transport requests. In these contexts, DES provides a coherent backbone for modelling time-dependent resource contention and operational rules, while enabling systematic experimentation with fleet size, priority rules, and other decision variables (Banks, 1998; Law, 2015).

Agent-based modelling (ABM) complements this view when individual entities exhibit autonomous behaviour and local interactions are central to system performance (Law, 2015; Macal and North, 2005). In intralogistics, an agent-based perspective is useful when mobile resources make decentralised decisions based on their own state and situational awareness, such as selecting tasks, negotiating access to shared space, or adapting behaviour under resource conflicts. In practice, many agent-based models still evolve through discrete changes in time and can be implemented with event-driven logic, but the modelling emphasis shifts from process flow structure to the behavioural rules and interaction mechanisms that produce system-level outcomes (Law, 2015).

For complex operations problems, combining paradigms can be more effective than committing to a single approach. Hybrid simulation is defined as the combination of two or more of DES, system dynamics, and agent-based simulation, and has shown strong growth over the last two decades, with increasing relevance for operational research applications (Brailsford et al., 2019). Brailsford et al. (2019) organise the hybrid simulation literature around a simulation lifecycle framework and highlight the need for stronger support for conceptual modelling and validation in multi-method studies. Under a hybrid logic, a DES backbone can represent outbound process stages and capacity constraints, while agent-like decision rules represent resource behaviour, movement coordination, and interaction dynamics at the level of detail required by the study objectives (Brailsford et al., 2019; Macal and North, 2005).

For simulation to support credible decision-making, model quality must be ensured through verification, validation, and a disciplined experimentation protocol. Verification assesses the correctness of the model implementation against its specification, focusing on the absence of logic and coding errors, while validation concerns whether the model represents the real system with sufficient accuracy for the intended purpose (Kleijnen, 1995; Sargent, 2010). Because simulation models are abstractions, validity is not absolute: it is assessed relative to the questions the model is meant to answer and the performance measures of interest (Sargent, 2010; Law, 2015). A model may be adequate for comparing alternative operating policies under the same assumptions, yet inadequate for predicting absolute cycle times or resource utilisation if key mechanisms or input data are missing. Credibility therefore depends not only on quantitative checks but also on transparency of assumptions and systematic involvement of domain experts during model formulation and review (Law, 2015).

In practice, verification is supported by modular development, trace-based debugging, and checks of intermediate outputs at the level of individual subflows (Kleijnen, 1995). Validation can combine face-validity assessments with subject-matter experts, comparisons between simulated and observed data when an existing system is available, and sensitivity analysis to test whether the model behaves plausibly under parameter variation (Kleijnen, 1995; Sargent, 2010). Scenario-based experimentation then requires a careful design of experiments, including the selection of decision variables such as fleet size, operating rules, and dock resource levels, the definition of performance measures such as completion time, dock occupation, and resource utilisation, and the use of replications to quantify output variability (Law, 2015). Where the modelled system includes transient effects, warm-up treatment and steady-state considerations become important for unbiased estimation of long-run performance measures (Law, 2015). In multi-method models, additional care is needed because variability arises not only from stochastic inputs but also from the way different modelling paradigms are coupled and synchronised (Brailsford et al., 2019).

The methodological framework described above, based on DES as the primary modelling backbone, complemented by resource-specific behavioural logic, and disciplined through verification, validation, and scenario-based experimentation, provides the foundation for the simulation model developed in this dissertation. The following sections review the substantive literature on the operational systems and decisions that this model represents: beginning with the robotized and autonomous systems that characterise the TO-BE scenario under study.

2.3 Intralogistics Automation with AMRs

Robotized intralogistics systems have expanded rapidly in response to e-commerce growth and the need for higher throughput, greater operational flexibility, and reduced reliance on manual labour. Reviews of the field synthesize the main technology families, including mobile-robot-based fulfilment systems, automated storage and retrieval solutions, and robotized sorting systems, while consistently highlighting that operational performance depends strongly on the control policies used to coordinate resources and material flows, rather than on technology choice alone (Azadeh et al., 2019).

Among these technologies, AMRs are increasingly adopted for transport-intensive intralogistics tasks. Compared with more traditional fixed-path automated transport systems, AMRs can use onboard sensors and maps to navigate autonomously and adapt their paths within predefined operating zones, without requiring a fixed physical guide path. This capability makes them particularly attractive in operations where transport demands are variable, resource availability needs to be monitored systematically, and operational flexibility is a priority (Fragapane et al., 2021). Fragapane et al. (2021) define AMRs as industrial robots that use a decentralized decision-making process for collision-free navigation to provide a platform for material handling and related services within a bounded area, and propose a planning and control framework covering the main decision areas affected by AMR adoption.

The planning and control framework identifies decisions at different planning horizons. Some choices, such as the degree of control decentralization, are fundamentally strategic, as they define the architecture within which all other decisions are made. Others, such as fleet size, charging infrastructure, and the definition of operating areas, are medium-term configuration decisions that define the range of feasible operating conditions available to the system. At shorter horizons, scheduling, dispatching, path planning, and resource management govern how the available fleet responds to demand and resolves conflicts over shared resources. A key insight from the framework is that these levels are not independent: configuration choices constrain and condition the effectiveness of operational control policies, and the performance consequences of any given configuration can only be evaluated by considering the operating rules applied within it (Fragapane et al., 2021).

In practice, many distribution centres currently rely on conventional forklifts for internal transport tasks such as moving pallets from storage areas or conveyor systems to loading docks. Forklifts are well-understood and versatile, but their throughput depends directly on operator availability, shift patterns, and human factors. The transition to an AMR-assisted operation introduces configuration choices and operating conditions that have no direct equivalent in forklift-based systems: fleet size, charging infrastructure, task assignment logic, and battery management collectively shape effective robot availability and transport capacity in ways that cannot be inferred from prior forklift-based experience. Evaluating this transition therefore requires a structured approach capable of representing the interdependencies between these variables and their joint effect on end-to-end outbound performance.

Section 2.4 reviews the literature on the specific decisions that shape AMR fleet performance in this context, covering fleet sizing and capacity planning, operating rules and task assignment, and charging behaviour.

2.4 AMR fleet configuration, operating rules, and charging behaviour

Fleet sizing and capacity planning address the problem of selecting the number and type of mobile robots required to meet a target service level under variable workload. Within the planning and control framework proposed by Fragapane et al. (2021), the number and type of vehicles is a tactical-level decision that must be determined together with related choices such as zoning, service point locations, and resource management, since these configuration decisions collectively define the range of feasible operating conditions within which all shorter-horizon control policies run.

Recent research illustrates this from several angles. Fleet sizing can be embedded in scheduling decisions when robot capabilities depend on the type and equipment of individual vehicles in the fleet. In reconfigurable fleets, where robots vary by load capacity or handling equipment, demand is translated into transportation tasks defined by load type, number of loads, and vehicle type, and tasks are assigned to discrete time periods to minimise investment and running costs (Albert et al., 2025). Capacity planning can also be supported by analytical models that estimate the required number of robots and human resources for alternative operating modes, and assess sensitivity to throughput targets and key system parameters, with simulation used to validate the estimates (Ghelichi and Kilaru, 2021).

A substantial share of fleet-sizing evidence in the AMR literature is developed in robotic mobile fulfilment systems (RMFS), which, despite differing structurally from outbound pallet transport, provide useful insights on the interaction between fleet size, control logic, and system configuration. Simulation studies in high-density RMFS settings enable quantification of trade-offs between space utilisation, throughput time, and energy consumption under different fleet sizes and layout configurations (Li et al., 2021). Modular DES frameworks have similarly shown that the fleet size required to reach peak output depends on the operating mode and conflict-resolution dynamics assumed in the model (Luo and Zhao, 2022). More recent work further suggests that fleet deployment should not be treated as a static configuration. In a simulation-based deep learning study of RMFS, Kim et al. (2026) define each scenario as a specific combination of workload, layout, and operating-strategy conditions, and evaluate alternative fleet sizes within that setting. Their results show that deploying the full available fleet is not always optimal: under some conditions, a smaller fleet may achieve a shorter completion time by avoiding unnecessary robot use, congestion, and lower utilisation. Together, these studies reinforce that fleet sizing is sensitive to modelling assumptions about workload generation, travel-time variability, and operating rules. Therefore, conclusions drawn in one system archetype may not transfer directly to another.

Dispatching and task assignment define how transport tasks are distributed across individual robots over time. Within the Fragapane et al. (2021) framework, scheduling and dispatching are tactic-operational decisions that govern where robots are sent, in what sequence they handle tasks, and how conflicts over shared resources such as charging stations and service points are resolved. The literature addresses both centralised and decentralised approaches. Centralised systems rely on global state information to generate and distribute tasks, which supports system-wide coordination but may become computationally demanding as fleet size and task volume grow. Decentralised approaches distribute decision-making across robots using local information, improving scalability and reactivity at the potential cost of globally suboptimal solutions (Fragapane et al., 2021). Both architectures can be implemented with varying levels of algorithmic sophistication, and the performance differences between them are consistently shown to depend on the operational context, system scale, and the objectives being optimised.

The task planning approach proposed by Bolu and Korçak (2021) illustrates how centralised assignment can be structured and evaluated in an RMFS-type simulation environment: demand is converted into executable tasks and distributed across robots while tracking system state, with the simulation incorporating charging and replenishment decisions alongside order execution. The resource-based task allocation algorithm proposed by Lee (2018) illustrates the decentralised case through simulation of a generic multi-robot delivery mission, not an industrial warehouse setting. The approach extends auction-based task allocation, where robots bid for tasks according to criteria such as proximity and estimated execution cost, by incorporating robot resource levels and potential charging detours into the bid calculation. Both studies are simulation-based and developed in contexts that differ from outbound pallet loading, but they share a consistent insight: task assignment effectiveness depends not only on the assignment rule itself, but also on how operational constraints such as energy availability and resource contention are incorporated into the evaluation. Ignoring these constraints during allocation can bias conclusions about the benefits of a given dispatching policy.

Resource management, and in particular charging behaviour, is a distinctive feature of AMR systems that has received growing attention in the planning and control literature. Fragapane et al. (2021) identify resource management as a tactical decision area, noting that the timing and nature of charging events directly affect robot availability and that different charging policies can produce substantially different throughput outcomes. In practice, this translates into a choice between charging strategies that trigger robots to recharge when battery levels fall below an operational threshold, and strategies that allow robots to charge opportunistically when stations are available and demand is temporarily low. The thresholds and logic governing these behaviours affect both the number of robots effectively available at any given moment and the variability of task completion times, particularly under sustained or uneven workloads.

Across these decision areas, the recurring insight is that fleet performance is the joint outcome of configuration choices and operating rules, not of any single variable in isolation. The interaction between fleet size, task assignment logic, and charging behaviour makes simulation a natural tool for evaluating these trade-offs, since analytical or decomposed approaches cannot easily capture the dynamic coupling between resource availability, task queuing, and charging-related unavailability. Figure 2.1 represents these influence paths schematically, linking the main decision areas to the system dynamics and performance outcomes evaluated in the simulation framework developed in this dissertation.

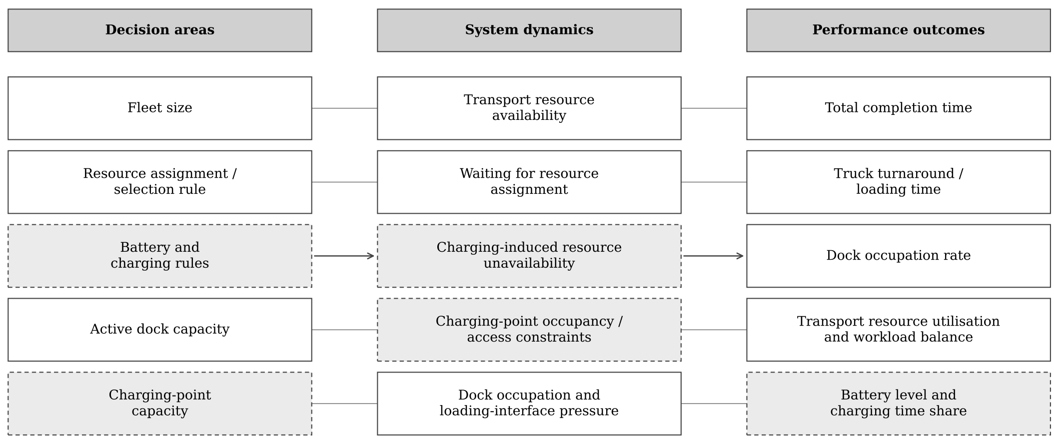

Figure 2.1: Influence diagram linking fleet configuration and operating decisions to outbound performance outcomes. Highlighted elements (dashed borders) apply exclusively to the TO-BE (AMR-based) configuration.

Accordingly, the contribution of this dissertation lies in using simulation to evaluate how selected fleet configurations and control rules, including task assignment logic and charging thresholds, affect outbound performance under representative system conditions, rather than in developing new optimisation algorithms for dispatching or charging. The following section examines the dock and loading interface, where transport resource availability interacts with dock capacity and loading requirements to determine system throughput.

2.5 Outbound dock operations, truck loading, and end-to-end outbound performance

The availability of internal transport resources described in the previous section determines not only how quickly loads can be moved through the warehouse, but also when and in what sequence they reach the dock area and become available for loading. This coupling between the internal transport layer and the dock and loading interface is the focus of this section, which reviews the literature on outbound dock operations, truck loading, and the physical constraints that govern how loads are placed inside vehicles.

Outbound shipping performance is shaped significantly by decisions at the interface between the warehouse and transportation. In practice, this interface is governed by limited dock-door capacity, finite staging space, and time-window constraints imposed by carriers and customers. Even when upstream handling is efficient, poor coordination at the docks can increase truck waiting times, delay departures, and reduce overall throughput. A first stream of work formulates outbound dispatch planning as a joint loading and scheduling problem. Tadumadze and Emde (2022) study a dispatch warehouse where shipments must be assigned to outbound trucks subject to trailer capacities, and trucks must be scheduled at dock doors while respecting time windows and limited terminal resources such as workers and doors. Their model explicitly captures the operational trade-off between dispatching as late as possible to align with just-in-time objectives and ensuring that shipment deadlines are not violated. The study also highlights that resource scarcity at the terminal can directly compromise delivery punctuality.

A complementary perspective focuses on loading strategies when goods are not all available when trucks arrive. Wei and Leung (2011) analyse truck loading under uncertain product readiness, which can create long waiting times and dock blocking. They compare alternative loading strategies that leverage real-time status information enabled by radio-frequency identification (RFID), using simulation to quantify impacts on time in system, truck throughput, and tardiness. Their approach is relevant for outbound areas where synchronisation between upstream preparation and dock operations is a primary driver of performance. Related cross-dock work emphasises how door assignment flexibility and buffering policies affect outbound outcomes under time windows. Bodnar et al. (2017) consider inbound and outbound truck scheduling with mixed-service dock doors and temporary buffering, minimising operational costs driven by handling in temporary storage and tardiness of outbound processing. Although the context differs from a dedicated dispatch warehouse, the study provides useful concepts on how door flexibility and buffering decisions interact with truck time windows and service priorities.

Beyond scheduling and coordination, outbound planning models can fail to represent the physical constraints governing how unit loads are arranged inside vehicles. These constraints include geometric packing, unloading sequence, static and dynamic stability, and weight distribution. Pollaris et al. (2015) systematise this research stream, and the core insight from subsequent work is that sequencing and unloading accessibility interact with stability and axle-weight feasibility throughout a multi-drop route, not only at departure (Pollaris et al., 2016). More recent integrated formulations extend these requirements to cover stability, axle weights, and sequential loading jointly across large real-world instances (Alonso et al., 2017, 2019, 2022; Gajda et al., 2022).

The dissertation does not address the general loading feasibility problem or develop new loading algorithms. However, the pallet positioning and sequencing logic implemented in the simulation model is based on two specific contributions from this stream. Ramos et al. (2016) propose a physical packing sequence algorithm for the container loading problem that enforces static mechanical equilibrium conditions, generating feasible load sequences in a computationally tractable way. Ramos et al. (2018) address load balance in road transport by applying vehicle-specific load distribution constraints that ensure axle-weight limits are respected without sacrificing volume utilisation. These two works provide the algorithmic foundation for the loading representation adopted in the simulation, where pallet positioning inside trucks is represented according to feasibility and balance principles rather than treated as an unconstrained input.

Taken together, these works establish that dock coordination and physical loading constraints are well-studied domains. However, the studies reviewed here typically model the internal transport layer as an external input or treat it in an aggregated way. As a result, dock scheduling, staging, and loading decisions are analysed largely separately from the internal transport dispatching logic and resource availability that determine when and in what sequence loads arrive at the shipping area. This abstraction limits the ability to assess how fleet configuration and operating rules affect end-to-end outbound performance, and motivates the simulation-based approach developed in this dissertation, which represents the transport and dock layers jointly and evaluates their coupled effects on completion time, dock occupation, and resource utilisation.

2.6 Knowledge gap and research contribution

The preceding sections have reviewed the literature across five interconnected areas: simulation methods for operations decision support, AMR systems and the transition from forklift-based operations, fleet configuration and operating rules, outbound dock operations and truck loading, and loading feasibility constraints. Across these streams, the evidence base is substantive and the modelling approaches are well developed. Reviews on AMR planning and control establish a structured framework of decision levels and consistently show that performance depends on the operating conditions and rules applied within a given configuration rather than on technology adoption alone (Azadeh et al., 2019; Fragapane et al., 2021). Work on dock scheduling, truck loading, and physical loading constraints provides solid formulations for resource allocation, time-window management, and loading feasibility (Tadumadze and Emde, 2022; Pollaris et al., 2015).

The limitation that emerges from this review is not an absence of relevant work, but a persistent gap in how these streams are connected. Studies on AMR fleet sizing, dispatching, and charging behaviour are predominantly developed in robotic mobile fulfilment and order-picking contexts, where system objectives, workload patterns, and resource types differ substantially from outbound pallet transport. Conclusions about fleet sizes, effective operating rules, and charging strategies are therefore sensitive to the assumptions and system archetypes from which they are derived, and their direct applicability to an outbound warehouse performing palletised truck loading under dock-capacity constraints is uncertain. As the fleet sizing literature reviewed in Section 2.4 illustrates, results depend on the operating mode, conflict-resolution logic, and demand variability assumed in the model, making cross-context transfer unreliable.

A related limitation concerns the coupling between the transport and dock layers. The dock scheduling and loading literature typically treats the internal transport layer as an external input, assuming that loads arrive at the shipping area in a given sequence and at given times, without representing the operational variability and resource contention that internal transport generates. Conversely, work on AMR planning and control rarely couples fleet configuration and dispatching decisions to the downstream handling constraints that govern how those decisions affect shipping stage feasibility and truck cycle time. Recent work on warehouse-vehicle routing integration shows that treating these operational layers independently can create measurable inefficiencies, and that shared evaluation frameworks are needed to identify policies that improve outbound process performance (Rijal et al., 2023). This decoupling is particularly consequential for the AS-IS to TO-BE transition studied in this dissertation: replacing forklifts with AMRs introduces configuration choices and operating conditions that have no direct equivalent in the existing system, and the performance consequences of those choices propagate through the dock and loading interface in ways that isolated evaluations of either layer cannot fully capture.

A third consideration is the limited availability of structured, parameterised scenario comparisons for change decisions of this type. The question facing an operator considering AMR adoption is not which system is superior in the abstract, but how the transition affects performance under the specific workload conditions, fleet sizes, and operating rules of the existing operation. Most studies in the reviewed literature evaluate systems in isolation or compare design variants within a single archetype, without providing systematic evidence on how performance shifts when the transport technology and operating rules change simultaneously.

The simulation-based decision-support framework developed in this dissertation addresses these gaps from a practical, application-oriented perspective. As represented in Figure 2.1, the framework captures the main decision areas alongside the system dynamics and performance outcomes they shape. By modelling the AS-IS forklift-based configuration and the TO-BE AMR-assisted configuration within a common simulation environment, the framework enables direct, consistent comparison under equivalent loading demands and process conditions. The contribution does not lie in developing new dispatching algorithms, loading optimisation procedures, or traffic control mechanisms. Rather, it lies in using simulation to evaluate how selected fleet configurations and operating rules affect outbound performance under representative process dynamics, providing structured decision support for the transition from AS-IS to TO-BE in a specific outbound logistics context.

Chapter 3

Case Study

This chapter presents the case study that provides the operational basis for the simulation model. Sections 3.1-3.4 introduce the outbound expedition context, describe the process under analysis, and characterise the AS-IS forklift-based configuration and the TO-BE AMR-based alternative. The remaining sections consolidate this description through a process flow representation and define the load typologies, system scope, operational boundaries and performance indicators used to frame the simulation study. Together, these elements establish the empirical and operational assumptions that guide the methodology developed in Chapter 4.

3.1 Operational Context

This study is conducted in collaboration with an industrial partner operating an outbound expedition facility for finished goods distribution. The outbound process encompasses the reception of trucks arriving to collect loads, the assignment of trucks to expedition docks, the preparation and release of pallets from an upstream automated warehouse, their transport to the expedition area, and the physical loading of trucks at the docks. During the considered operating period, trucks arrive to collect loads of varying composition, which are dispatched to different destinations.

For the purposes of this case study, the outbound expedition area is represented through six loading docks. Each dock is served by two conveyors, each consisting of two belt segments, through which pallets are progressively made available at the dock interface as a load order is released. A schematic representation of this layout is provided in Figure A.1 in Appendix A.1.

In the current operational configuration, truck loading is performed by human-operated forklifts, with execution relying on operator experience and judgement. The industrial partner has identified AMR-based loading as a candidate change within a broader automation effort aimed at reducing dependence on manual resource availability in repetitive and safety-sensitive activities, while improving execution standardisation and delivery quality. This study was developed in response to that hypothesis, with the objective of providing a simulation-based evaluation of the two configurations under representative operating conditions.

Building on the loading feasibility and load-balance principles reviewed in Section 2.5, this study uses an existing load scheduling logic, based on the contributions of Ramos et al. (2016, 2018), to generate a loading plan for each load order. The loading plan defines the sequence, position and orientation of each pallet inside the vehicle, accounting for pallet and vehicle dimensions, weight distribution and balance constraints. In the real AS-IS operation, loading decisions do not necessarily follow a pre-computed plan. In this study, however, the same algorithm-generated loading plan is applied in both the AS-IS and TO-BE configurations, supporting a controlled comparison in which loading requirements remain constant across configurations.

3.2 Outbound Expedition Process

The outbound expedition process begins with the arrival of a truck at the facility entrance and concludes with the truck's departure after its load has been fully placed inside the vehicle. The process can be understood as a sequence of operational phases involving truck arrival and dock assignment, travel to the assigned dock, inspection and loading preparation, pallet release to the dock conveyors, loading execution and truck departure.

Trucks arrive at the facility entrance, where they are registered and wait until an expedition dock is assigned. Dock assignment is coordinated at the entrance based on the load order information available and on the operational availability of the docks. The observed elapsed time between entrance registration and dock assignment is interpreted in this study as a waiting and coordination outcome rather than as a fixed process delay. For representation purposes, dock assignment can be modelled either through a fixed-dock logic, in which the dock is pre-determined by the load order data, or through a first-available logic, in which the truck is directed to an available dock. These representations are formalised later as part of the modelling scope.

Following dock assignment, the truck travels from the facility entrance to the assigned dock. This displacement takes approximately ten minutes. Upon arriving at the dock, the truck occupies the docking position and an inspection and validation step is carried out before loading begins. This step verifies the readiness of the truck and its container for loading and lasts approximately one minute. For the purposes of this study, the inspection is assumed to always result in clearance, and the dock remains occupied throughout this period.

Once the truck/load order is associated with a dock, the pallet release process is initiated. Pallets become available progressively at the conveyor entry points associated with the assigned dock. After the load order is released to the conveyor feed process, the first pallet becomes available after approximately five minutes. The remaining pallets are then released progressively to the conveyors according to the pallet feeding logic associated with the load order. Pallets are presented on the belt segments associated with the assigned dock in the sequence defined by the loading plan.

The upstream automated warehouse is therefore represented through timed pallet availability at the conveyors, rather than as a detailed physical process. This representation allows the analysis to focus on the outbound expedition and loading area, where the main operational differences between the configurations under study are expected to arise.

Once pallets begin arriving at the conveyor and the inspection is complete, the transport resource performs repeated loading cycles until all pallets in the load order have been placed inside the truck. Each cycle consists of collecting a pallet from the conveyor pickup point, transporting it to the truck, and placing it in the position defined by the loading plan. Where the loading plan requires a pallet to be rotated before placement, the rotation is carried out as part of that cycle. The transport resource returns to the conveyor after each placement and begins the next cycle.

Loading is considered complete when the last pallet has been placed inside the truck. At that point, the truck departs and the dock becomes available for the next load order. The specific transport resource involved in the loading execution determines the operational characteristics of each cycle. In the AS-IS configuration, this role is performed by forklifts; in the TO-BE configuration, it is performed by AMRs. These two configurations are described in Sections 3.3 and 3.4.

3.3 AS-IS Configuration: Forklift-Based Loading

In the AS-IS configuration, the common outbound process described in Section 3.2 is executed using human-operated forklifts as the loading resource, with each load order assigned to a single forklift during loading. The pre-loading inspection is performed by an operator or supervisor, who verifies the readiness of the truck and container before loading begins. Once clearance is given, pallet handling and placement are carried out by forklift operators. In the model, one pallet is transported per forklift cycle, and pallet rotation is performed whenever required by the loading plan.

As discussed in Section 3.1, real AS-IS loading execution relies on the experience and judgement of forklift operators. In this study, however, the AS-IS configuration is represented using the same algorithm-generated loading plan applied to the TO-BE configuration. This ensures that the comparison focuses on the operational performance of the transport resource rather than on differences in pallet sequencing or positioning requirements.

During the loading phase, a forklift may be actively handling or transporting a pallet, travelling between the truck and the conveyor area, or temporarily idle between loading tasks. The distribution of time across these operational states provides a basis for evaluating resource utilisation and workload balance in the AS-IS configuration.

3.4 TO-BE Configuration: AMR-Based Loading

The TO-BE configuration represents the AMR-based alternative under evaluation, in which Autonomous Mobile Robots replace human-operated forklifts as the loading resource.

In the TO-BE configuration, the pre-loading inspection described in Section 3.2 is performed by an AMR rather than by a human operator or supervisor. The assigned AMR travels to the dock, approaches the truck, and performs the inspection using its onboard camera-based sensing capabilities. The sensing and validation mechanism itself is abstracted in the model, since the study focuses on the operational impact of assigning this task to the AMR rather than on the technical implementation of the inspection system. The inspection duration and outcome are consistent with those assumed in the AS-IS configuration: approximately one minute, always resulting in clearance.

AMRs are assumed to be homogeneous and transport one pallet per loading cycle. The loading requirements remain those defined for the common process, including pallet rotation when required.

The principal operational element introduced by the TO-BE configuration is battery management. AMR battery level decreases with activity, and two charging behaviours coexist: mandatory charging, triggered when the battery level falls below a defined minimum level, and opportunistic charging, in which an AMR may charge when temporarily idle. These behaviours affect resource availability during the loading phase and introduce a dependency between fleet-level battery state and outbound performance that is absent from the AS-IS configuration.

At any point during the simulation, an AMR may be actively handling or transporting a pallet, travelling between the conveyor area and the truck, waiting for a task, or charging. This broader set of operational states provides the basis for evaluating AMR utilisation, charging time share and the impact of battery management on the TO-BE configuration. Routing decisions are not part of the operational problem under study and are therefore not treated as a decision variable.

Both configurations share the same dock layout, conveyor feed logic, truck arrival process, loading requirements and one-pallet-per-cycle handling logic, ensuring that the comparison isolates the effect of the loading resource and its associated operating logic. Table 3.1 summarises the main operational dimensions that distinguish the AS-IS and TO-BE configurations.

Table 3.1: Main operational differences between the AS-IS and TO-BE configurations.

Dimension	**AS-IS**	**TO-BE**
Configuration role	Current baseline configuration	Alternative configuration under evaluation
Loading resource	Human-operated forklifts	Autonomous Mobile Robots (AMRs)
Inspection execution	Operator/supervisor validation	AMR-based validation using onboard sensing
Battery management	Not applicable	Battery level decreases with activity
Charging behaviour	Not applicable	Mandatory and opportunistic charging coexist
Relevant operational states	Handling/transporting, travelling, idle	Handling/transporting, travelling, idle, charging

3.5 Outbound Process Flow Representation

Figure 3.1 presents the flowchart of the outbound expedition process described in Section 3.2, providing a structured visual representation of the sequence of operational phases from truck arrival to truck departure and dock release.

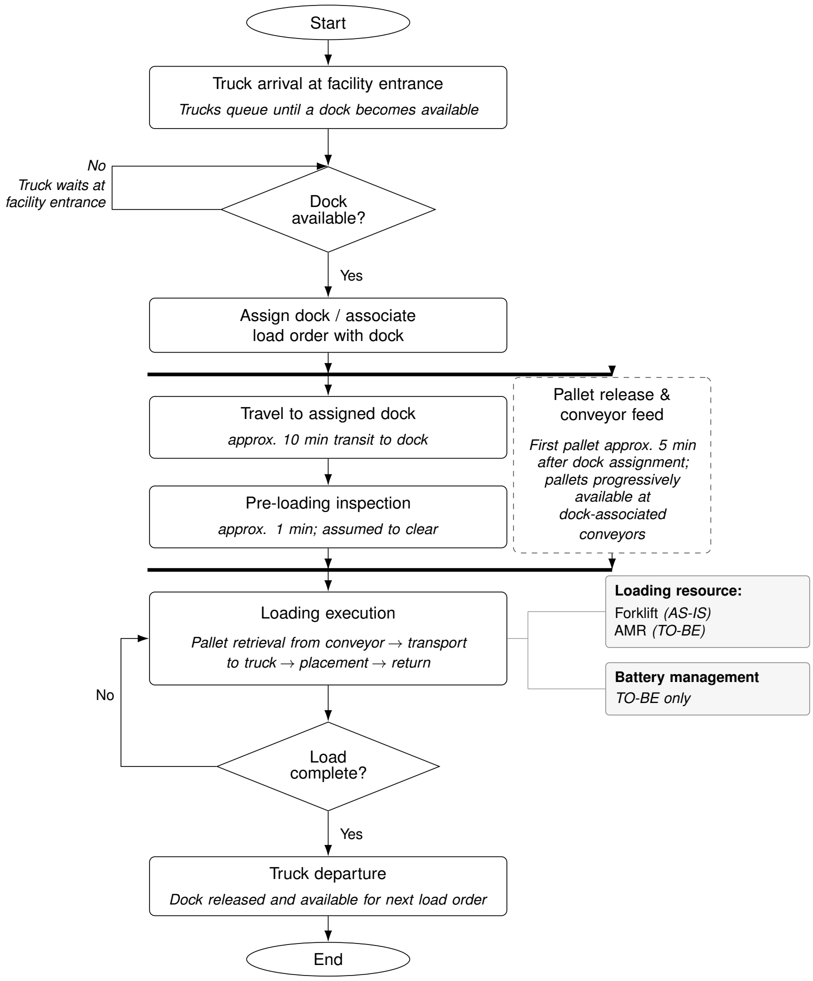

Figure 3.1: Outbound expedition process flowchart.

The flowchart follows a top-to-bottom main sequence, with one structural feature worth highlighting: once a dock is assigned, the truck travel and the pallet release processes are initiated in parallel, reflecting the operational reality in which upstream pallet retrieval does not depend on the truck having reached the dock. The loading execution block is the primary point at which the two configurations differ, the loading resource is a forklift in the AS-IS configuration and an AMR in the TO-BE configuration, while battery management, shown as a lateral annotation, is exclusive to the TO-BE configuration and does not alter the structure of the common process flow.

3.6 Workload Representation and Load Typologies

The case study combines structural data, describing the physical configuration of the expedition area, with workload data, describing the load orders processed by the system. This section focuses on the workload representation, which defines the demand placed on the simulation model.

Each load order is represented as a LoadNum, corresponding to one complete truck loading operation. A LoadNum may contain one or more LoadLines, each corresponding to a distinct product line within the same load order. When a load order includes more than one LoadLine, the pallets belonging to different lines differ in product type or destination, but all pallets within a single LoadNum share the same pallet type. Each LoadNum has an associated loading layout, generated by the load scheduling algorithm, which defines the placement sequence, position and orientation of all pallets across its LoadLines.

The workload used in this study is defined through eight representative load typologies provided by the industrial partner, selected to represent relevant loading configurations of the real operation. Table 3.2 summarises the eight typologies, designated T1 through T8, along with the characteristics that differentiate them operationally.

Table 3.2: Representative load typologies used in the study.

Typology	Vehicle Type	Pallet Type	No. of LoadLines	Total Pallets	Loading Layout	Rotation Required
T1	`TK02`	`A7P`	3	19	`L19`	No
T2	`TK02`	`A7P`	2	24	`L24`	Yes
T3	`TK02`	`A7P`	3	29	`L29`	No
T4	`TK02`	`A7P`	2	31	`L31`	No
T5	`TK11`	`G10P`	1	27	`L27`	No
T6	`TK11`	`S11`	1	21	`L21`	No
T7	`TK11`	`H1P`	1	28	`L28`	Yes
T8	`TK11`	`A7P`	1	33	`L33`	No

The typologies vary in total pallets per load order (from 19 to 33), number of LoadLines (from one to three), and vehicle type (TK02 or TK11), with two of them, T2 and T7, requiring pallet rotation. The loading layouts, identified by the codes L19 through L33, correspond to the algorithm-generated placement plans associated with each typology and remain fixed across all simulation runs. Figures A.2 and A.3 in Appendix A.1 illustrate the difference between load typologies with rotated pallets, as in Typology 2, and typologies without prior pallet rotation, as in Typology 8.

The total number of LoadNums and their distribution across typologies are treated as experimental design choices, defined later in the methodology. What is fixed at this stage is the set of eight typologies and their loading layouts, which define the space of possible load orders in the model.

The use of representative typologies provides controlled demand conditions for comparing the two configurations, rather than reproducing a specific historical period. Detailed operational parameters, such as resource speeds, handling times and battery-related values, are defined later as part of the simulation parameterisation.

3.7 System Scope and Operational Boundaries

This section defines the operational boundaries of the case study and clarifies the elements included in, and excluded from, the outbound expedition process analysed in this study.

System boundaries: The study is limited to the outbound expedition and loading area. The following boundaries define the beginning, end and external interfaces of the process:

- The process starts when a truck arrives at the facility entrance and ends when the dock is released after load completion.
- Activities occurring after dock release, including onward distribution, are outside the scope of the study.
- Upstream warehouse operations are not represented as detailed physical processes; their effect is captured through the timed availability of pallets at the dock-associated conveyors.
- Pallets are transported directly from the conveyor interface to their placement position inside the truck, except when a pallet must be rotated, in which case it may be temporarily set down before final placement.

Truck arrival and dock assignment: Truck movement from the facility entrance to the loading area depends on dock assignment. The following boundaries apply to this stage:

- A truck only proceeds to the loading area after a dock has been assigned to its load order.
- Any elapsed time between truck arrival at the facility entrance and dock assignment is associated with dock availability and operational coordination.
- Once a dock is assigned, it is committed to the corresponding truck and load order. It remains unavailable to other trucks until the loading operation has been completed and the dock is released.
- Trucks that cannot immediately be assigned to a dock remain waiting at the facility entrance. No alternative truck handling arrangements are considered within the scope of the study.

Conveyor feed and pallet availability: The conveyor feed represents the interface between the upstream warehouse and the outbound loading area. The following conditions are considered:

- The pallet release process is initiated once the load order is associated with an assigned dock.
- For each assigned dock, the dock-associated conveyors are dedicated to the corresponding load order during its loading process. Pallets from different LoadNums are not mixed on the conveyors associated with the same dock.
- The sequence in which pallets become available at the conveyors follows the loading sequence defined by the loading plan.
- Pallet availability at the conveyors is assumed not to be the limiting factor of the outbound loading process. This boundary ensures that the analysis remains focused on dock availability, loading resource behaviour and, in the TO-BE configuration, AMR battery and charging constraints.

3.8 Key Performance Indicators

The performance evaluation is structured around five dimensions that together support the comparison between the AS-IS and TO-BE configurations. One of the main objectives of the analysis is to support the assessment of how many AMRs would be required to achieve performance comparable to the forklift-based configuration under different operational conditions. The indicators defined here operationalise that question at system, truck, dock, and resource level.

Because the experimental implementation is based on fixed workload scenarios with stochastically generated arrivals, as detailed in Chapter 4, the primary system-output indicator is the total completion time required to process the full workload, complemented where relevant by additional timing indicators such as the arrival time of the last truck. Table 3.3 presents the main indicators by performance dimension.

Table 3.3: Key performance indicators by dimension.

Performance dimension	**Main indicator(s)**	**Analytical purpose**
System output	Total completion time; loaded trucks over time	Assess workload completion and output progression over time.
Truck flow performance	Truck turnaround time; waiting time for dock assignment	Evaluate truck progression and the impact of dock availability.
Dock performance	Dock utilisation / dock occupation	Assess dock use and identify under-utilisation or congestion.
Loading resource performance	Resource utilisation; average pallet loading cycle time	Compare workload, availability and cycle efficiency of forklifts and AMRs.
AMR battery management (TO-BE only)	Charging time share; battery level / AMR availability	Evaluate the impact of battery and charging behaviour on AMR availability.

Auxiliary metrics support interpretation where relevant, including pallet loading cycle time by rotation requirement, rotation share, completed loads by dock, pallets loaded per resource, load participation by resource and detailed resource-state distributions. These metrics are not treated as primary comparison dimensions but provide additional diagnostic information. The calculation and reporting of all indicators, including their aggregation across replications and confidence-interval treatment, are defined in Chapter 4 as part of the simulation methodology and experimental design.

Chapter 4

Simulation-based Methodology

This chapter describes the simulation-based methodology used to compare the AS-IS and TO-BE outbound loading configurations and to generate the results reported in Chapter 5. It begins by justifying the modelling approach and the selection of FlexSim as the simulation platform in Section 4.1, before presenting the model architecture and process logic in Section 4.2. The following sections define the operational parameters, input data pipeline and arrival process that make the model executable, while Sections 4.6 and 4.7 explain how outputs are collected and how the model was verified. The chapter concludes with the experimental design in Section 4.8, which structures the AS-IS and TO-BE comparison across the scenarios analysed in this dissertation.

4.1 Simulation Approach and Tool Selection

The comparative evaluation of forklift-based and AMR-based loading configurations requires a modelling approach capable of representing the dynamic and resource-dependent nature of the outbound expedition process. The system involves concurrent truck arrivals, contention for dock access, parallel pallet release processes, repeated loading cycles with resource-specific operational characteristics, and, in the TO-BE configuration, battery-driven availability constraints that evolve throughout the simulation horizon. The interaction of these elements produces system-level behaviour that is not tractable through closed-form analytical models. Queueing network models can provide approximate steady-state performance estimates for simplified configurations, but they cannot readily capture the loading sequence constraints imposed by predefined loading plans, the time-varying availability of resources driven by battery management, or the spatial dynamics of resource movement within the expedition area. These structural characteristics make discrete-event simulation, rather than analytical approximation, the appropriate basis for representing the system.

A defining characteristic of the outbound process is that it is subject to variability and uncertainty. Truck interarrival times follow a stochastic arrival process, and operational times such as travel to the assigned dock are represented by probability distributions rather than fixed values. A single deterministic calculation based on average times would therefore fail to capture the resulting performance variability and the risk of congestion during periods of higher demand. DES addresses this limitation by reproducing the stochastic behaviour of the system across multiple independent replications, allowing performance to be characterised through expected values and confidence intervals.

The model developed in this study is designed as a scenario-comparison and decision-support tool rather than a predictive replica of the real operation. The objective is to produce reliable estimates of the relative performance of the AS-IS and TO-BE configurations under representative and controlled conditions, not to reproduce a specific historical operational period with high fidelity. This framing is consistent with the limited availability of disaggregated time-stamped operational data and with the comparative nature of the research question, as discussed in Chapter 3.

Although other simulation tools, such as AnyLogic, Arena or Simio, could also be applied, FlexSim was selected as the simulation platform because it combines discrete-event logic with a three-dimensional physical model, allowing the control logic and spatial layout of the expedition area to be represented in a single environment. A key advantage for this study is the direct connection between Process Flow and 3D objects, which makes it possible to observe, verify and interpret the behaviour defined in the control flow within the spatial representation of the system. Two FlexSim modules are central to this study: the Process Flow module, which structures the event-driven control logic, and the Experimenter module, which supports scenario execution and output collection across replications. Both are detailed in the sections that follow.

The model is primarily a discrete-event simulation, since system state changes occur at discrete events such as truck arrivals, dock assignments and loading completions. At the same time, forklifts and AMRs are represented as mobile resources moving through the 3D environment, maintaining operational states that evolve during the simulation. These resources also respond dynamically to task availability and system conditions, with their behaviour integrated into the event-driven control logic through FlexSim's mobile resource and task-execution mechanisms.

4.2 Model Architecture and Process Logic

The simulation model combines a three-dimensional physical representation of the outbound expedition area with a Process Flow layer that controls the behaviour of trucks, pallets and loading resources. These layers are directly coupled: Process Flow activities trigger actions on 3D objects, while the spatial model defines the travel paths and distances followed by mobile resources.

Both operational configurations, AS-IS and TO-BE, are implemented within a single model file. The active configuration is selected through the parameter `ScenarioType`, set via the FlexSim Experimenter before each run. When `ScenarioType` is set to `AS_IS`, the forklift-based loading logic is activated; when it is set to `TO_BE`, the AMR-based logic is activated instead. This structure ensures that both configurations are evaluated under identical structural and input conditions, while avoiding the need to maintain separate model versions.

4.2.1 Physical Model and 3D Representation

The 3D model represents the outbound expedition area of the case study facility, including the loading docks, conveyor systems, transport resources and truck loading spaces. The spatial layout was configured to reflect the relevant physical characteristics of the real operation, such as dock positions, conveyor lengths and vehicle loading spaces. Table 4.1 summarises the main object types present in the model and their roles.

Table 4.1: Main object types in the 3D simulation model.

Object type	**Role in the model**
Trucks	Entities that enter the system, travel to the assigned dock, are loaded, and depart after completion.
Docks	Fixed loading positions where trucks remain during loading; each dock is associated with two conveyors.
Conveyors / belt segments	Pallet delivery interfaces between the upstream system and the loading area.
Pallets	Entities gradually released to conveyors and transported to the assigned truck by the loading resource.
Forklifts	Mobile loading resources active in the AS-IS configuration.
AMRs	Mobile loading resources active in the TO-BE configuration.
Park / charging points	Idle positions for forklifts and charging positions for AMRs.

Resource movement in the 3D environment is managed through a path network connecting the main operational locations. Mobile resources navigate this network using shortest-path logic, which resolves movement conflicts between resources and fixed infrastructure without requiring manual programming of individual travel paths.

Figure 4.1 shows a view of the 3D model, illustrating the spatial arrangement of the main elements in the expedition area.

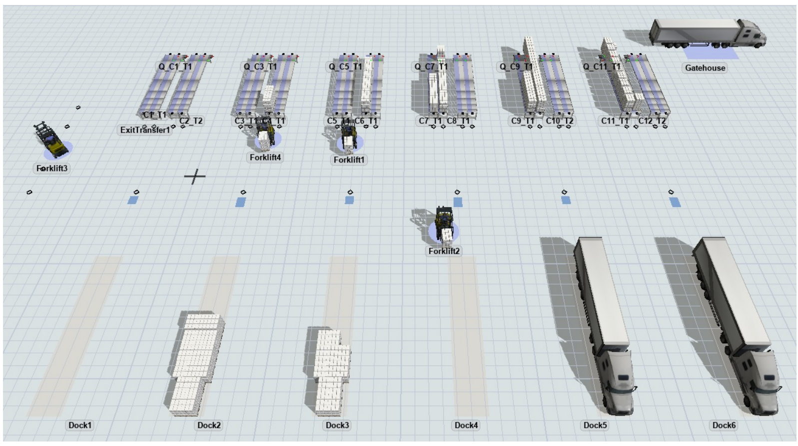

Figure 4.1: Three-dimensional representation of the FlexSim outbound expedition model.

4.2.2 Process Flow Structure and Control Logic

The Process Flow layer implements the control logic of the simulation model through structured blocks that coordinate the movement and processing of entities. In FlexSim, tokens are logical entities that carry labels and trigger actions on 3D objects, lists and global variables. In this model, load-order progression is represented through tokens that move across the Process Flow until loading is complete and the dock is released.

As shown in Figure 4.2, the model receives two main input categories: load order data, stored in the `LoadNumData` table, and scenario parameters defined in the FlexSim Experimenter. The construction of the `LoadNumData` input table is described in Section 4.4, while the main experimenter parameters are summarised in Table 4.3.

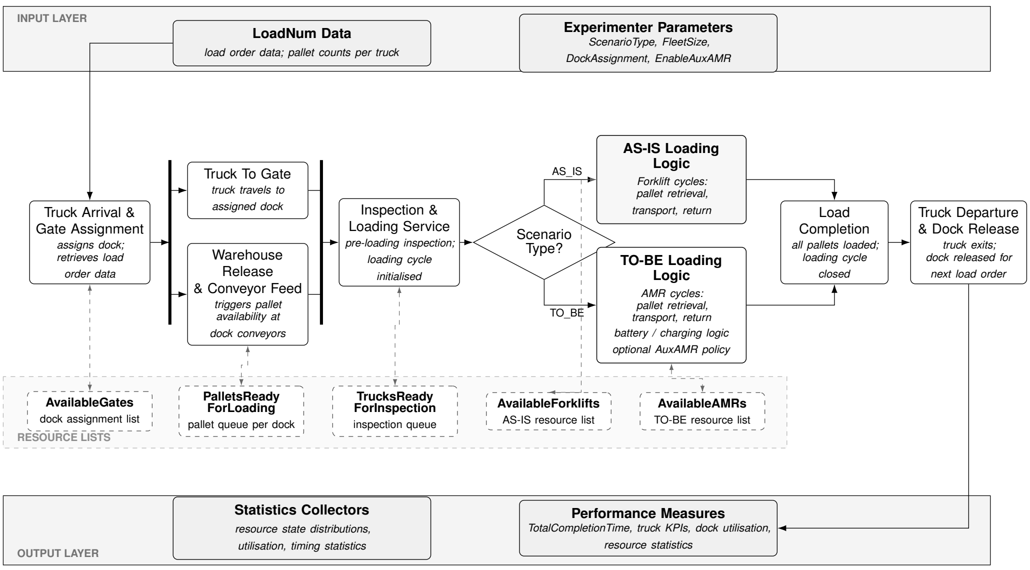

Figure 4.2: Architecture of the FlexSim Process Flow control model.

The common Process Flow logic handles the phases that are independent of the active configuration. A token is generated for each arriving truck, associated with the next load order, and assigned to a dock. Dock assignment then triggers two concurrent processes shown in Figure 4.2: 'Truck To Gate', which moves the truck through the 3D path network to the assigned dock, and 'Warehouse Release & Conveyor Feed', which begins the timed release of pallets to the associated conveyor segments. The two processes are coordinated before loading begins, ensuring that the truck is positioned and that pallets are available for retrieval. Throughout these phases, coordination between tokens, 3D objects and resources is maintained through a set of supporting lists, summarised in Table 4.2; a detailed view of the common Process Flow implementation is provided in Figure B.2 in Appendix B.

Table 4.2: Main supporting lists used for entity and resource coordination.

List	Role
AvailableGates	Loading docks available for assignment to an incoming load order.
TrucksReadyForInspection	Trucks that have arrived at their assigned dock and are awaiting the pre-loading inspection step.
PalletsReadyForLoading	Pallets that have been released to the conveyor and are ready for pickup by a loading resource.
AvailableForklifts	Forklifts not currently engaged in a loading task in the AS-IS configuration.
AvailableAMRs	AMRs not currently engaged in a loading task or mandatory charging event in the TO-BE configuration.

The main model parameters used to control scenario configuration are summarised in Table 4.3.

Table 4.3: Main model parameters used to control simulation scenarios.

Parameter	Role
ScenarioType	Selects the active configuration: AS_IS or TO_BE.
ActiveDockCount	Defines the number of docks available for assignment in a given scenario.
FleetSizeForklifts	Defines the number of active forklifts in the AS-IS configuration.
FleetSizeAMRs	Defines the number of active AMRs in the TO-BE configuration.
WorkloadTarget	Defines the number of load orders/trucks to be generated in the scenario.
ArrivalPattern	Defines the truck arrival behaviour used in the scenario.
EnableAuxAMR	Activates the auxiliary AMR policy in TO-BE scenarios.

After the Inspection & Loading Service block, the token reaches the ScenarioType decision point, which routes AS_IS runs to the forklift-based loading logic and TO_BE runs to the AMR-based logic, including battery management and, when enabled, the auxiliary AMR policy. In the FlexSim implementation, available docks are stored in AvailableGates; in this dissertation, dock and gate refer to the same loading-position concept. Figure B.1 in Appendix B.1 shows this convergence and decision block, while the configuration-specific branches are described in Sections 4.2.3 and 4.2.4.

4.2.3 AS-IS Process Logic

In the AS-IS configuration, loading operations are executed by human-operated forklifts. After the common Process Flow logic has assigned a dock and coordinated the truck and pallet-release processes, the token is routed to the AS-IS branch through the ScenarioType decision block.

This branch covers three main functions: pre-loading inspection, forklift-based loading cycles and idle positioning.

The pre-loading inspection is modelled through the 'AS_IS Inspection' subflow and does not require the acquisition of a forklift resource, since it is assumed to be carried out by a supervisor or operator. Once the inspection is complete, the token proceeds to the 'AS_IS Loading' subflow, where a forklift is acquired from the `AvailableForklifts` list and assigned to the corresponding load order.

Each loading cycle follows the same structure: the forklift travels to the conveyor pickup area, retrieves the next pallet, transports it to the truck and places it according to the loading plan. When rotation is required, the forklift first unloads the pallet onto the floor, repositions itself, reloads the pallet from the required side and then completes the final placement. After each placement, the forklift returns to the pickup area and the next cycle is initiated until all pallets associated with the load order have been placed.

When not assigned to an active loading task, the forklift is directed to a designated park point through the 'Initialize ParkPoint' and 'Idle Forklift To Park Point' subflows, preserving spatial consistency and avoiding unnecessary occupation of operational areas. The AS-IS configuration does not include battery management or charging logic, since forklift batteries are assumed to last for a full working shift and to be replaced at the end of the shift.

Forklift behaviour is therefore tracked through a simpler set of operational states: empty travel, loaded travel, pallet loading, pallet unloading and idle time. These states provide the reference baseline for evaluating resource utilisation in the comparative analysis. A more detailed view of the AS-IS Process Flow implementation is shown in Figure B.3 in Appendix B.1.

4.2.4 TO-BE Process Logic

In the TO-BE configuration, forklifts are replaced by AMRs as the loading resource. The high-level loading cycle follows the same structure as in the AS-IS configuration, but introduces AMR-specific movement and handling characteristics, as well as battery management logic that affects resource availability throughout the simulation.

After entering the TO-BE branch, the token acquires an AMR from the `AvailableAMRs` list and assigns it to the corresponding load order. The assigned AMR first travels to the dock and performs the inspection before proceeding to the 'TO_BE Loading subflow', where it executes the loading cycles associated with the load order.

During each cycle, the AMR retrieves the next pallet from the conveyor, transports it to the truck, places it according to the loading plan and returns to the pickup area. When rotation is required, the AMR unloads the pallet, repositions itself, reloads it from the required side and then completes the final placement. Loading is complete when all pallets associated with the load order have been placed.

4.2.4.1 Battery and charging logic

Unlike forklifts, AMRs maintain a battery state that changes according to their operational activity. Different depletion rates are associated with states such as empty travel, loaded travel, pallet loading, pallet unloading, inspection and idle time, while the numerical rates, charging rates and thresholds are presented later in Section 4.3.

The model distinguishes between opportunistic and mandatory charging. Opportunistic charging is initiated when an AMR becomes idle, has sufficient battery to remain operational, and a park/charging point is available; in this case, idle time is used for recharging without removing the AMR from the available fleet. Mandatory charging is triggered after task completion whenever the battery level falls below the lower operational threshold, and the AMR only returns to availability after reaching the upper threshold.

The battery state is stored as a label on each AMR object and updated at relevant events using the 'Update Battery to Current Time' function. This function recalculates the battery level from the elapsed time since the previous update and the AMR state during that interval, avoiding the need for continuous polling.

The resulting AMR states include empty travel, loaded travel, pallet loading, pallet unloading, inspection, idle time, opportunistic charging and mandatory charging. These states support AMR-specific indicators such as charging time share, idle time, active utilisation and battery level evolution.

Figure 4.3 presents the AMR state diagram. Core loading-cycle states (idle/available, empty travel, pallet loading, loaded travel and pallet unloading) correspond to logic shared by both configurations, whereas inspection and charging-related states are TO-BE-specific.

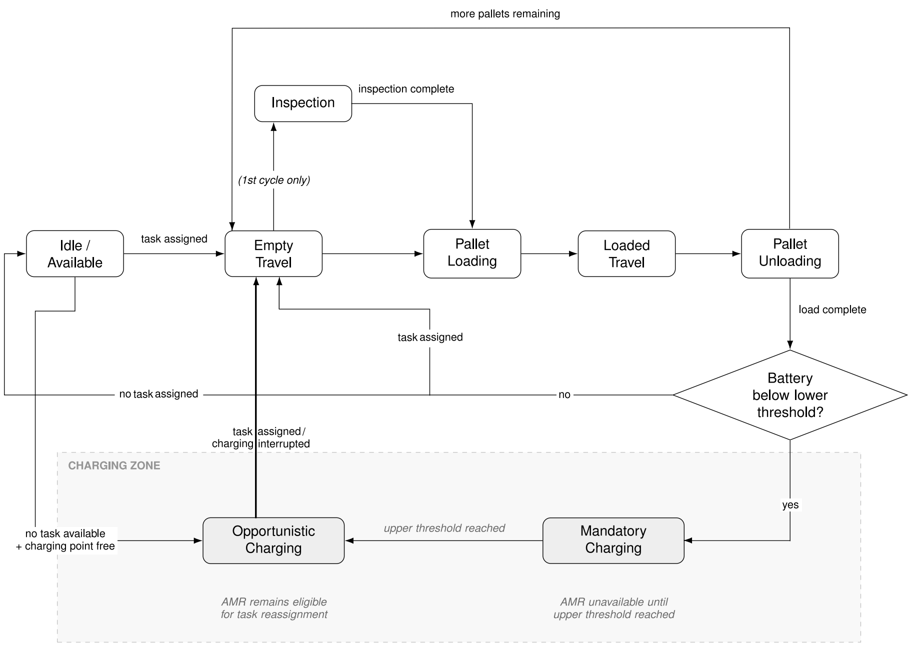

Figure 4.3: AMR operational state transitions in the TO-BE configuration.

4.2.4.2 Auxiliary AMR policy

The model also includes an optional auxiliary AMR policy, activated by the parameter `EnableAuxAMR`. This policy was introduced as an experimental task-assignment rule to evaluate whether assigning spare AMR capacity to active loads could reduce loading times by allowing two AMRs to cooperate on the same load.

When active, a second AMR may be assigned at the start of the loading phase, provided that no other load order is waiting for a primary AMR assignment. This ensures that starting new loads remains the first use of available AMRs, while auxiliary assignment is only allowed when spare fleet capacity exists. At most one auxiliary AMR can be assigned to each load, so a load can be served by a maximum of two AMRs. Once assigned, the auxiliary AMR remains committed to that load until completion, even if another truck later becomes available and requires a primary AMR assignment.

When two AMRs are assigned to the same load, the `PalletClaimLock` mechanism coordinates pallet retrieval. Before retrieving the next pallet from the conveyor, each AMR must acquire this lock, preventing both resources from claiming the same pallet simultaneously. Once a pallet is claimed and the AMR begins its transport cycle, the lock is released, allowing the other AMR to claim the following pallet in sequence. This ensures that the loading plan sequence is respected even when two resources operate concurrently on the same load.

The primary AMR assigned to a load is recorded in the label `LoaderName`; the auxiliary AMR, when present, is recorded in `AuxLoaderName`. The label `LoadShared` indicates whether a load was completed with auxiliary AMR participation. These labels support the post-simulation analysis of resource utilisation and load participation. A more detailed view of the TO-BE Process Flow implementation is shown in Figure B.4 in Appendix B.1.

4.3 Operational Parameters and Resource Specifications

This section presents the operational parameters used to instantiate the simulation model described in Section 4.2. Parameters are organised by scope: common process parameters that apply to both configurations, followed by resource-specific specifications for forklifts and AMRs, and finally the battery and charging parameters governing AMR energy management. Fleet sizes and the number of active docks are not fixed here; they constitute experimental factors defined in Section 4.8.

4.3.1 Common Operational Parameters

Table 4.4 presents the process-level parameters shared by both the AS-IS and TO-BE configurations. Travel time to the assigned dock was parameterised through a triangular distribution using the lower, typical, and upper values identified during the operational characterisation of the case study. Pallet availability parameters represent the progressive upstream feeding of pallets to the conveyor interface, without modelling the upstream warehouse as a limiting resource.

Table 4.4: Common operational parameters.

Parameter	Value	Unit	Description
Travel time to assigned dock	Triangular(8, 10, 12)	min	Based on lower, typical, and upper operational values
Pre-loading inspection time	60	s	Fixed duration applied to the inspection step
First pallet available after dock assignment	5	min	Represents initiation of upstream pallet feeding
All pallets available after dock assignment	10	min	Represents progressive pallet availability at the conveyor interface

4.3.2 Forklift Resource Specifications

The forklift parameters were defined based on technical specifications of the vehicle model used at the case study facility, with values adjusted to reflect operational conditions rather than peak technical limits. Table 4.5 summarises the main resource parameters.

Table 4.5: Forklift resource specifications (AS-IS configuration).

Parameter	Value	Unit	Description
`Forklift_Speed_Empty`	2.5	m/s	Travel max speed when not carrying a pallet
`Forklift_Speed_Loaded`	2.0	m/s	Travel max speed when carrying a pallet
`Forklift_Max_Accel`	1.2	m/s^2	Maximum acceleration
`Forklift_Max_Decel`	1.2	m/s^2	Maximum deceleration
`ForkliftLoadTime`	5.0	s	Time to pick up one pallet
`ForkliftUnloadTime`	5.0	s	Time required to set down or place one pallet

When a loading plan requires pallet rotation, the forklift executes an additional unload-reposition-load sequence before proceeding to the truck. The total time added by this rotation step is not a separate parameter but emerges from the combination of `ForkliftUnloadTime`, the repositioning travel, and `ForkliftLoadTime`.

4.3.3 AMR Resource Specifications

AMR movement and handling parameters are defined separately from the forklift values, reflecting the different operational behaviour of the TO-BE loading resource. As with the forklift, pallet rotation adds an unload-reposition-load sequence to the cycle, with the duration determined by `AMRUnloadTime`, repositioning travel, and `AMRLoadTime`. Table 4.6 summarises the main AMR resource specifications.

Table 4.6: AMR resource specifications (TO-BE configuration).

Parameter	**Value**	**Unit**	**Description**
`AMR_SpeedEmpty`	1.5	m/s	Travel speed when not carrying a pallet
`AMR_SpeedLoaded`	1.0	m/s	Travel speed when carrying a pallet
`AMR_Max_Accel`	0.3	m/s^2	Maximum acceleration
`AMR_Max_Decel`	0.3	m/s^2	Maximum deceleration
`AMRLoadTime`	15.0	s	Time to pick up one pallet
`AMRUnloadTime`	10.0	s	Time required to set down or place one pallet

4.3.4 Battery and Charging Parameters

Battery depletion is modelled at the state level: each AMR operational state is associated with a distinct depletion rate, expressed in percentage points per minute. The original AMR specifications provided battery capacity in ampere-hours (Ah) and state-dependent current values in amperes (A). These values were converted into percentage variation per minute using Equation 4.1, where I represents the current draw or charging current and C_{bat} represents the nominal battery capacity.

$$r_{\%/\text{min}} = \frac{I}{C_{bat}} \times \frac{100}{60} \tag{4.1}$$

For example, the loaded travel depletion rate was obtained using a current draw of 15 A and a battery capacity of 150 Ah, resulting in:

$$r_{\text{loaded}} = \frac{15}{150} \times \frac{100}{60} \approx 0.17\%/\text{min} \tag{4.2}$$

The same conversion logic was applied to the remaining depletion and charging rates. Table 4.7 summarises the main battery and charging parameters; the complete set of state-specific depletion rates is reported in Table B.1 in Appendix B.2.

Table 4.7: Main AMR battery and charging parameters (TO-BE configuration).

Parameter	**Value**	**Unit**	**Description**
`InitialBatteryLevel`	100	%	Battery level at simulation start
`AMR_LowerBatteryThreshold`	20	%	Triggers mandatory charging below this level
`AMR_UpperBatteryThreshold`	80	%	Target level for mandatory charging
`AMR_ChargeRate_0_80`	0.73	%/min	Charging rate up to the upper threshold
`AMR_ChargeRate_80_100`	0.29	%/min	Charging rate above the upper threshold

This representation allows the model to update the AMR battery level consistently across operational states without requiring a detailed electrical model. Charging follows two rate stages:

a faster rate below 80% and a slower rate above it, consistent with the charging profile assumed for the AMR battery model. The upper threshold of 80% serves as both the target for mandatory charging and the transition point between the two charging rates. The complete parameter structure implemented in FlexSim is shown in Figures B.5, B.6 and B.7 in Appendix B.2.

4.4 Input Data Pipeline and Workload Construction

The eight load typologies introduced in Section 3.6 serve as the basis for generating the workload instances used in the simulation. This section describes how those typologies were converted into structured input data for the FlexSim model and how workloads of different sizes were constructed. In relation to the input layer presented in Figure 4.2, the focus of this section is the construction of the `LoadNumData` table, the main operational input consumed by the Process Flow.

The representative load order data for the eight typologies were initially defined in JSON format, capturing the main attributes required to characterise each typology, such as vehicle type, pallet type, number of load lines and total pallet count. A Python script converted this JSON representation into a structured Excel table containing the eight representative typologies, from which a workload-generation script created experimental workload instances by replicating the typologies according to the target workload size. The resulting Excel file was imported into FlexSim and stored as a global table named `LoadNumData`. Each LoadNum identifies one truck loading operation, while the corresponding rows provide the load-line information required to create the truck, generate the pallets, assign the relevant dock and conveyor inputs, and link the load to the appropriate loading plan. A representative excerpt of the `LoadNumData` table is provided in Table B.4 in Appendix B.4.

In addition to the `LoadNumData` table, predefined loading-plan information was stored in FlexSim tables. Since the study is based on eight representative load typologies, the corresponding loading plans were computed in advance through the load scheduling algorithm and introduced into FlexSim as input tables. These tables specify the pallet placement sequence, the position inside the truck, and whether pallet rotation is required. During the simulation, the Process Flow uses `LoadNumData` to identify the load being processed and the loading-plan table to execute the corresponding pallet placement sequence. The complete loading-plan table for Typology 2, corresponding to layout L24, is provided in Table B.5 in Appendix B.4, and the overall input data preparation pipeline is summarised in Figure 4.4.

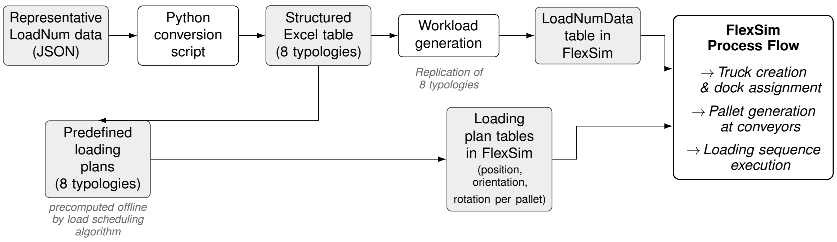

Figure 4.4: Input data preparation pipeline for the FlexSim simulation model.

Workloads of three sizes are used in this study: 100, 120 and 140 trucks. These values were selected as plausible representations of different daily operational demand levels. In the baseline workload construction, instances are distributed approximately uniformly across the eight typologies. For workloads that are exact multiples of eight, such as 120 trucks, the distribution is perfectly balanced; for 100 and 140 trucks, the remaining instances are distributed as evenly as possible, with a difference of at most one instance between the most and least represented typology. This ensures that no typology dominates and that comparisons across workload levels remain interpretable. An additional workload of 200 trucks, constructed under the same balancing principle, is used in the high-demand sensitivity scenario introduced in Section 4.8.

The workload table defines the set of load orders to be processed in a given simulation run. The temporal release of trucks into the system, that is, the interarrival times governing when each truck arrives at the facility entrance, is defined separately and is described in Section 4.5. The `LoadNumData` table is therefore concerned with what is processed, not when.

4.5 Arrival Process Modelling

Truck arrivals were modelled as independent events, with the time between consecutive arrivals represented by exponential interarrival times. Two arrival patterns are considered: a stationary Poisson process, used as the controlled baseline, and a piecewise non-stationary Poisson process, used to represent a structured temporal concentration of arrivals. In both cases, the total number of trucks generated in each run is controlled by the `WorkloadTarget` parameter.

The arrival rate and mean interarrival time are derived from the workload target W and the arrival window H, according to Equation 4.3:

$$\lambda = \frac{W}{H}, \qquad \bar{t} = \frac{60}{\lambda} = \frac{60H}{W} \tag{4.3}$$

where λ is expressed in trucks per hour, $\bar{t}$ is expressed in minutes, H is expressed in hours, and W represents the number of trucks in the workload. The 16-hour arrival window is used as a parametrisation reference for the expected arrival cadence; it does not impose a hard arrival deadline. Since interarrival times are sampled stochastically, the last truck in a given replication may arrive before or after the 16-hour mark. Table 4.8 presents the resulting stationary arrival parameters for the three workload levels.

Table 4.8: Stationary Poisson arrival parameters by workload level.

Workload	**λ (trucks/h)**	**Mean interarrival (min)**
100 trucks	6.25	9.60
120 trucks	7.50	8.00
140 trucks	8.75	6.86

Under the non-stationary pattern, the expected arrival intensity varies across the operational day. The 16-hour window is divided into eight two-hour blocks, each assigned a workload share p_b. Within each block, interarrival times continue to be sampled from an exponential distribution, but the mean changes according to the active block. The block-specific arrival rate and mean interarrival time are computed using Equation 4.4:

$$\lambda_b = \frac{W p_b}{2}, \qquad \bar{t}_b = \frac{60}{\lambda_b} = \frac{120}{W p_b} \tag{4.4}$$

where λ_b is expressed in trucks per hour, $\bar{t}_b$ is expressed in minutes, W is the workload target, and p_b is the workload share assigned to block b. This formulation changes the expected arrival intensity over time while preserving the stochastic nature of individual arrivals. As a result, individual replications may deviate from the expected block shares, while the aggregate pattern is expected to reflect the defined temporal profile.

The non-stationary profile follows a bimodal structure, with higher expected arrival intensity in the h4-h6 and h12-h14 blocks. Table 4.9 summarises the block shares and corresponding parameters for the 120-truck workload as a representative case.

Table 4.9: Non-stationary Poisson arrival profile for the 120-truck workload.

Block	**Workload share**	**E[trucks]**	**λ (trucks/h)**	**Mean interarrival (min)**
h0-h2	8%	9.6	4.8	12.5
h2-h4	15%	18.0	9.0	6.7
h4-h6	19%	22.8	11.4	5.3
h6-h8	10%	12.0	6.0	10.0
h8-h10	8%	9.6	4.8	12.5
h10-h12	13%	15.6	7.8	7.7
h12-h14	17%	20.4	10.2	5.9
h14-h16	10%	12.0	6.0	10.0
Total	**100%**	**120.0**		

The block shares were defined as scenario-design assumptions to represent a structured and operationally plausible temporal concentration of arrivals, rather than to reproduce a specific observed day. The same proportional profile can be applied to different workload levels by recalculating the expected number of trucks and the corresponding mean interarrival time for each block. The non-stationary arrival profile is shown in Figure B.8 and the corresponding arrival-parameter tables for the remaining workload levels are presented in Tables B.2 and B.3 in Appendix B.3.

In FlexSim, the arrival process is implemented through the Source activity, with the ArrivalPattern parameter selecting the stationary or non-stationary distribution. Under the non-stationary setting, the mean interarrival time is selected according to the active two-hour block when each arrival is generated, while `WorkloadTarget` controls the total number of trucks.

4.6 Output Measurement and Data Collection

Alongside the process logic, parameters, and input data described in the preceding sections, the model includes a structured output measurement layer designed to support both visual inspection during development and systematic data collection for experimental comparison. Two complementary mechanisms are used: dashboards for run-level monitoring and diagnostic support, and Performance Measures for structured output collection across scenarios and replications.

Three groups of dashboards were developed within the model. The loading dashboard summarises truck, dock, loading, arrival, rotation and weight-related indicators at system level. The AMR and forklift dashboards present the corresponding resource-level information for each configuration: operational states, utilisation, and, for AMRs, battery level and charging behaviour. These dashboards supported model development, debugging, verification and exploratory inspection of model behaviour during individual runs. They also provide a visual interface through which managers can understand the model, build trust in its behaviour and identify operational insights. Representative layouts of the loading, AMR and forklift dashboards are shown in Figures B.9, B.10 and B.11 in Appendix B.5.

Performance Measures constitute the formal output collection mechanism used for scenario comparison. Unlike dashboard indicators, which are observed interactively during a run, Performance Measures are recorded automatically at the end of each replication and can be aggregated consistently across runs in the FlexSim Experimenter. The implemented measures are organised into three groups: System_Performance, Resource_Performance, and Energy_Performance. These groups cover system-level throughput, truck flow and loading efficiency, resource activity, and, for the TO-BE configuration, battery and charging behaviour. Table 4.10 summarises the main measures in each group and their analytical purpose. For scenario selection, TotalCompletionTime is the primary system-level measure, since all scenarios process the same workload.

Table 4.10: Main Performance Measures by output group.

Output group	Main measures	Analytical purpose
`System_Performance`	TotalCompletionTime; average truck total time in system; average waiting time for dock assignment; average loading time	Assess the primary scenario-comparison outcome and truck-level flow performance across configurations and workloads.
`Resource_Performance`	Average loading cycle time, including overall, non-rotated and rotated pallets;	Compare loading-cycle efficiency and the overall movement effort required by forklifts and AMRs.
`Energy_Performance`	Average AMR battery level; average charging state percentage; mandatory charging event count	Evaluate the impact of battery depletion and charging behaviour on AMR availability.

Among the resource-level measures, the share of time each AMR spends in its operational states is also recorded. These states are grouped into working states, comprising travel, loading, unloading and inspection, and non-working states, comprising idle and charging. As a practical guideline, a combined idle-and-charging share below approximately 20 per cent, equivalently a working-state share above 80 per cent, is commonly used as an indication that a fleet has limited spare capacity and that an additional unit may be beneficial.

The main quantitative comparisons reported in this dissertation are derived from Performance Measure outputs aggregated across replications. The scenario execution and replication structure are described in Section 4.8. The implemented system-level, resource-level and battery-related Performance Measures are listed in Tables B.6, B.7 and B.8 in Appendix B.6.

4.7 Model Verification and Operational Plausibility

Consistent with the model's purpose, verification focused on internal consistency, correct implementation of the process logic, and operational plausibility, rather than on formal statistical validation against historical observations.

Process logic verification was carried out through incremental Process Flow testing, using small-workload runs to inspect token progression step by step. These tests confirmed that the main operational sequence was executed correctly, from truck generation and dock assignment to pallet generation, inspection, loading completion, truck departure and dock release. They also verified that load completion was triggered only after all pallets associated with each LoadNum had been placed, and that both `ScenarioType` branches behaved as expected.

Resource behaviour was verified by observing state transitions for forklifts and AMRs across representative runs. In the AS-IS configuration, forklifts were confirmed to execute loading cycles and returned to park points when no task was available. In the TO-BE configuration, AMRs transitioned correctly between travel, loading, idle and charging states. Additional checks confirmed that mandatory charging removed AMRs from `AvailableAMRs` until the upper threshold was reached, that the auxiliary AMR policy was activated only under the intended conditions, and that `PalletClaimLock` prevented simultaneous pallet claims. Figure 4.5 illustrates this verification process through the AMR battery-level evolution over a representative run.

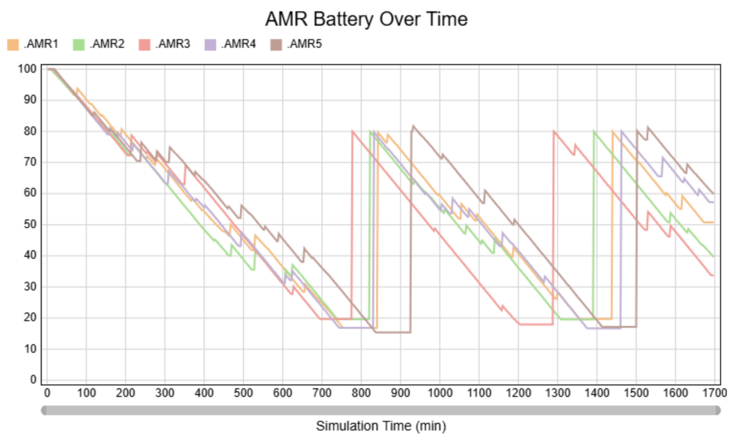

Figure 4.5: Battery level evolution of the AMR fleet over a representative run.

The figure confirms that the charging logic behaves as intended: battery levels decrease during operation, AMRs enter mandatory charging after reaching the lower threshold, and charging continues until the upper threshold is reached. It also shows that an AMR may complete an ongoing task slightly below the lower threshold before charging, which is consistent with the implemented rule that the threshold prevents new task assignment rather than interrupting work already in progress.

Input-output consistency was monitored using counters, dashboards and statistics collectors. The main checks confirmed that generated trucks matched `WorkloadTarget`, that `CompletedLoads` reached `ExpectedLoads` without residual blocked trucks, that TotalCompletionTime was recorded only after all load orders were completed, and that pallet counts per load were consistent with the `LoadNumData` table across the representative typologies.

Operational plausibility was assessed by comparing the model behaviour with the case study characterisation presented in Chapter 3. Workload sizes and arrival parameters were defined as controlled scenario inputs grounded in available operational information, while resource specifications were based on technical references and scenario-design assumptions, as described in Section 4.3. The validity claim is therefore based on an internally consistent and symmetrically parametrised comparison between AS-IS and TO-BE configurations under equivalent demand and operational conditions, rather than on the exact prediction of a specific historical day. This approach is consistent with the view that the appropriate scope and rigour of model verification depend on the model's intended purpose (Sargent, 2010).

4.8 Experimental Design and Replication Strategy

The experimental strategy was structured to compare the AS-IS and TO-BE configurations under controlled and progressively more demanding conditions. Across the primary experiments, the reference AS-IS configuration uses four forklifts and the base dock configuration assumes six active docks. The general principle governing the design is that comparable scenarios maintain equivalent conditions whenever possible, changing only the factor or set of factors under evaluation. This supports a controlled comparison between the existing forklift-based system and the proposed AMR-based configuration.

The main experimental factors considered are `ScenarioType` (AS-IS or TO-BE), loading-resource fleet size, `EnableAuxAMR`, workload level, arrival pattern, and load composition. These factors are not varied through a full factorial design. Instead, the experiments are organised into sequential blocks, each targeting a specific dimension of analysis while maintaining controlled comparison conditions. This structure keeps the experimental space manageable by narrowing the configurations evaluated in later blocks, rather than testing all possible factor combinations.

The experimental design is organised into five experimental blocks and one additional stress-test scenario. Block A establishes the controlled baseline comparison between AS-IS and TO-BE without the auxiliary AMR policy. Block B evaluates the auxiliary AMR policy under the same stationary arrival and balanced workload conditions. Block C tests robustness under the non-stationary Poisson arrival pattern described in Section 4.5. Block D evaluates sensitivity to a more demanding load composition, with a higher share of loads requiring pallet rotation, while Block E assesses whether the conclusions remain consistent under a substantially higher workload. The additional stress-test scenario is treated separately because it changes the truck-release logic rather than only varying operational demand or configuration parameters. Figure 4.6 summarises the workflow from scenario definition to output analysis.

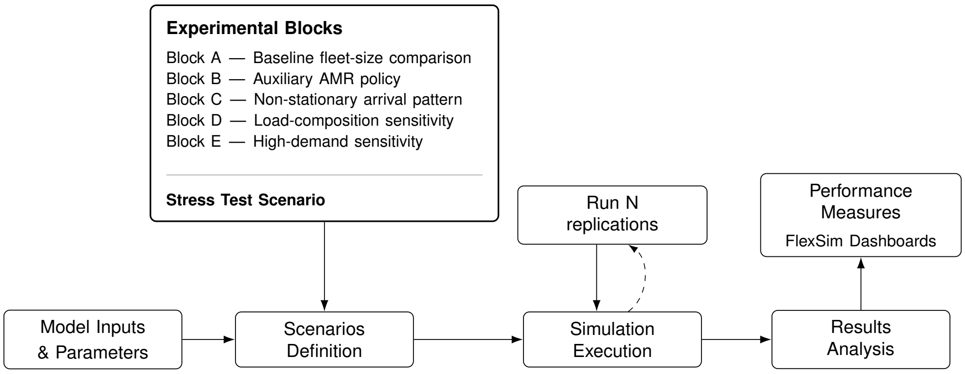

Figure 4.6: Experimental workflow from scenario definition to results analysis.

In Block A, the TO-BE configuration is evaluated without the auxiliary AMR policy, using a stationary Poisson arrival process and a balanced distribution of the eight representative load typologies. The AS-IS reference uses four forklifts, while the TO-BE fleet sizes range from four to eight AMRs. This block establishes the baseline fleet-size comparison when each load is served by a single AMR.

In Block B, the same stationary arrival pattern and balanced load composition are maintained, but the auxiliary AMR policy is enabled. TO-BE fleet sizes from five to nine AMRs are tested, alongside the same four-forklift AS-IS reference. This block evaluates whether allowing spare AMR capacity to assist an ongoing load changes the behaviour of the TO-BE configuration under otherwise equivalent conditions.

In Block C, the auxiliary AMR policy remains enabled and the arrival process is changed to the non-stationary Poisson pattern. The load composition remains balanced across the eight typologies. The bimodal arrival profile concentrates demand in two higher-intensity blocks separated by lower-intensity periods, allowing the model to test whether the scenario comparison remains robust under a more temporally concentrated arrival pattern. Fleet sizes of six to nine AMRs are evaluated, selected based on the outcomes of Blocks A and B.

In Block D, the arrival process returns to the stationary Poisson pattern, but the workload composition is changed. Instead of a balanced typology distribution, a rotation-heavy workload is used, in which approximately 80 per cent of loads require pallet rotation and the remaining 20 per cent correspond to the other typologies. This block is applied to workloads of 120 and 140 trucks, with TO-BE fleet sizes from seven to eleven AMRs and the auxiliary AMR policy enabled. The purpose is to evaluate the sensitivity of the AS-IS versus TO-BE comparison to a more demanding load mix.

In Block E, a high-demand sensitivity scenario is used to test whether the conclusions remain consistent when demand increases substantially. This block considers a workload of 200 trucks within the same 16-hour base arrival window, under stationary Poisson arrivals, balanced load composition and the auxiliary AMR policy enabled. The AS-IS reference uses four forklifts, while the TO-BE configuration is evaluated across fleet sizes from six to twelve AMRs. This block reflects a situation in which the outbound operation faces substantially higher daily demand, due to seasonal peaks, business growth or the consolidation of additional product flows.

An additional stress-test scenario is considered outside the five experimental blocks. Unlike the block-based scenarios, all LoadNum instances are made available at simulation time zero rather than released progressively according to the arrival process. By removing arrival cadence as a limiting factor, this scenario provides an upper-bound reference for system behaviour under continuous demand, rather than a normal operating pattern. It uses a balanced workload of 100 trucks, the four-forklift AS-IS reference, and TO-BE fleet sizes from eight to twelve AMRs with the auxiliary AMR policy enabled. Table 4.11 summarises the experimental blocks and the stress-test scenario.

Table 4.11: Summary of the five experimental blocks and the stress-test scenario.

Block	**Arrival pattern**	**Load composition**	**AuxAMR**	**Resource configurations**	**Workloads**
A	Stationary Poisson	Balanced typology mix	Disabled	AS-IS: 4 forklifts; TO-BE: 4-8 AMRs	100, 120, 140
B	Stationary Poisson	Balanced typology mix	Enabled	AS-IS: 4 forklifts; TO-BE: 5-9 AMRs	100, 120, 140
C	Non-stationary Poisson	Balanced typology mix	Enabled	AS-IS: 4 forklifts; TO-BE: 6-9 AMRs	100, 120, 140
D	Stationary Poisson	Rotation-heavy mix	Enabled	AS-IS: 4 forklifts; TO-BE: 7-11 AMRs	120, 140
E	Stationary Poisson	Balanced typology mix	Enabled	AS-IS: 4 forklifts; TO-BE: 6-12 AMRs	200
Stress test	All loads at $t = 0$	Balanced typology mix	Enabled	AS-IS: 4 forklifts; TO-BE: 8-12 AMRs	100

Each scenario was executed with 30 independent replications. The same replication count was used across all scenarios to avoid differences arising from unequal sampling effort. Its adequacy was assessed a posteriori through the precision of the obtained confidence intervals: for the main system-level measures, the 95% confidence-interval half-widths remained within a small fraction of the corresponding means, indicating sufficient precision to distinguish between the compared configurations.

Within each experiment, a common random-number structure was used so that scenarios sharing the same replication index were exposed to equivalent stochastic arrival conditions. Performance Measures were collected per replication through the FlexSim Experimenter and aggregated using means, confidence intervals and dispersion indicators.

Each replication ends when `CompletedLoads` reaches `ExpectedLoads`, at which point TotalCompletionTime is recorded. A safety stop time is defined only as a technical safeguard; scenarios are therefore compared by completion of the same workload rather than by a fixed time horizon.

Chapter 5

Results and Discussion

This chapter discusses the results obtained from the experimental design defined in Section 4.8. The analysis follows the five experimental blocks: the baseline fleet-size comparison (Block A), the auxiliary AMR policy evaluation (Block B), the arrival-pattern sensitivity analysis (Block C), the load-composition sensitivity analysis (Block D), and the high-demand scenario (Block E), followed by the stress-test scenario and an integrated discussion of the operational implications. The discussion focuses on the main system-level Performance Measures, with results reported across 30 replications and 95% confidence intervals. Detailed Performance Measure outputs are provided in Appendix C.

5.1 Block A: Baseline Fleet-size Comparison

Block A establishes the baseline comparison between the AS-IS and TO-BE configurations under stationary Poisson arrivals, a balanced load composition, six active docks, and the auxiliary AMR policy disabled. The AS-IS configuration uses four forklifts, while the TO-BE configuration is evaluated across fleet sizes of four to eight AMRs. The objective is to assess whether increasing the AMR fleet size alone is sufficient to approach the AS-IS baseline. Table 5.1 reports the main system-level Performance Measures across the three workload levels.

Table 5.1: Block A - Main system-level Performance Measures, reported as mean $\pm$ 95% confidence interval across 30 replications.

Measure	Workload	AS-IS	4 AMRs	5 AMRs	6 AMRs	7 AMRs	8 AMRs
Completion time (h)	100	16.35 ± 0.7	17.82 ± 0.3	16.77 ± 0.6	16.71 ± 0.7	16.72 ± 0.7	16.72 ± 0.7
	120	16.41 ± 0.6	21.04 ± 0.2	17.70 ± 0.3	17.34 ± 0.4	17.34 ± 0.4	17.34 ± 0.4
	140	16.47 ± 0.5	24.41 ± 0.1	20.05 ± 0.1	18.49 ± 0.2	18.49 ± 0.2	18.49 ± 0.2
Truck time in system (min)	100	31.8 ± 4.7	81.9 ± 13.2	55.3 ± 4.2	52.9 ± 3.3	52.9 ± 3.3	52.9 ± 3.3
	120	33.1 ± 8.0	149.6 ± 16.9	76.7 ± 9.9	68.6 ± 8.0	68.6 ± 8.0	68.6 ± 8.0
	140	37.0 ± 2.2	271.4 ± 17.4	144.1 ± 16.4	115.8 ± 14.8	115.7 ± 14.8	115.8 ± 14.8
Waiting for dock (min)	100	0.9 ± 0.4	28.4 ± 12.0	9.2 ± 4.1	7.7 ± 3.3	7.7 ± 3.3	7.7 ± 3.3
	120	2.0 ± 0.7	91.4 ± 16.5	30.0 ± 9.6	23.5 ± 8.0	23.5 ± 8.0	23.5 ± 8.0
	140	5.5 ± 2.1	210.7 ± 17.4	94.8 ± 16.3	70.6 ± 14.8	70.6 ± 14.8	70.6 ± 14.8

A first observation is that TO-BE performance reaches a saturation point from six AMRs onwards. Across all three workload levels, total completion time is effectively identical for fleets of six, seven and eight AMRs, indicating that additional AMRs no longer improve system throughput. This pattern results from the interaction between the six active docks and the single-AMR-per-load rule: once enough AMRs are available to serve all active docks, further fleet expansion does not create additional simultaneous loading capacity under the baseline TO-BE logic. The seven- and eight-AMR runs therefore act as a robustness check for the fleet-size range, confirming that the absence of improvement beyond six AMRs is structural rather than an artefact of stopping the test too early. This pattern is also visible in the replication-level TotalCompletionTime plots for the three workload levels, shown in Figures C.1, C.2 and C.3 in Appendix C.1.

A second observation concerns the sensitivity of each configuration to workload. The AS-IS configuration is remarkably stable across workload levels, with completion times ranging only from 16.35 to 16.47 hours between the 100- and 140-truck workloads. The TO-BE configuration, by contrast, degrades markedly as workload increases. At the 100-truck workload, the TO-BE configuration with six or more AMRs approaches the AS-IS baseline closely. For the 140-truck workload, the best TO-BE baseline configuration remains above the AS-IS reference, with total completion times of 18.49 h and 16.47 h, respectively, a statistically significant difference. This confirms that the relative advantage of the AS-IS configuration widens with increasing demand.

This difference is explained primarily by the loading time at truck level, which is insensitive to fleet size but reflects the operational characteristics of each resource type. The average loading time per truck is approximately 33.7 minutes in the TO-BE configuration and approximately 19.2 minutes in the AS-IS configuration, regardless of fleet size. Since each load is served by a single resource when the auxiliary policy is disabled, the slower handling and travel characteristics of the AMR translate directly into longer per-truck loading times, which accumulate into longer system completion times as workload increases. The complete loading-time outputs for Block A are provided in Table C.1 in Appendix C.1.

Battery-related constraints were relevant only in the smaller AMR fleets. Mandatory charging occurred in the four- and five-AMR configurations and increased with workload, but disappeared from six AMRs onwards. Since the competitive TO-BE configurations therefore operated without mandatory charging constraints, battery behaviour does not explain the remaining completion-time gap relative to the AS-IS benchmark. The full battery and charging outputs are reported in Table C.2 in Appendix C.1.

In summary, Block A shows that increasing the AMR fleet improves TO-BE performance only up to the point where all active docks can be served. The TO-BE configuration requires at least six AMRs to reach its baseline performance ceiling, but this ceiling remains above the AS-IS baseline at higher workloads. The gap is driven mainly by per-truck loading time rather than by fleet availability or battery constraints. These findings motivate the evaluation of the auxiliary AMR policy in Block B, which targets the per-truck loading process by allowing additional AMR capacity to be used within the same load.

5.2 Block B: Auxiliary AMR Policy Evaluation

Block B evaluates the effect of the auxiliary AMR policy on TO-BE performance under the same stationary Poisson arrivals, balanced load composition and six active docks used in Block A. When enabled, the policy allows a second AMR to assist an ongoing load under the conditions defined in Section 4.2.4. Since the AS-IS configuration is unaffected by this policy, the equivalent Block A results are used here as the reference benchmark. The main analysis focuses on fleet sizes from five to nine AMRs. A preliminary four-AMR auxiliary configuration was tested for the 100-truck workload as a screening check, but it was not extended to the remaining workloads because its completion time remained clearly above the AS-IS benchmark. It was therefore outside the competitive range, defined here as the set of fleet sizes whose performance is close enough to the benchmark to be relevant for the substitution decision. Table 5.2 reports the selected system-level Performance Measures.

Table 5.2: Block B - Selected Performance Measures, reported as mean ± 95% confidence interval across 30 replications.

Measure	W	AS-IS*	5 AMR	6 AMR	7 AMR	8 AMR	9 AMR
Completion time (h)	100	16.35	16.63 ± 0.6	16.45 ± 0.7	16.40 ± 0.7	16.36 ± 0.7	16.34 ± 0.7
	120	16.41	17.83 ± 0.3	16.80 ± 0.5	16.61 ± 0.6	16.53 ± 0.6	16.50 ± 0.6
	140	16.47	19.86 ± 0.1	17.56 ± 0.2	16.83 ± 0.4	16.58 ± 0.5	16.50 ± 0.5
Waiting for dock (min)	100	0.9	6.3 ± 3.4	2.7 ± 1.4	1.8 ± 1.0	1.2 ± 0.6	1.0 ± 0.5
	120	2.0	27.1 ± 9.7	10.8 ± 4.8	5.7 ± 2.4	3.3 ± 1.3	2.6 ± 1.0
	140	5.5	96.5 ± 16.1	39.9 ± 11.5	22.1 ± 8.3	12.2 ± 4.7	7.7 ± 3.2
Loading time (min)	100	19.2	27.1 ± 0.6	23.5 ± 0.5	22.3 ± 0.3	20.6 ± 0.2	20.2 ± 0.1
	120	19.2	29.0 ± 0.4	25.3 ± 0.5	23.3 ± 0.3	21.3 ± 0.2	20.5 ± 0.2
	140	19.2	30.6 ± 0.2	27.0 ± 0.3	24.4 ± 0.3	22.3 ± 0.2	21.0 ± 0.1

The main result of Block B is that the auxiliary policy changes the fleet-size effect observed in Block A. With the auxiliary policy enabled, total completion time continues to decrease up to nine AMRs across all workload levels, instead of stabilising after six AMRs. This indicates that AMR capacity previously unused under the single-AMR-per-load rule is now used within ongoing loads to reduce truck-level loading time.

As a consequence, the TO-BE configuration becomes competitive with the AS-IS benchmark. At the 100-truck workload, six or more AMRs already approach the AS-IS reference; at 120 trucks, this requires eight to nine AMRs; and at 140 trucks, nine AMRs effectively close the gap, reaching 16.50 hours against 16.47 hours for AS-IS. The fleet size required to reach parity therefore increases with workload. The replication-level results for the three workload levels are provided in Figures C.4, C.5 and C.6 in Appendix C.2.

The operational mechanism behind this improvement is the reduction in truck-level loading time, which directly captures the effect of the auxiliary policy. In Block A, loading time remained close to 33.7 minutes regardless of AMR fleet size, since each load was served by a single AMR. With the auxiliary policy enabled, loading time decreases as fleet size grows; at the 140-truck workload, it falls from 30.6 minutes with five AMRs to 21.0 minutes with nine AMRs, a 38% reduction relative to the non-auxiliary baseline. Faster loading releases docks sooner, reducing dock waiting, truck time in system and total completion time. Detailed loading-time outputs and reductions are reported in Tables C.4 and C.5, and truck-flow output in Table C.3, all in Appendix C.2.

Shared-load participation confirms the effective use of the auxiliary policy. It increases with fleet size in all workloads, showing that spare AMR capacity is increasingly assigned to ongoing loads. At the 120-truck workload, it rises from 31.7% with five AMRs to 90.3% with nine AMRs. For a given fleet size, participation is lower at higher workloads because primary AMR assignments take priority over auxiliary support.

Battery-related constraints remain confined to smaller fleets. Mandatory charging occurs mainly in five-AMR configurations and, marginally, in the six-AMR configuration at the 140-truck workload, while competitive seven- to nine-AMR configurations record none. The auxiliary policy therefore does not add battery pressure in the competitive range. Battery and charging outputs are reported in Table C.6 in Appendix C.2.

In summary, Block B shows that the auxiliary AMR policy uses spare fleet capacity to reduce per-truck loading time, allowing sufficiently large TO-BE fleets to match the AS-IS benchmark under stationary arrivals. These findings motivate Block C, which tests whether this conclusion remains robust when arrivals are temporally concentrated.

5.3 Block C: Sensitivity to Non-Stationary Arrivals

The TO-BE configurations in Block C retain the auxiliary AMR policy and balanced load composition, while arrivals follow the non-stationary Poisson process defined in Section 4.5. To ensure a fair comparison, the AS-IS configuration with four forklifts was re-simulated under the same arrival pattern, and the TO-BE configuration is evaluated across six to nine AMRs, the competitive range identified in Block B. Table 5.3 reports the selected system-level Performance Measures, while truck time in system is reported separately in Table C.8 in Appendix C.3.

Table 5.3: Block C - Selected system-level Performance Measures under non-stationary arrivals, reported as mean ± 95% confidence interval across 30 replications.

Measure	W	AS-IS	6 AMR	7 AMR	8 AMR	9 AMR
Completion time (h)	100	16.61 ± 0.8	16.75 ± 0.7	16.70 ± 0.7	16.65 ± 0.8	16.60 ± 0.8
	120	16.70 ± 0.6	17.10 ± 0.5	16.90 ± 0.6	16.81 ± 0.6	16.77 ± 0.6
	140	16.66 ± 0.6	18.00 ± 0.3	17.17 ± 0.4	16.90 ± 0.5	16.77 ± 0.6
Waiting for dock (min)	100	2.2 ± 1.1	7.6 ± 2.9	5.0 ± 2.3	3.3 ± 1.6	2.6 ± 1.4
	120	5.5 ± 1.9	21.5 ± 6.2	13.9 ± 4.0	9.6 ± 2.8	7.3 ± 2.3
	140	13.6 ± 3.6	60.0 ± 11.9	38.5 ± 9.2	25.4 ± 6.4	18.6 ± 4.8
Loading time (min)	100	19.2 ± 0.0	24.5 ± 0.5	22.8 ± 0.3	20.9 ± 0.2	20.4 ± 0.1
	120	19.1 ± 0.0	25.6 ± 0.4	23.4 ± 0.3	21.6 ± 0.3	20.7 ± 0.2
	140	19.2 ± 0.0	26.9 ± 0.3	24.6 ± 0.2	22.4 ± 0.2	21.2 ± 0.2

The main finding of Block C is that the conclusions of Block B remain robust under non-stationary arrivals. The relative ordering of configurations is unchanged: larger AMR fleets remain the most competitive, and the nine-AMR configuration remains the best-performing TO-BE option at every workload level. At the 140-truck workload, the nine-AMR configuration completes in 16.77 hours, compared with 16.66 hours for the AS-IS benchmark. The gap is therefore slightly larger than the near-parity observed in Block B, but the concentrated arrival pattern does not overturn the comparison.

The principal effect of the non-stationary pattern is higher waiting time for dock assignment, affecting both AS-IS and TO-BE configurations. At the 140-truck workload, waiting time more than doubles for both the AS-IS benchmark and the nine-AMR TO-BE configuration relative to the equivalent stationary scenarios in Block B. This is consistent with the purpose of the non-stationary pattern: although both Blocks B and C are stochastic, Block C concentrates a higher expected share of arrivals in peak intervals, increasing queue formation. Loading time, by contrast, remains essentially unchanged between the two blocks, confirming that the observed differences are driven mainly by dock-assignment waiting rather than by the loading process itself.

Shared-load participation confirms that the auxiliary policy continues to be used under non-stationary arrivals, following the same increasing trend with fleet size observed in Block B. Differences relative to the stationary case are small and not systematic, so this measure is treated as supporting evidence rather than as a primary explanation. Slightly lower participation in some configurations may contribute to small increases in AMR loading time, but dock-assignment waiting remains the main driver of the Block C results. Detailed loading-time and shared-load participation outputs are reported in Table C.7 in Appendix C.3.

Battery-related constraints remain confined to the smaller fleets. The competitive seven- to nine-AMR configurations record no mandatory charging events at any workload level, and in the six-AMR configuration mandatory charging remains marginal, decreasing under non-stationary arrivals, as the lower-intensity periods between arrival peaks provide additional opportunities for opportunistic charging. The full battery and charging outputs are reported in Table C.9 in Appendix C.3.

In summary, Block C confirms that non-stationary arrivals increase dock waiting and slightly reduce TO-BE competitiveness, but do not change the relative comparison between configurations or the effectiveness of the auxiliary AMR policy. Replication-level results are shown in Figures C.7, C.8 and C.9 in Appendix C.3.

5.4 Block D: Sensitivity to Rotation-Heavy Load Composition

Block D examines whether the AS-IS versus TO-BE comparison is sensitive to a rotation-heavy workload composition. The TO-BE configuration retains the auxiliary AMR policy and stationary Poisson arrivals; only the load composition changes, following the rotation-heavy distribution defined in Section 4.8. The analysis focuses on the 120- and 140-truck workloads, where the effect of increased rotation is expected to be more visible, while the 100-truck workload is excluded as a lower-demand case. Given the additional loading complexity, the six-AMR configuration from the balanced case was not re-evaluated; instead, ten- and eleven-AMR configurations were added to the seven- to nine-AMR range to assess whether the required fleet size shifted upward. Table 5.4 reports the selected system-level Performance Measures.

Table 5.4: Block D - Main system-level Performance Measures under the rotation-heavy workload, reported as mean ± 95% confidence interval across 30 replications.

Measure	W	AS-IS	7 AMR	8 AMR	9 AMR	10 AMR	11 AMR
Completion time (h)	120	16.44 ± 0.6	17.20 ± 0.4	16.79 ± 0.5	16.64 ± 0.5	16.54 ± 0.7	16.54 ± 0.7
	140	16.72 ± 0.5	18.84 ± 0.1	17.80 ± 0.3	17.31 ± 0.4	16.98 ± 0.4	16.92 ± 0.4
Truck time in system (min)	120	40.1 ± 1.9	71.2 ± 9.7	57.1 ± 6.8	49.6 ± 4.6	40.1 ± 1.5	40.2 ± 1.3
	140	48.5 ± 5.0	111.7 ± 15.0	83.6 ± 11.9	68.0 ± 9.3	57.9 ± 7.1	56.4 ± 6.6
Waiting for dock (min)	120	4.7 ± 1.8	26.5 ± 9.4	15.9 ± 6.6	10.3 ± 4.4	2.7 ± 1.3	2.6 ± 1.2
	140	12.9 ± 4.9	66.1 ± 14.6	41.5 ± 11.8	28.3 ± 9.2	20.0 ± 7.0	18.8 ± 6.6
Loading time (min)	120	22.7 ± 0.0	31.5 ± 0.3	28.7 ± 0.3	27.3 ± 0.2	25.8 ± 0.1	26.0 ± 0.1
	140	22.6 ± 0.0	32.2 ± 0.2	29.6 ± 0.2	27.7 ± 0.2	26.1 ± 0.2	26.2 ± 0.1

The rotation-heavy workload widens the performance gap between AS-IS and TO-BE relative to the balanced case. Because pallet rotation occurs more frequently, loading times increase for both configurations, with a stronger effect on TO-BE due to the additional handling requirements imposed on AMR loading cycles. As a result, the seven- to nine-AMR range that achieved near-parity under the balanced workload no longer does so: at the 140-truck workload, the nine-AMR configuration completes in 17.31 hours against 16.72 hours for AS-IS, leaving a residual gap of approximately 0.59 hours.

Recovering proximity to the AS-IS benchmark under the rotation-heavy workload requires a larger AMR fleet. The ten-AMR configuration is the key result of this block, reducing the gap at both workload levels. At the 140-truck workload, completion time falls from 17.31 hours with nine AMRs to 16.98 hours with ten AMRs, narrowing the gap to approximately 0.26 hours. At the 120-truck workload, truck time in system falls to 40.1 minutes, effectively matching the AS-IS value. This improvement is mainly driven by lower dock waiting, which drops sharply between the nine- and ten-AMR configurations.

The eleven-AMR configuration was included as a saturation check. It does not provide a conclusive improvement over ten AMRs: completion time is identical at the 120-truck workload and only 0.06 hours lower at 140 trucks, with overlapping confidence intervals. The ten-AMR configuration is therefore the point at which the rotation-heavy workload regains proximity to the AS-IS benchmark, while the eleven-AMR result indicates diminishing returns.

Shared-load participation supports this saturation interpretation. At the 120-truck workload, it reaches 97.5% with ten AMRs, meaning that almost all loads already receive auxiliary support. The limited remaining opportunity for additional assistance helps explain why eleven AMRs do not provide a conclusive improvement. Detailed loading-time and shared-load participation outputs are reported in Table C.10 in Appendix C.4.

Battery-related constraints remain negligible across the evaluated range. Mandatory charging occurs only marginally in the seven-AMR configuration, while the competitive configurations record none at either workload level. The higher prevalence of rotation therefore does not place the relevant TO-BE fleets under battery pressure. Battery and charging outputs are reported in Table C.11 in Appendix C.4.

In summary, Block D shows that a rotation-heavy workload increases TO-BE loading complexity and shifts the AMR fleet size required to approach the AS-IS benchmark from the seven- to nine-AMR range to around ten AMRs. The eleven-AMR configuration does not yield a conclusive further improvement, indicating that additional fleet capacity saturates once shared-load participation approaches its practical limit. Replication-level results are shown in Figures C.10 and C.11 in Appendix C.4.

5.5 Block E: High-Demand Sensitivity Scenario

Block E tests whether the conclusions of the previous blocks remain valid under a substantially higher workload. The scenario uses 200 trucks within the same 16-hour base arrival window, with stationary Poisson arrivals, balanced load composition and the auxiliary AMR policy enabled. The AS-IS configuration uses four forklifts, while the TO-BE configuration is evaluated from six to twelve AMRs. Table 5.5 reports the main system-level Performance Measures.

Table 5.5: High-demand scenario (W200) - Main system-level Performance Measures, reported as mean ± 95% confidence interval across 30 replications.

Measure	AS-IS	6 AMR	7 AMR	8 AMR	9 AMR	10 AMR	11 AMR	12 AMR
Completion time (h)	18.44 ± 0.12	24.50 ± 0.08	21.71 ± 0.09	20.32 ± 0.10	19.31 ± 0.11	18.57 ± 0.12	18.39 ± 0.13	18.01 ± 0.17
Truck time in system (min)	95.8 ± 13.0	259.9 ± 14.9	190.1 ± 15.0	148.4 ± 14.3	119.3 ± 14.4	98.6 ± 13.1	92.8 ± 12.9	80.8 ± 11.5
Waiting for dock (min)	63.9 ± 12.9	217.0 ± 14.9	152.1 ± 15.0	113.0 ± 14.3	85.7 ± 14.4	66.5 ± 13.0	61.2 ± 12.9	50.1 ± 11.5
Loading time (min)	19.13 ± 0.01	28.18 ± 0.11	25.22 ± 0.10	23.25 ± 0.08	21.72 ± 0.11	20.34 ± 0.12	20.08 ± 0.05	19.19 ± 0.00

The overall pattern established in the previous blocks is preserved under the higher workload. TO-BE completion time decreases with fleet size and approaches the AS-IS benchmark around ten to eleven AMRs: the ten-AMR configuration completes in 18.57 hours against 18.44 hours for AS-IS, while the eleven- and twelve-AMR configurations match or slightly exceed the benchmark. Compared with the 140-truck workload, where parity was reached around nine AMRs, the 43% workload increase raises the required fleet by only one to two AMRs. This suggests that fleet size does not scale proportionally with demand because dock capacity becomes the binding constraint. Waiting for dock accounts for roughly two thirds of truck time in system in both configurations, confirming that AS-IS and TO-BE are constrained in parallel by shared dock capacity. Replication-level results are provided in Figure C.12 in Appendix C.5.

AMR utilisation helps explain why fleets below the competitive range are insufficient. The non-working share, combining idle and charging time, increases with fleet size, from 19.1% with six AMRs to 21.1% with seven AMRs and 25.7% with eight AMRs. The six-AMR fleet therefore operates above 80% working time, indicating very limited spare capacity, while the seven-AMR configuration remains close to this threshold. From eight AMRs onwards, spare capacity becomes more evident and can support the auxiliary policy more effectively. This should be read as evidence of undersizing rather than efficiency, since the six- and seven-AMR configurations remain far from the AS-IS benchmark in completion time. Full per-state utilisation outputs are reported in Table C.12 in Appendix C.5.

Battery management does not constrain the competitive configurations. Mandatory charging is concentrated in the six-AMR configuration and becomes marginal with seven AMRs; from eight AMRs onwards, no mandatory charging occurs. This confirms that battery-induced unavailability affects only non-competitive fleets. Battery and charging outputs are reported in Table C.13 in Appendix C.5.

In summary, the high-demand scenario confirms that the conclusions of the main blocks remain valid when demand grows well beyond the reference levels. The relative comparison between the configurations is preserved, the AMR fleet required to match the forklift-based operation grows at a lower rate than workload, as dock capacity acts as the main system constraint and battery management does not limit the competitive configurations.

5.6 Stress-Test Scenario

This section presents an additional stress-test scenario designed to evaluate system behaviour under an extreme capacity condition. Unlike the previous experiments, in which trucks arrive gradually, all load orders are available at the start of the simulation, removing the arrival process as a mediating factor. This does not represent a realistic operating day; instead, it provides an upper-bound reference in which demand is always available and spare loading capacity is immediately used. The scenario uses the 100-truck workload with balanced load composition and the auxiliary AMR policy enabled, comparing AS-IS with TO-BE fleets of eight to twelve AMRs. Table 5.6 reports the main Performance Measures.

Table 5.6: Stress-test scenario - Main Performance Measures with all load orders available at simulation start, reported as mean ± 95% confidence interval across 30 replications.

Measure	AS-IS	8 AMR	9 AMR	10 AMR	11 AMR	12 AMR
Completion time (h)	10.58 ± 0.0	12.25 ± 0.0	11.59 ± 0.0	11.07 ± 0.1	10.96 ± 0.0	10.35 ± 0.0
Loading time (min)	23.0 ± 0.0	30.7 ± 0.1	28.6 ± 0.2	26.9 ± 0.3	26.9 ± 0.0	25.4 ± 0.0
Shared loads (%)	—	66.6 ± 0.9	78.8 ± 1.1	90.1 ± 1.7	90.5 ± 0.3	100.0 ± 0.0

Because all load orders are available from the start, the system operates under continuous backlog: a queue of trucks waiting for dock access persists throughout the simulation, and the waiting time for dock assignment is consequently very large for all configurations. These waiting times are an artefact of the stress-test design rather than a measure of operational performance, and the analysis therefore focuses on total completion time and loading time, which reflect the processing capacity of each configuration.

Under continuous backlog, increasing the AMR fleet produces clearer improvements than in the Poisson-arrival blocks. Completion time decreases from 12.25 hours with eight AMRs to 10.35 hours with twelve AMRs, because additional loading capacity is immediately absorbed by the waiting queue. This differs from the previous blocks, where extra fleet capacity eventually brought little benefit once queueing was no longer the dominant constraint.

A notable result is that TO-BE can match or exceed AS-IS when the system is limited mainly by processing capacity. The eleven- and twelve-AMR configurations complete the workload in 10.96 and 10.35 hours, respectively, against 10.58 hours for AS-IS. Shared-load participation supports this interpretation, reaching 100% with twelve AMRs and showing that cooperative loading is fully used. No relevant mandatory charging constraints arose for the fleet sizes tested in this scenario; the full battery and charging outputs are reported in Table C.14 in Appendix C.6.

In summary, the stress-test scenario shows that, under continuous demand, additional AMRs continue to improve performance across the evaluated range, allowing a sufficiently large TO-BE fleet to match or exceed AS-IS throughput. These observations are revisited in the integrated discussion in Section 5.7. Replication-level results are shown in Figure C.13 in Appendix C.6.

5.7 Synthesis of Findings and Operational Implications

This section synthesises the findings that span the experimental blocks and discusses their operational implications for the forklift-to-AMR transition. Rather than restating each block separately, the focus is on the conditions under which the TO-BE configuration becomes competitive, the interpretation of the main Performance Measures, and the implications for fleet-sizing decisions.

The first cross-cutting finding is that a one-to-one replacement of forklifts by AMRs is not viable in the studied context. Matching the four-forklift AS-IS benchmark consistently required a larger AMR fleet, with the required size depending on demand level, arrival pattern and load composition. Under balanced conditions, the competitive range varied with workload; under rotation-heavy loads and high demand, additional AMRs were required. This shows that the substitution decision cannot rely on a fixed fleet-conversion ratio, but must be treated as a fleet-sizing problem in which the decision-maker evaluates the performance gap, fleet size and operating conditions together. This numerical difference should not be interpreted only as a disadvantage of the AMR-based configuration. The AS-IS fleet was kept fixed to preserve a stable operational benchmark, and increasing the number of forklifts would represent a different operational strategy. Unlike AMRs, which can be coordinated through a shared control logic within an automated fleet-management system, adding more human-operated forklifts in a constrained loading area may increase safety, coordination and labour-availability challenges.

The second finding concerns the interpretation of mean total completion time. Under progressive arrivals, this measure is partly conditioned by the arrival window and by the timing of the last trucks entering the system, rather than by processing capacity alone. Configurations with similar completion times may therefore differ in their operational behaviour, as shown by waiting for dock assignment, truck time in system and loading time. The stress-test scenario illustrates this contrast: when demand is available from the start, completion time more directly reflects processing capacity. Total completion time should therefore be read alongside the other Performance Measures, since relying on it alone can mask relevant differences in flow.

The third finding relates to the auxiliary AMR policy, which is the main mechanism through which the TO-BE configuration becomes competitive. By allowing spare AMR capacity to be assigned within ongoing loads, the policy reduces per-truck loading time and helps close the gap to the AS-IS benchmark. However, its system-level benefit depends on the demand regime: under progressive arrivals, once sustained queues are reduced, further loading-time improvements have a smaller effect on mean total completion time; under continuous backlog, as shown by the stress-test scenario, additional AMRs are absorbed more directly by the system. Fleet sizing should therefore account not only for loading capacity, but also for whether demand conditions allow that capacity to be fully used.

The fourth finding concerns battery and charging behaviour. Across the tested competitive configurations, battery management did not constrain performance, and mandatory charging was largely confined to the smaller fleets. This outcome is partly a consequence of the modelled operational logic: the interval between dock assignment and the next truck reaching the dock, represented by a triangular distribution of $(8, 10, 12)$ minutes, creates windows during which idle AMRs can charge opportunistically. In configurations with smaller fleets, or in an operation where a dock were reoccupied without such an interval, energy availability could become more relevant. The robustness observed here should therefore be understood as holding for the implemented operational logic and the competitive fleet sizes tested, rather than as a general property of the system.

Overall, these findings support the use of the simulation model as a decision-support tool for comparing feasible AMR fleet configurations under different operational conditions, rather than as a single deterministic replacement rule.

Chapter 6

Conclusions and Future Work

This chapter closes the study by summarising the main findings of the simulation analysis, discussing their implications for the evaluation of alternative outbound loading strategies. It also outlines directions for future work and potential extensions to the simulation model.

6.1 Conclusions

This dissertation evaluated the transition from a forklift-based outbound loading operation to an AMR-based alternative through a FlexSim discrete-event simulation model. The model was used as a scenario-comparison and decision-support tool, testing representative operating conditions such as fleet size, auxiliary AMR use, arrival patterns, load composition, high demand and continuous backlog. Its main value was enabling controlled analysis of configurations that would be difficult or disruptive to test directly in the real system.

In response to the main research question, the results indicate that replacing forklifts with AMRs is operationally viable in the studied context, but that doing so requires a fleet numerically larger than the forklift fleet it replaces. The size of this difference is not constant: matching four forklifts required a fleet roughly one and a half to two and a half times larger, depending on the workload level, the arrival pattern, and the prevalence of rotation in the load composition. The transition is thus best understood as a trade-off rather than a direct equivalence, with the required AMR fleet growing as operational complexity and demand pressure increase.

It is also important to note that the comparison developed here focused primarily on time-based operational performance, while the decision to adopt AMRs involves dimensions that the model does not fully capture. Autonomous mobile robots may reduce human error in the loading process, improve safety by removing operators from the loading area, and offer greater traceability, reconfiguration flexibility, and integration with digital warehouse management systems. These advantages do not offset the need for a larger fleet to match forklift time performance, but they may support the transition on grounds beyond throughput alone, particularly in connection with Sustainable Development Goal 8, Decent Work and Economic Growth, and Sustainable Development Goal 9, Industry, Innovation and Infrastructure.

Beyond the specific comparison addressed here, a broader contribution of this work lies in the simulation model itself. The main technical challenge was to design and verify a model that coherently represents the key operational conditions of the outbound loading process. Once this architecture was established, evaluating further configurations became a matter of adjusting inputs and parameters rather than redesigning the model. While each new scenario still requires validation and interpretation in light of its operational context, the model can serve as a reusable decision-support framework for analysing outbound loading configurations beyond those examined here.

The scope of the study reflects a set of deliberate modelling choices consistent with its purpose as a scenario-comparison and decision-support tool, rather than a predictive reproduction of a specific operational period. The model was built on representative and operationally plausible inputs, structured around eight load typologies that capture the main variability of the loading operation in a controlled manner. Resource movement was represented through the path network as a supporting infrastructure rather than as a decision variable, and the arrival patterns were defined by design to provide controlled and well-defined demand conditions. These boundaries make the comparison tractable and the results interpretable, while defining the scope within which the findings should be understood.

6.2 Future Work

Future work should extend the model in directions that are both operationally relevant for the case study and connected to extensions commonly discussed in the literature. A first and relatively direct extension would be to test alternative truck-release, dock-assignment and dock-reoccupation policies. Truck scheduling and dock-door assignment have been studied in cross-docking contexts, where decisions such as truck time windows, flexible dock-door use and temporary buffering affect operational performance (Bodnar et al., 2017). In the present outbound-loading model, related experiments could vary the truck-release rule, the delay between dock assignment and dock occupation, or the logic used to allocate trucks to docks, assessing how these decisions affect dock waiting, AMR utilisation, battery availability and the AS-IS versus TO-BE comparison. This extension would be direct to implement because the current model already represents truck arrivals, dock queues, dock assignment and truck-to-dock travel times.

A second extension would be to move from the current first-available assignment logic towards more advanced AMR task-assignment rules. Fragapane et al. (2021) identifies scheduling, dispatching and resource management as central planning and control decisions for AMRs in intralogistics, while Lee (2018) proposes a resource-based task allocation approach in which robot resource levels and charging requirements are considered during task assignment. In this model, this could be translated into resource-aware dispatching policies where AMR battery level, expected charging needs, dock availability and load complexity are considered when selecting the robot for each task. This would be moderately direct to implement, since the model already tracks

AMR availability, battery state and auxiliary assignment, but it would require more sophisticated decision rules than the current rule-based allocation.

Finally, the model could be complemented with an economic cost-benefit analysis of the AMR fleet, including investment cost, operating cost, maintenance, labour implications and potential safety or quality benefits. This would not require major changes to the simulation logic, but would require reliable cost inputs and assumptions to translate the operational Performance Measures into investment indicators. Such an extension would move the model from an operational performance comparison towards an investment-level decision-support tool, helping decision-makers assess not only whether an AMR configuration can match the forklift-based benchmark, but also whether the required fleet size is economically justified.

Other extensions discussed in the AMR planning and control literature, such as detailed routing or advanced congestion management, would also be relevant but less direct to implement, as they would require a more detailed representation of AMR traffic interactions than the one developed in this dissertation.

Bibliography

Albert, B., Chaikovskaia, M., Gayon, J. P., and Quilliot, A. (2025). Sizing and scheduling for a fleet of reconfigurable mobile robots. *International Journal of Production Research*, 63(18):6758–6775.

Alonso, M. T., Alvarez-Valdes, R., Iori, M., and Parreño, F. (2019). Mathematical models for multi container loading problems with practical constraints. *Computers & Industrial Engineering*, 127:722–733.

Alonso, M. T., Álvarez-Valdés, R., Iori, M., Parreño, F., and Tamarit, J. M. (2017). Mathematical models for multicontainer loading problems. *Omega*, 66:106–117.

Alonso, M. T., Martinez-Sykora, A., Álvarez-Valdés, R., and Parreño, F. (2022). The pallet-loading vehicle routing problem with stability constraints. *European Journal of Operational Research*, 302(3):860–873.

Azadeh, K., de Koster, R., and Roy, D. (2019). Robotized and automated warehouse systems: Review and recent developments. *Transportation Science*, 53(4):917–945.

Banks, J., editor (1998). *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*. John Wiley & Sons.

Bodnar, P., de Koster, R., and Azadeh, K. (2017). Scheduling trucks in a cross-dock with mixed service mode dock doors. *Transportation Science*, 51(1):112–131.

Bolu, A. and Korçak, Ö. (2021). Adaptive task planning for multi-robot smart warehouse. *IEEE Access*, 9:27346–27358.

Brailsford, S. C., Harper, P. R., Patel, B., and Pitt, M. (2019). Hybrid simulation modelling in operational research: A state-of-the-art review. *European Journal of Operational Research*, 278:721–737.

Descartes Systems Group (2024). Descartes' study reveals 76% of supply chain and logistics operations are experiencing notable workforce shortages. Press release. Accessed: 2026-06-23.

Fragapane, G., De Koster, R., Sgarbossa, F., and Strandhagen, J. O. (2021). Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. *European Journal of Operational Research*, 294(2):405–426.

Gajda, M., Trivella, A., Mansini, R., and Pisinger, D. (2022). An optimization approach for a complex real-life container loading problem. *Omega*, 107:102559.

Ghelichi, Z. and Kilaru, S. (2021). Analytical models for collaborative autonomous mobile robot solutions in fulfillment centers. *Applied Mathematical Modelling*, 91:438–457.

International Federation of Robotics (2025). World robotics 2025 report – service robots: Service robots see global growth boom. Press release. Accessed: 2026-06-23.

Kim, J., Kang, J., and Jung, H. (2026). A simulation-based deep learning approach to operational robot fleet sizing in robotic mobile fulfillment systems. *Computers & Industrial Engineering*, 214:111910.

Kleijnen, J. P. C. (1995). Verification and validation of simulation models. *European Journal of Operational Research*, 82(1):145–162.

Law, A. M. (2015). *Simulation Modeling and Analysis*. McGraw-Hill Education, 5 edition.

Lee, D.-H. (2018). Resource-based task allocation for multi-robot systems. *Robotics and Autonomous Systems*, 103:151–161.

Li, X., Yang, X., Zhang, C., and Qi, M. (2021). A simulation study on the robotic mobile fulfillment system in high-density storage warehouses. *Simulation Modelling Practice and Theory*, 112:102366.

Luo, L. and Zhao, N. (2022). An efficient simulation model for layout and mode performance evaluation of robotic mobile fulfillment systems. *Expert Systems with Applications*, 203:117492.

Macal, C. M. and North, M. J. (2005). Tutorial on agent-based modeling and simulation. In Kuhl, M. E., Steiger, N. M., Armstrong, F. B., and Joines, J. A., editors, *Proceedings of the 2005 Winter Simulation Conference*, pages 2–15.

Pollaris, H., Braekers, K., Caris, A., Janssens, G. K., and Limbourg, S. (2015). Vehicle routing problems with loading constraints: state-of-the-art and future directions. *OR Spectrum*, 37(2):297–330.

Pollaris, H., Braekers, K., Caris, A., Janssens, G. K., and Limbourg, S. (2016). Capacitated vehicle routing problem with sequence-based pallet loading and axle weight constraints. *EURO Journal on Transportation and Logistics*, 5(2):231–255.

Ramos, A. G., Oliveira, J. F., and Lopes, M. P. (2016). A physical packing sequence algorithm for the container loading problem with static mechanical equilibrium conditions. *International Transactions in Operational Research*, 23(1-2):215–238. First published online 2014-10-08.

Ramos, A. G., Silva, E., Oliveira, J. F., Gonçalves, J. F., and Lopes, M. P. (2018). A new load balance methodology for container loading problem in road transportation. *European Journal of Operational Research*, 266(3):1140–1152.

Rijal, A., Bijvank, M., and de Koster, R. B. M. (2023). Dynamics between warehouse operations and vehicle routing. *Production and Operations Management*, 32(11):3575–3593.

Sargent, R. G. (2010). Verification and validation of simulation models. In *Proceedings of the 2010 Winter Simulation Conference*, pages 166–183. IEEE.

Tadumadze, G. and Emde, S. (2022). Loading and scheduling outbound trucks at a dispatch warehouse. *IISE Transactions*, 54(8):770–784.

Wei, J. and Leung, S. C. H. (2011). A simulation modeling and analysis for rfid-enabled mixed-product loading strategy for outbound logistics: A case study. *Computers & Industrial Engineering*, 61(1):209–215.

Appendix A

Layout and Load Typology Visuals

A.1 Outbound Dock and Conveyor Layout

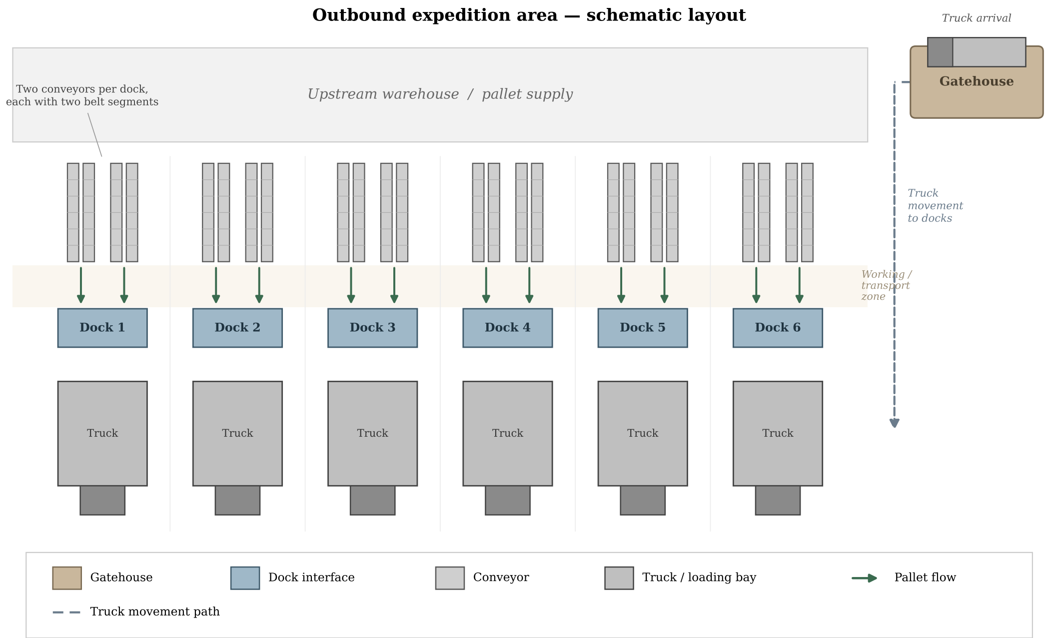

Figure A.1: Schematic representation of the outbound expedition area with six loading docks, each served by two conveyors composed of two belt segments.

A.2 Illustrative Load Typologies

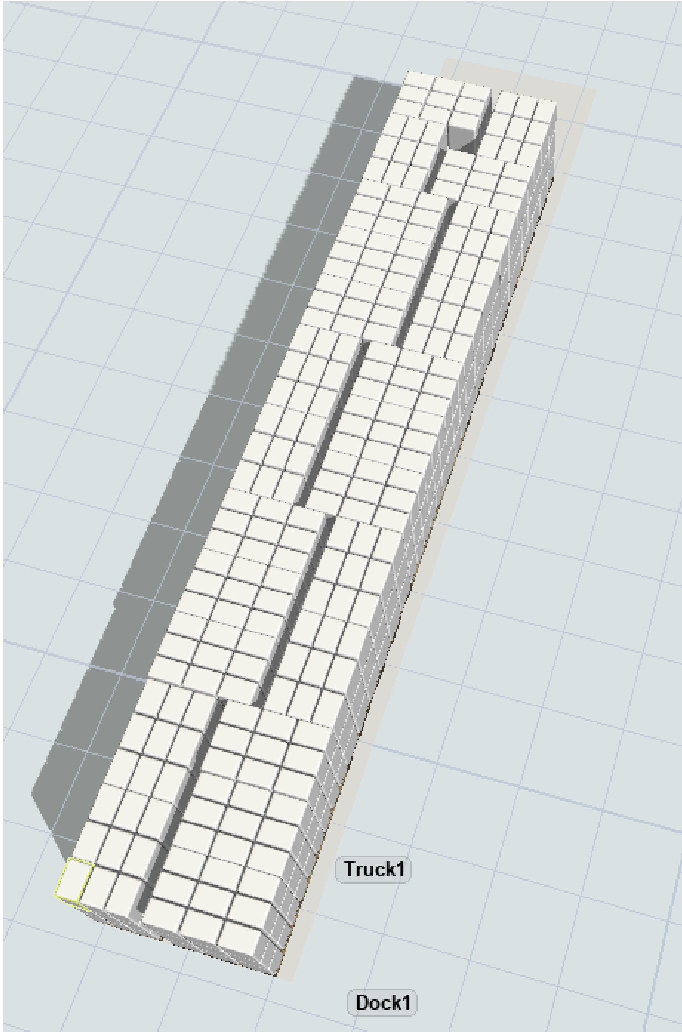

Figure A.2: Illustrative 3D representation of load Typology 2, including both rotated and non-rotated pallets.

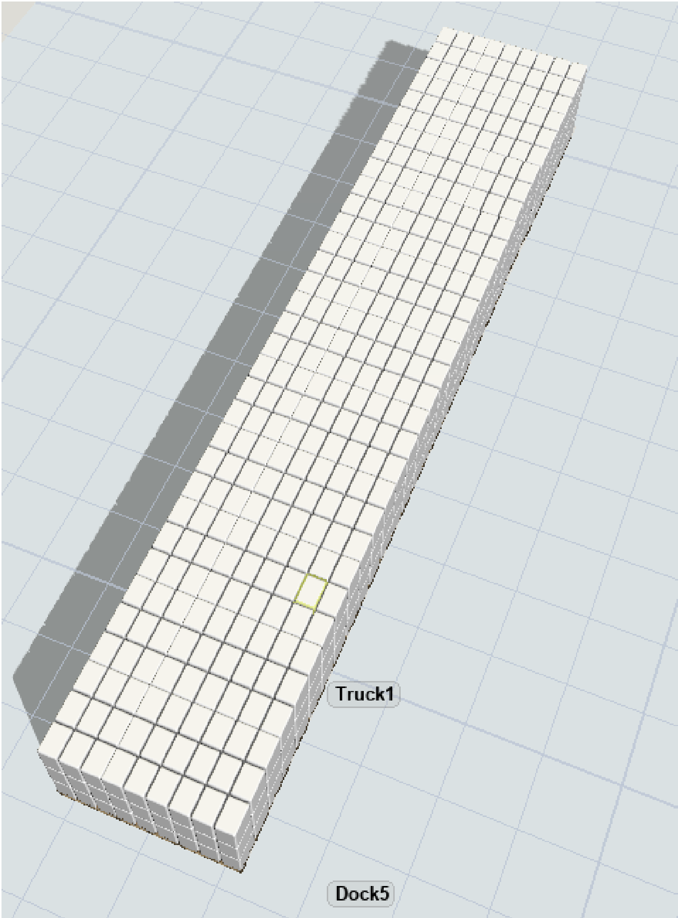

Figure A.3: Illustrative 3D representation of load Typology 8, in which all pallets preserve the same orientation in which they leave the conveyor.

Appendix B

FlexSim Model Supporting Materials

B.1 Detailed Process Flow Implementations

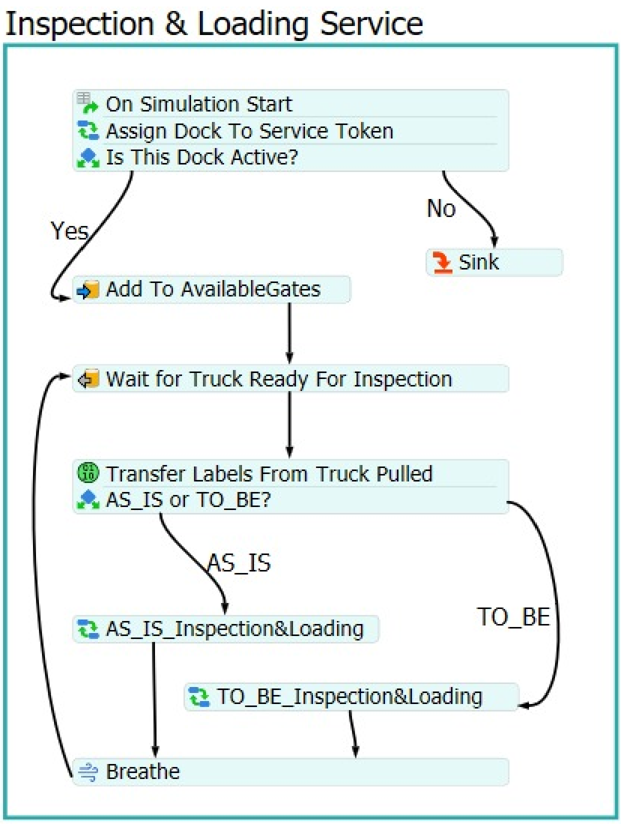

Figure B.1: FlexSim Process Flow implementation detail showing the convergence at the inspection step and the `ScenarioType` decision block.

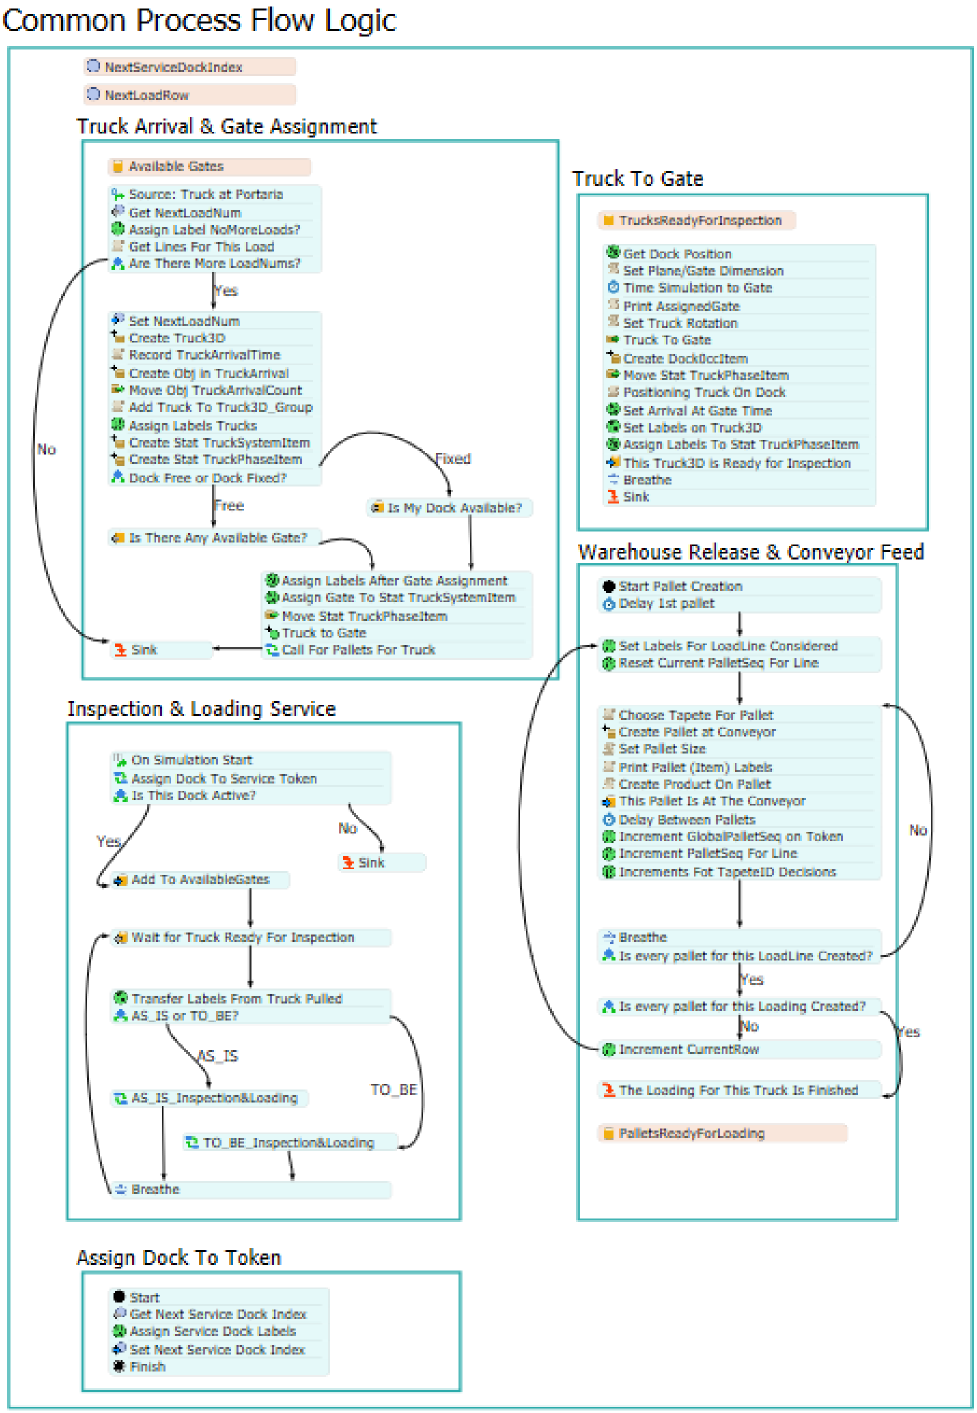

Figure B.2: Detailed common Process Flow implementation.

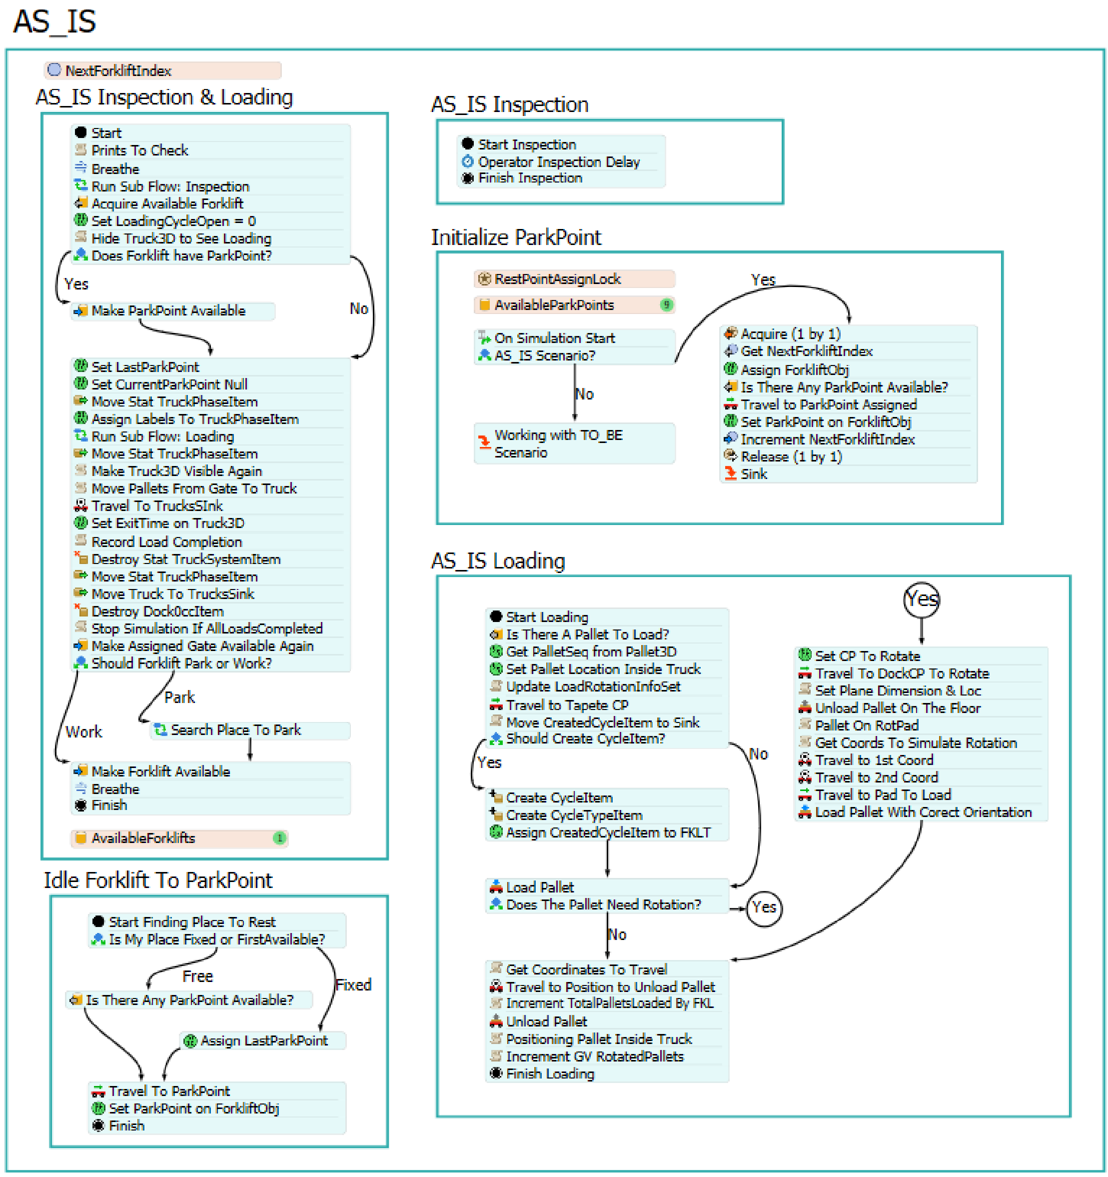

Figure B.3: Detailed AS-IS Process Flow implementation.

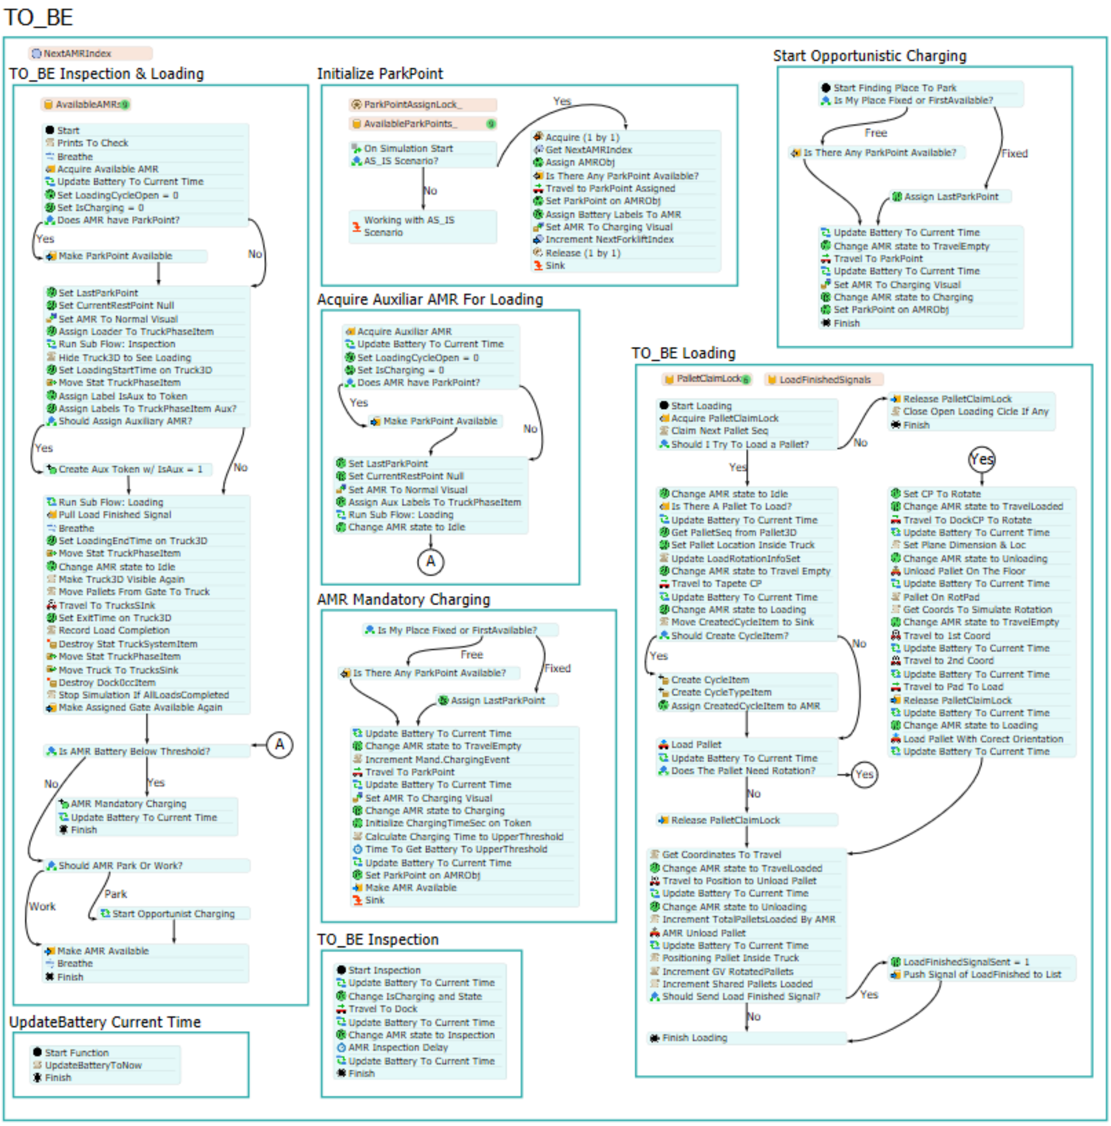

Figure B.4: Detailed TO-BE Process Flow implementation.

B.2 FlexSim Parameter Structure

Table B.1: State-specific AMR battery depletion rates (TO-BE configuration).

Parameter	Value	Unit	Description
AMR_Batt_TravelLoaded	0.17	%/min	Depletion rate during loaded travel
AMR_Batt_TravelEmpty	0.11	%/min	Depletion rate during empty travel
AMR_Batt_Inspection	0.11	%/min	Depletion rate during inspection
AMR_Batt_Loading	0.17	%/min	Depletion rate during pallet pickup
AMR_Batt_Unloading	0.17	%/min	Depletion rate during pallet placement
AMR_Batt_Idle	0.04	%/min	Depletion rate when idle

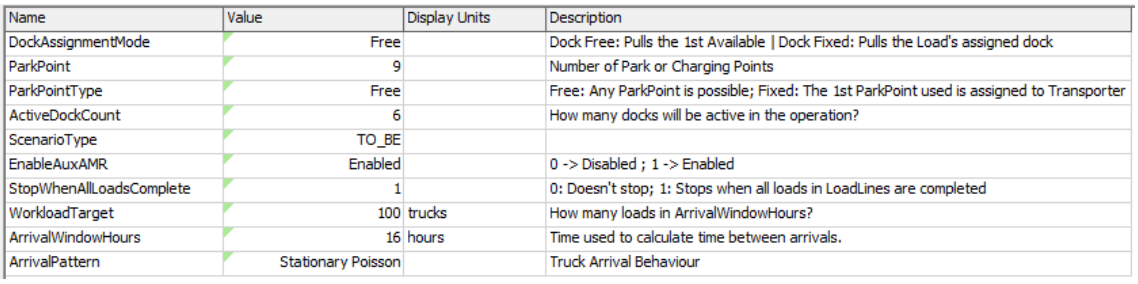

Name	Value	Display Units	Description
DockAssignmentMode	Free		Dock Free: Pulls the 1st Available | Dock Fixed: Pulls the Load's assigned dock
ParkPoint	9		Number of Park or Charging Points
ParkPointType	Free		Free: Any ParkPoint is possible; Fixed: The 1st ParkPoint used is assigned to Transporter
ActiveDockCount	6		How many docks will be active in the operation?
ScenarioType	TO_BE		
EnableAuxAMR	Enabled		0 -> Disabled ; 1 -> Enabled
StopWhenAllLoadsComplete	1		0: Doesn't stop; 1: Stops when all loads in LoadLines are completed
WorkloadTarget	100	trucks	How many loads in ArrivalWindowHours?
ArrivalWindowHours	16	hours	Time used to calculate time between arrivals.
ArrivalPattern	Stationary Poisson		Truck Arrival Behaviour

Figure B.5: General model parameters implemented in FlexSim.

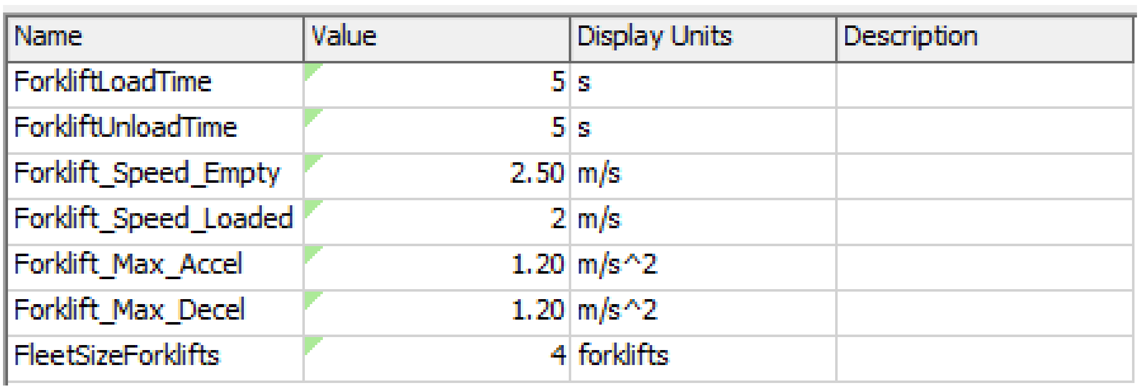

Name	Value	Display Units	Description
ForkliftLoadTime	5	s	
ForkliftUnloadTime	5	s	
Forklift_Speed_Empty	2.50	m/s	
Forklift_Speed_Loaded	2	m/s	
Forklift_Max_Accel	1.20	m/s^2	
Forklift_Max_Decel	1.20	m/s^2	
FleetSizeForklifts	4	forklifts	

Figure B.6: Forklift resource parameters implemented in FlexSim.

Name	Value	Display Units	Description
AMRLoadTime	15	s	
AMRUnloadTime	10	s	
InitialBatteryLevel	100	%	
AMR_LowerBatteryThreshold	20	%	
AMR_UpperBatteryThreshold	80	%	
AMR_ChargeRate_0_80	0.73	%/min	
AMR_ChargeRate_80_100	0.29	%/min	
AMR_Batt_TravelLoaded	0.17	%/min	
AMR_Batt_TravelEmpty	0.11	%/min	
AMR_Batt_Inspection	0.11	%/min	
AMR_Batt_Loading	0.17	%/min	
AMR_Batt_Unloading	0.17	%/min	
AMR_Batt_Idle	0.04	%/min	
AMR_SpeedEmpty	1.50	m/s	
AMR_SpeedLoaded	1	m/s	
AMR_Max_Accel	0.30	m/s^2	
AMR_Max_Decel	0.30	m/s^2	
FleetSizeAMRs	9	AMRs	

Figure B.7: AMR resource, battery and charging parameters implemented in FlexSim.

B.3 Arrival Process Parameters

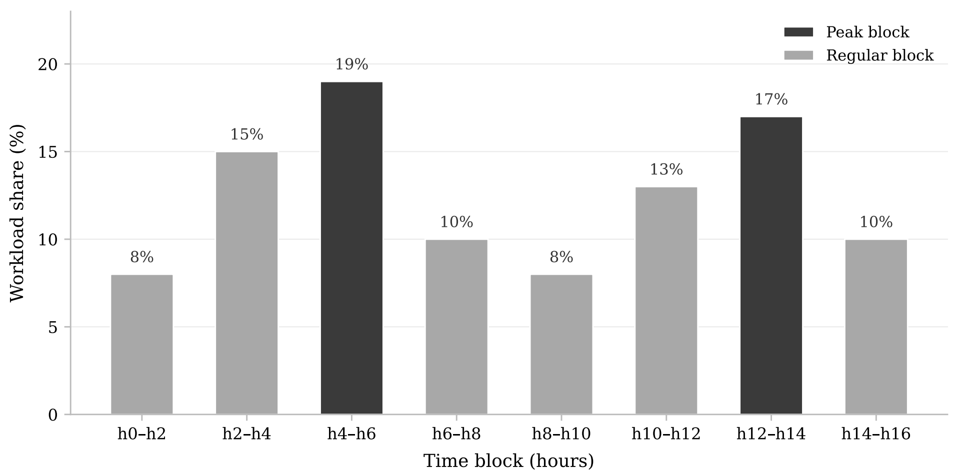

Figure B.8: Non-stationary arrival profile used to distribute truck arrivals across the 16-hour operational window.

Table B.2: Non-stationary Poisson arrival profile for the 100-truck workload.

Block	Workload share	E[trucks]	λ (trucks/h)	Mean interarrival
h0–h2	8%	8.0	4.0	15.0 min
h2–h4	15%	15.0	7.5	8.0 min
h4–h6	19%	19.0	9.5	6.3 min
h6–h8	10%	10.0	5.0	12.0 min
h8–h10	8%	8.0	4.0	15.0 min
h10–h12	13%	13.0	6.5	9.2 min
h12–h14	17%	17.0	8.5	7.1 min
h14–h16	10%	10.0	5.0	12.0 min
Total	**100%**	**100.0**		

Table B.3: Non-stationary Poisson arrival profile for the 140-truck workload.

Block	Workload share	E[trucks]	λ (trucks/h)	Mean interarrival
h0–h2	8%	11.2	5.6	10.7 min
h2–h4	15%	21.0	10.5	5.7 min
h4–h6	19%	26.6	13.3	4.5 min
h6–h8	10%	14.0	7.0	8.6 min
h8–h10	8%	11.2	5.6	10.7 min
h10–h12	13%	18.2	9.1	6.6 min
h12–h14	17%	23.8	11.9	5.0 min
h14–h16	10%	14.0	7.0	8.6 min
Total	**100%**	**140.0**		

B.4 FlexSim Input Data Tables

Table B.4: Representative excerpt of the `LoadNumData` table used as input to the FlexSim model.

LoadNum	Vehicle	Truck dims.	Max wt.	Dock	Conv.	Line	Product	ReqQty	Pallet	Pallet dims.	Layout
LD001	TK11	250×1350×242	24800	4	1	1	XXY	28	H1P	780×980×140	L28
LD002	TK11	250×1350×242	24800	2	1	1	XYY	21	S11	660×920×135	L21
LD003	TK11	250×1350×242	24800	3	1	1	YXX	27	G10P	690×880×139	L27
LD004	TK02	250×1350×242	24800	6	1	1	YYX	8	A7P	800×1200×151	L29
LD004	TK02	250×1350×242	24800	6	2	2	YXY	1	A7P	800×1200×151	L29
LD004	TK02	250×1350×242	24800	6	3	3	XXY	20	A7P	800×1200×151	L29

Table B.5: Complete loading-plan table for Typology 2, corresponding to layout `L24`. The X and Y coordinates are defined relative to the truck loading space and are expressed in millimetres.

Seq.	X	Y	Rotation
1	0	100	0
2	0	1040	1
3	800	1040	1
4	1200	100	0
5	1600	1040	1
6	2400	100	1
7	2400	1440	0
8	3200	100	1
9	3600	1440	0
10	4000	100	1
11	4800	100	0
12	4800	1040	1
13	5600	1040	1
14	6000	100	0
15	6400	1040	1
16	7200	100	1
17	7200	1440	0
18	8000	100	1
19	8800	1440	0
20	8800	100	1
21	9600	100	0
22	9600	1040	1
23	10400	1440	0
24	10800	100	1

B.5 Dashboard Layouts

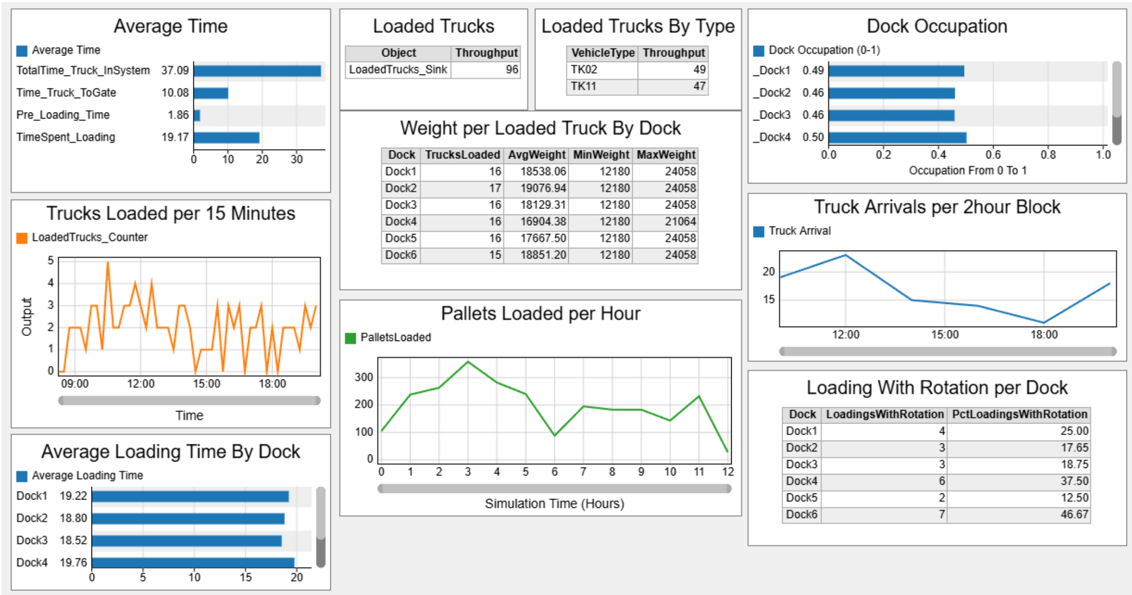

Figure B.9: Loading dashboard layout used for system-level monitoring.

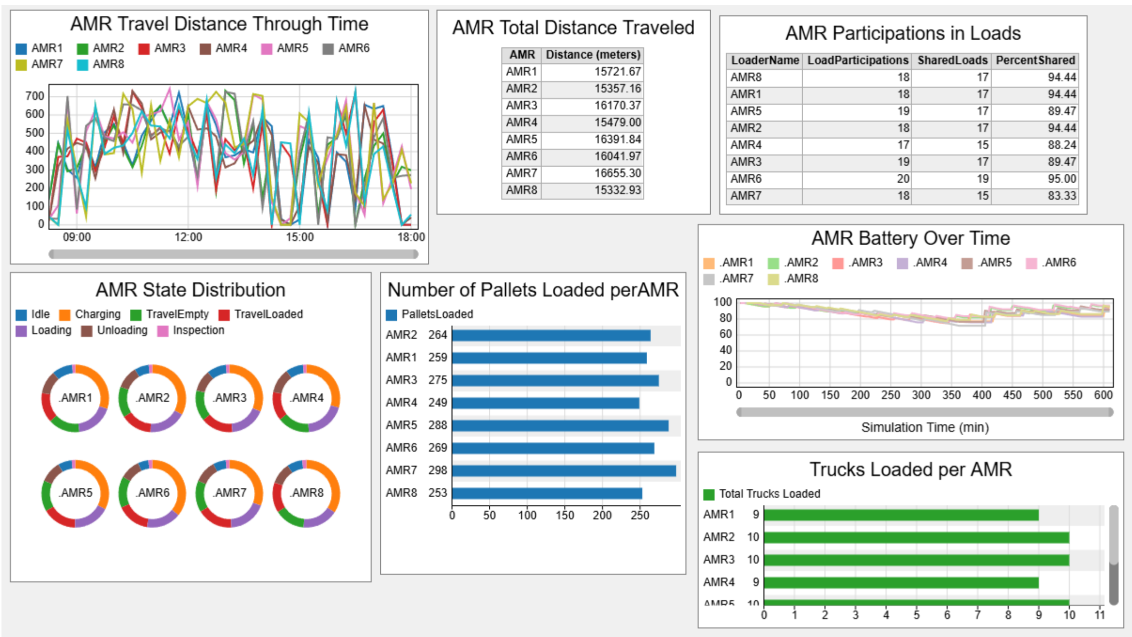

Figure B.10: AMR dashboard layout used for monitoring TO-BE resource behaviour.

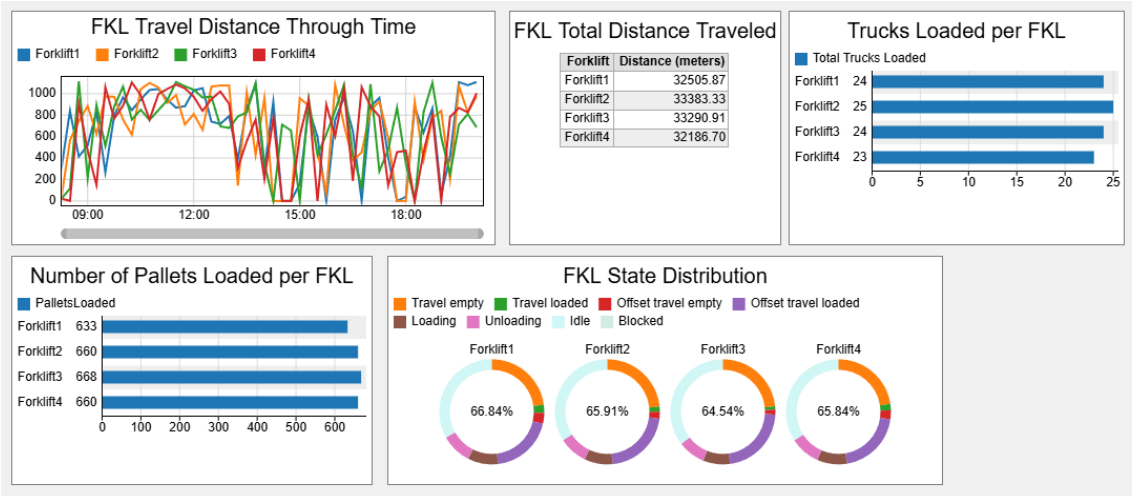

Figure B.11: Forklift dashboard layout used for monitoring AS-IS resource behaviour.

B.6 Implemented Performance Measures

Table B.6: System-level Performance Measures implemented in FlexSim.

Performance Measure	Display units	Description
TotalCompletionTime	h	Time required to complete all load orders in the scenario.
TrucksCompleted	–	Number of trucks/load orders completed during the replication.
Average Truck Total Time In System	s	Average time spent by trucks in the system, from arrival to completion.
Average Truck Total Time In System (hours)	h	Average truck time in system converted to hours.
Average Waiting Time for Dock Assignment	s	Average time spent by trucks waiting for dock assignment.
Average Time To Get To Gate	s	Average time required for the truck to travel to the assigned dock/gate after dock assignment.
Average Gate Pre Loading Time	s	Average time between truck arrival at the dock/-gate and start of loading.
Average Loading Time	s	Average duration of the loading execution phase at truck level.
Average Dock Occupation	ratio	Average dock occupation level, where 0 corresponds to 0% and 1 corresponds to 100%.
TotalPalletsLoaded	pallets	Total number of pallets loaded during the replication.
Pct of Pallets Rotated	%	Percentage of loaded pallets requiring rotation.
LastTruckArrivalTime	h	Arrival time of the last truck generated in the replication.
Average SharedLoads Percentage	%	Percentage of loads assisted by more than one AMR when the auxiliary AMR policy is enabled.

Table B.7: Resource-level and pallet-cycle Performance Measures implemented in FlexSim.

Performance Measure	Display units	Description
Resource Total Distance	meters	Total distance travelled by the resources considered in the scenario.
Average Distance per Resource	meters	Average distance travelled per resource.
Resource Pallets Loaded Balance	–	Difference between the maximum and minimum number of pallets loaded per resource.
Resource Average Pallets Loaded	loaded pallets	Average number of pallets loaded per resource.
Resource Trucks Loaded Balance	–	Difference between the maximum and minimum number of trucks/load orders served per resource.
Average Pallet Loading Time Cycle	s	Average time from reaching the conveyor, unloading at the truck and returning to the conveyor.
Average NonRotated Pallet Loading Time Cycle	s	Average pallet loading cycle time for pallets that do not require rotation.
Average Rotated Pallet Loading Time Cycle	s	Average pallet loading cycle time for pallets that require rotation.

Table B.8: Battery-related Performance Measures implemented in FlexSim.

Performance Measure	Display units	Description
Min Battery Level Observed	%	Minimum battery level observed among the AMRs during the replication.
Final Average Battery Level	%	Average AMR battery level at the end of the replication.
Average Battery Level	%	Average battery level of the AMR group during the replication.
Average State Charging Percentage	%	Average percentage of time spent by AMRs in the charging state.
Mandatory Charging Events	events	Number of mandatory charging events triggered during the replication.

Appendix C

Detailed Results and Performance Measure Outputs

C.1 Block A: Baseline Fleet-size Comparison

C.1.1 Replication-Level Completion Time Outputs

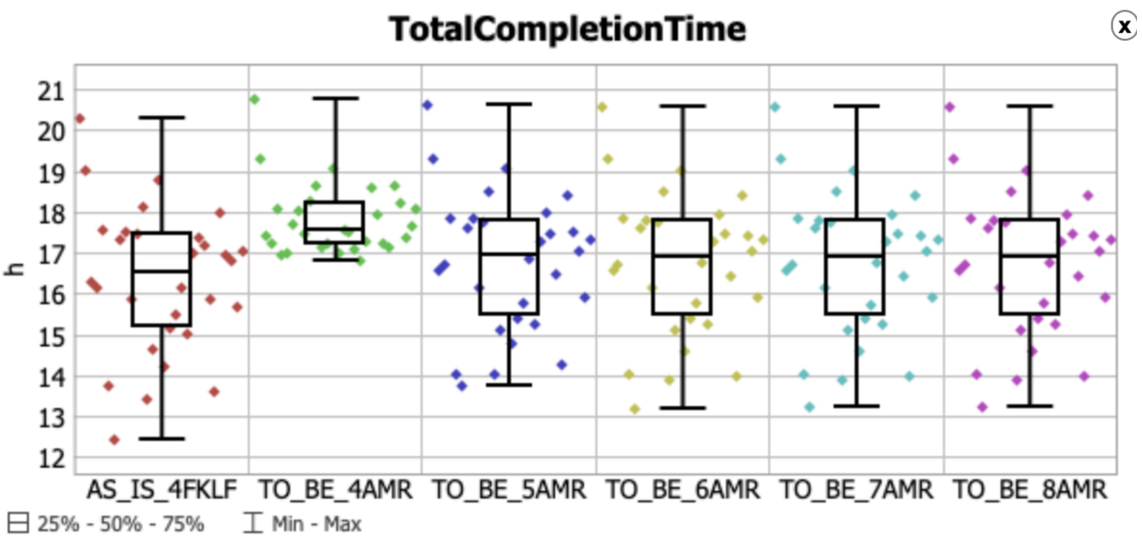

Figure C.1: Replication-level TotalCompletionTime results for Block A under the 100-truck workload.

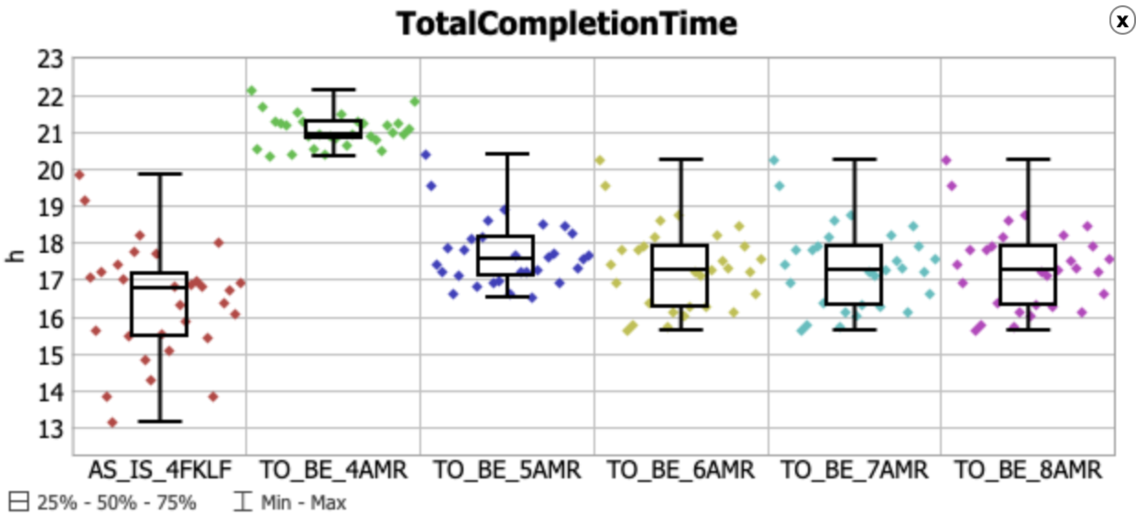

Figure C.2: Replication-level TotalCompletionTime results for Block A under the 120-truck workload.

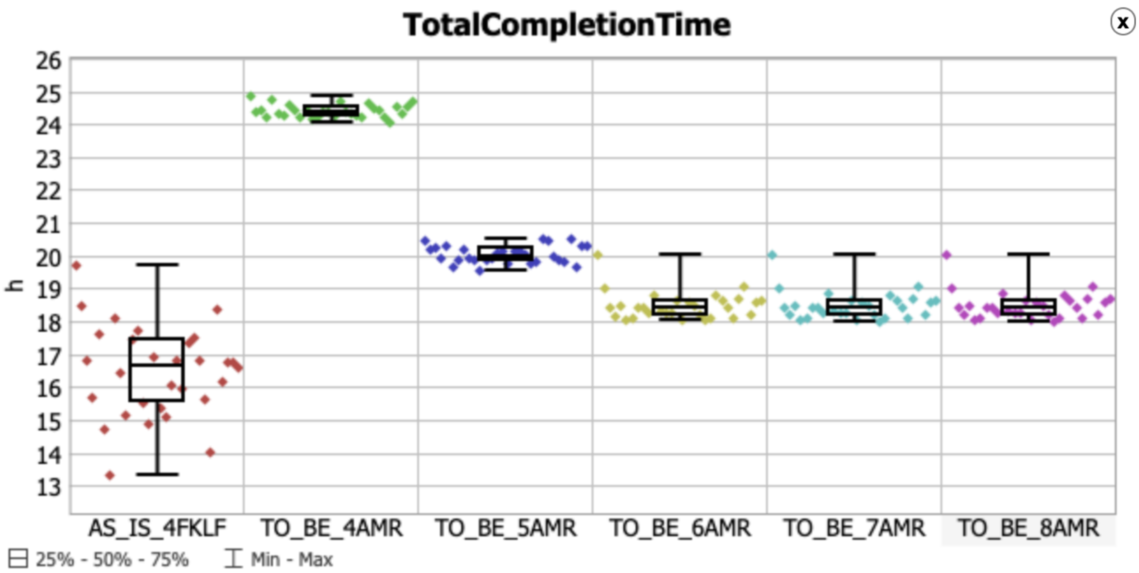

Figure C.3: Replication-level TotalCompletionTime results for Block A under the 140-truck workload.

C.1.2 Detailed Loading-Time Outputs

Table C.1: Detailed loading-time outputs for Block A, reported as mean ± 95% confidence interval (minutes) across 30 replications.

Scenario	W	Loading time	Pallet cycle	Non-rotated cycle	Rotated cycle
AS_IS_4FKLF	100	19.25 ± 0.01	0.72 ± 0.00	0.65 ± 0.00	1.01 ± 0.00
TO_BE_4AMR	100	33.67 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_5AMR	100	33.66 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_6AMR	100	33.64 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_7AMR	100	33.65 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_8AMR	100	33.65 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
AS_IS_4FKLF	120	19.18 ± 0.01	0.72 ± 0.00	0.65 ± 0.00	1.01 ± 0.00
TO_BE_4AMR	120	33.53 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_5AMR	120	33.51 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_6AMR	120	33.50 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_7AMR	120	33.50 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_8AMR	120	33.50 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
AS_IS_4FKLF	140	19.24 ± 0.01	0.72 ± 0.00	0.65 ± 0.00	1.01 ± 0.00
TO_BE_4AMR	140	33.67 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_5AMR	140	33.66 ± 0.01	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_6AMR	140	33.65 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_7AMR	140	33.65 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00
TO_BE_8AMR	140	33.65 ± 0.00	1.27 ± 0.00	1.13 ± 0.00	1.85 ± 0.00

C.1.3 Battery and Charging Outputs

Table C.2: Battery and charging outputs for Block A, reported as mean ± 95% confidence interval across 30 replications.

Scenario	W	Avg battery level (%)	Avg charging state (%)	Mandatory charging events
AS_IS_4FKLF	100	–	–	–
TO_BE_4AMR	100	59.8 ± 3.3	16.6 ± 1.4	2.2 ± 0.7
TO_BE_5AMR	100	78.7 ± 3.6	27.1 ± 2.8	0.0 ± 0.0
TO_BE_6AMR	100	89.8 ± 1.7	38.5 ± 2.5	0.0 ± 0.0
TO_BE_7AMR	100	94.3 ± 0.6	47.0 ± 2.2	0.0 ± 0.0
TO_BE_8AMR	100	95.3 ± 0.1	53.4 ± 1.9	0.0 ± 0.0
AS_IS_4FKLF	120	–	–	–
TO_BE_4AMR	120	55.0 ± 0.7	16.7 ± 0.5	7.0 ± 0.5
TO_BE_5AMR	120	65.4 ± 2.5	18.7 ± 1.3	1.3 ± 0.7
TO_BE_6AMR	120	83.7 ± 1.2	29.6 ± 1.7	0.0 ± 0.0
TO_BE_7AMR	120	91.7 ± 0.9	39.4 ± 1.5	0.0 ± 0.0
TO_BE_8AMR	120	95.2 ± 0.1	46.7 ± 1.3	0.0 ± 0.0
AS_IS_4FKLF	140	–	–	–
TO_BE_4AMR	140	52.1 ± 0.3	14.8 ± 0.3	8.0 ± 0.0
TO_BE_5AMR	140	59.3 ± 0.5	16.4 ± 0.4	4.8 ± 0.3
TO_BE_6AMR	140	80.5 ± 0.3	23.7 ± 0.6	0.0 ± 0.0
TO_BE_7AMR	140	88.6 ± 0.5	34.4 ± 0.5	0.0 ± 0.0
TO_BE_8AMR	140	95.3 ± 0.1	42.4 ± 0.5	0.0 ± 0.0

C.2 Block B: Auxiliary AMR Policy Evaluation

C.2.1 Replication-Level Completion Time Outputs

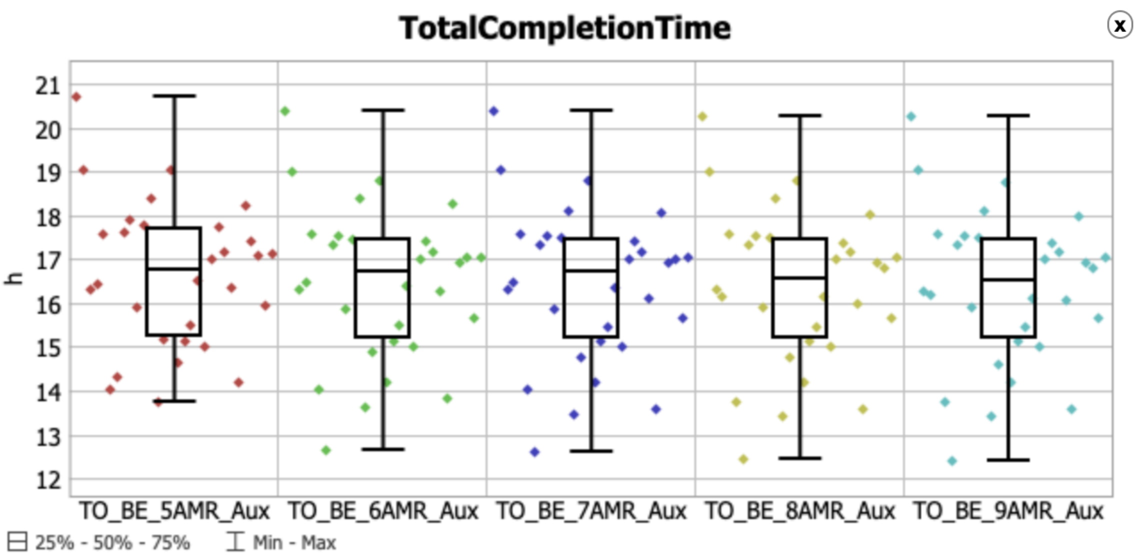

Figure C.4: Replication-level completion time results for Block B under the 100-truck workload.

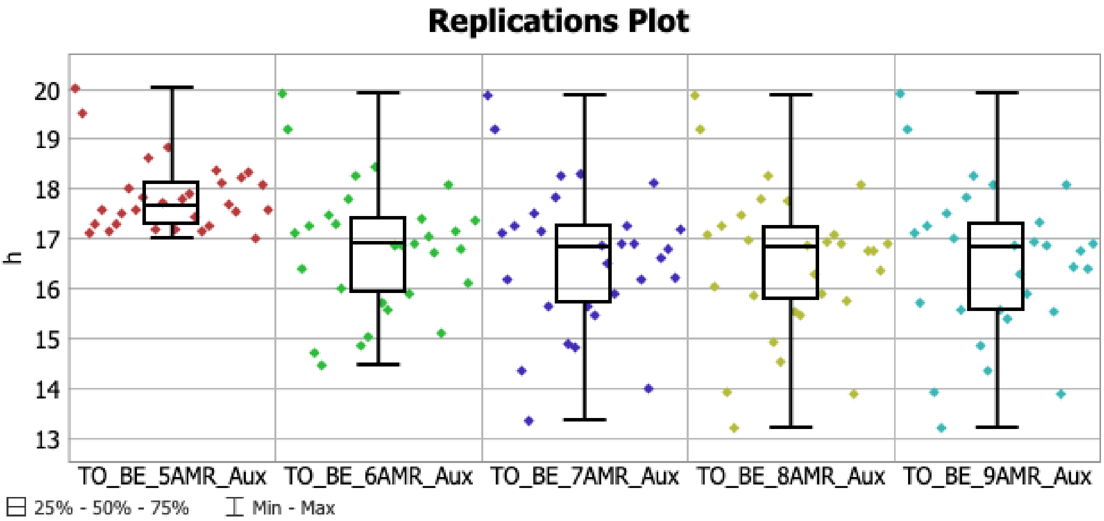

Figure C.5: Replication-level completion time results for Block B under the 120-truck workload.

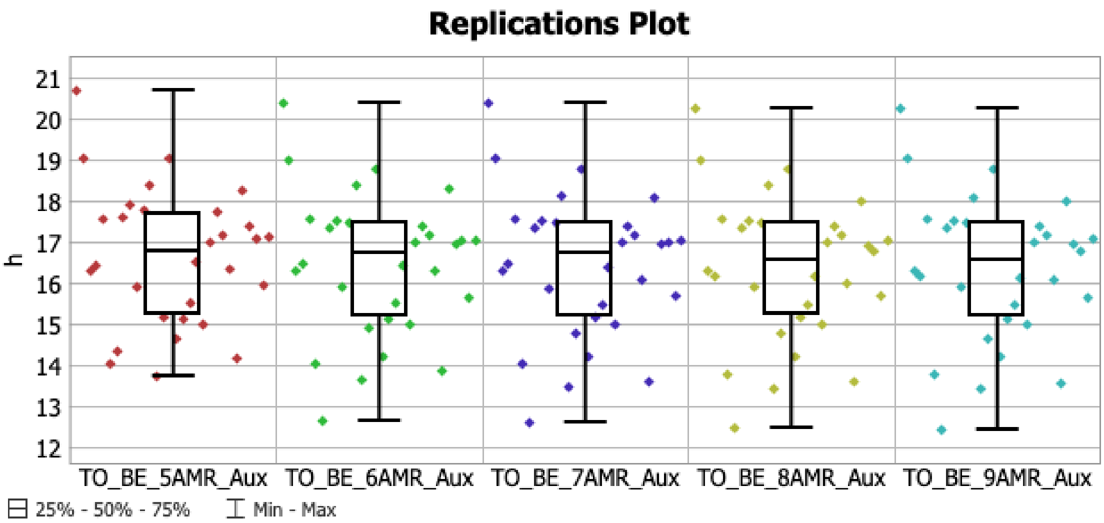

Figure C.6: Replication-level completion time results for Block B under the 140-truck workload.

C.2.2 Additional Truck-Flow Outputs

Table C.3: Additional truck-flow output for Block B, reported as mean ± 95% confidence interval across 30 replications.

Measure	W	AS-IS	5 AMR	6 AMR	7 AMR	8 AMR	9 AMR
	100	31.8	47.5 ± 4.1	39.0 ± 1.8	36.1 ± 1.2	33.6 ± 0.7	32.8 ± 0.6
Truck time in system (min)	120	33.1	72.1 ± 10.6	49.3 ± 5.2	41.3 ± 2.7	36.5 ± 1.5	34.7 ± 1.0
	140	37.0	145.1 ± 16.2	80.7 ± 11.9	59.1 ± 8.5	46.5 ± 4.9	40.3 ± 3.3

C.2.3 Detailed Loading-Time Outputs

Table C.4: Detailed loading-time outputs for Block B, reported as mean ± 95% confidence interval (minutes) across 30 replications. Shared-load participation is reported in per cent.

Scenario	W	Loading time	Pallet cycle	Non-rotated cycle	Rotated cycle	Shared loads (%)
TO_BE_5AMR_Aux	100	27.06 ± 0.59	1.33 ± 0.01	1.17 ± 0.00	1.98 ± 0.02	45.7
TO_BE_6AMR_Aux	100	23.53 ± 0.48	1.36 ± 0.01	1.19 ± 0.00	2.05 ± 0.02	70.1
TO_BE_7AMR_Aux	100	22.27 ± 0.30	1.37 ± 0.00	1.20 ± 0.00	2.07 ± 0.01	78.9
TO_BE_8AMR_Aux	100	20.57 ± 0.20	1.39 ± 0.00	1.21 ± 0.00	2.10 ± 0.01	90.8
TO_BE_9AMR_Aux	100	20.20 ± 0.11	1.39 ± 0.00	1.22 ± 0.00	2.10 ± 0.00	93.4
TO_BE_5AMR_Aux	120	29.02 ± 0.44	1.31 ± 0.00	1.16 ± 0.00	1.93 ± 0.02	31.7
TO_BE_6AMR_Aux	120	25.26 ± 0.46	1.34 ± 0.01	1.18 ± 0.00	2.01 ± 0.01	57.4
TO_BE_7AMR_Aux	120	23.32 ± 0.32	1.36 ± 0.00	1.20 ± 0.00	2.04 ± 0.01	71.1
TO_BE_8AMR_Aux	120	21.28 ± 0.23	1.38 ± 0.00	1.21 ± 0.00	2.07 ± 0.01	85.3
TO_BE_9AMR_Aux	120	20.54 ± 0.15	1.39 ± 0.00	1.22 ± 0.00	2.10 ± 0.01	90.3
TO_BE_5AMR_Aux	140	30.60 ± 0.15	1.30 ± 0.00	1.15 ± 0.00	1.90 ± 0.01	21.3
TO_BE_6AMR_Aux	140	27.00 ± 0.31	1.33 ± 0.00	1.17 ± 0.00	1.98 ± 0.01	45.9
TO_BE_7AMR_Aux	140	24.41 ± 0.28	1.36 ± 0.00	1.19 ± 0.00	2.02 ± 0.01	63.9
TO_BE_8AMR_Aux	140	22.25 ± 0.22	1.38 ± 0.00	1.20 ± 0.00	2.06 ± 0.01	78.9
TO_BE_9AMR_Aux	140	20.95 ± 0.15	1.39 ± 0.00	1.21 ± 0.00	2.08 ± 0.01	88.0

C.2.4 Loading-Time Reductions Relative to Block A

Table C.5: Loading-time reductions of Block B (auxiliary AMR policy) relative to the equivalent Block A configuration without the policy, at the same workload and fleet size.

Fleet	W	Block A (min)	Block B (min)	Reduction (%)
5 AMR	100	33.66	27.06	19.6
6 AMR	100	33.64	23.53	30.1
7 AMR	100	33.65	22.27	33.8
8 AMR	100	33.65	20.57	38.9
9 AMR	100	–	20.20	–[a]
5 AMR	120	33.51	29.02	13.4
6 AMR	120	33.50	25.26	24.6
7 AMR	120	33.50	23.32	30.4
8 AMR	120	33.50	21.28	36.5
9 AMR	120	–	20.54	–[a]
5 AMR	140	33.66	30.60	9.1
6 AMR	140	33.65	27.00	19.7
7 AMR	140	33.65	24.41	27.5
8 AMR	140	33.65	22.25	33.9
9 AMR	140	–	20.95	–[a]

[a]No 9-AMR configuration was simulated in Block A. Since Block A loading time is essentially constant across fleet sizes (≈33.6 min), the 9-AMR auxiliary configuration corresponds to a reduction of approximately 40% (W100), 39% (W120) and 38% (W140) relative to that constant baseline.

C.2.5 Battery and Charging Outputs

Table C.6: Battery and charging outputs for Block B, reported as mean ± 95% confidence interval across 30 replications.

Scenario	W	Avg battery level (%)	Avg charging state (%)	Mandatory charging events
TO_BE_5AMR_Aux	100	75.0 ± 4.1	23.1 ± 2.4	0.3 ± 0.4
TO_BE_6AMR_Aux	100	87.2 ± 2.3	32.6 ± 2.6	0.0 ± 0.0
TO_BE_7AMR_Aux	100	93.5 ± 1.2	41.2 ± 2.4	0.0 ± 0.0
TO_BE_8AMR_Aux	100	95.7 ± 0.5	47.6 ± 2.2	0.0 ± 0.0
TO_BE_9AMR_Aux	100	96.7 ± 0.3	52.9 ± 2.0	0.0 ± 0.0
TO_BE_5AMR_Aux	120	61.6 ± 2.3	16.4 ± 0.9	2.5 ± 0.8
TO_BE_6AMR_Aux	120	75.2 ± 3.3	22.8 ± 2.0	0.0 ± 0.1
TO_BE_7AMR_Aux	120	86.4 ± 2.1	31.5 ± 2.2	0.0 ± 0.0
TO_BE_8AMR_Aux	120	92.3 ± 1.3	38.5 ± 2.1	0.0 ± 0.0
TO_BE_9AMR_Aux	120	95.3 ± 0.8	44.4 ± 2.0	0.0 ± 0.0
TO_BE_5AMR_Aux	140	54.3 ± 0.8	13.8 ± 0.4	5.1 ± 0.1
TO_BE_6AMR_Aux	140	62.7 ± 2.3	16.0 ± 0.9	1.6 ± 0.7
TO_BE_7AMR_Aux	140	77.0 ± 2.3	22.1 ± 1.7	0.0 ± 0.1
TO_BE_8AMR_Aux	140	86.1 ± 1.7	29.3 ± 2.0	0.0 ± 0.0
TO_BE_9AMR_Aux	140	91.2 ± 1.6	35.8 ± 2.0	0.0 ± 0.0

C.3 Block C: Sensitivity to Non-Stationary Arrivals

C.3.1 Loading-Time and Shared-Load Participation Outputs

Table C.7: Detailed loading-time outputs for Block C, reported as mean ± 95% confidence interval (minutes) across 30 replications. Shared-load participation is reported in per cent.

Scenario	W	Loading time	Non-rotated cycle	Rotated cycle	Shared loads (%)
TO_BE_6AMR_Aux	100	24.52 ± 0.47	1.19 ± 0.00	2.04 ± 0.01	63.1
TO_BE_7AMR_Aux	100	22.76 ± 0.35	1.20 ± 0.00	2.07 ± 0.01	75.4
TO_BE_8AMR_Aux	100	20.94 ± 0.24	1.21 ± 0.00	2.09 ± 0.01	88.2
TO_BE_9AMR_Aux	100	20.44 ± 0.14	1.22 ± 0.00	2.10 ± 0.01	91.7
TO_BE_6AMR_Aux	120	25.59 ± 0.36	1.18 ± 0.00	2.00 ± 0.01	55.3
TO_BE_7AMR_Aux	120	23.41 ± 0.31	1.20 ± 0.00	2.03 ± 0.01	70.4
TO_BE_8AMR_Aux	120	21.62 ± 0.27	1.21 ± 0.00	2.06 ± 0.01	83.0
TO_BE_9AMR_Aux	120	20.65 ± 0.17	1.22 ± 0.00	2.09 ± 0.01	89.6
TO_BE_6AMR_Aux	140	26.90 ± 0.30	1.17 ± 0.00	1.99 ± 0.01	46.7
TO_BE_7AMR_Aux	140	24.60 ± 0.25	1.19 ± 0.00	2.03 ± 0.01	62.6
TO_BE_8AMR_Aux	140	22.40 ± 0.21	1.20 ± 0.00	2.06 ± 0.01	78.1
TO_BE_9AMR_Aux	140	21.17 ± 0.16	1.21 ± 0.00	2.09 ± 0.01	86.4

C.3.2 Additional Truck-Flow Outputs

Table C.8: Additional truck-flow output for Block C, reported as mean ± 95% confidence interval across 30 replications.

Measure	W	AS-IS	6 AMR	7 AMR	8 AMR	9 AMR
Truck time in system (min)	100	33.2 ± 1.1	45.0 ± 3.3	40.0 ± 2.6	36.2 ± 1.8	34.7 ± 1.5
	120	36.7 ± 1.9	60.4 ± 6.5	49.7 ± 4.3	43.2 ± 3.1	39.7 ± 2.5
	140	45.1 ± 3.7	100.5 ± 12.0	75.7 ± 9.4	60.0 ± 6.6	51.6 ± 4.9

C.3.3 Battery and Charging Outputs

Table C.9: Battery and charging outputs for Block C, reported as mean ± 95% confidence interval across 30 replications.

Scenario	W	Avg battery level (%)	Avg charging state (%)	Mandatory charging events
TO_BE_6AMR_Aux	100	86.3 ± 2.1	34.3 ± 2.6	0.0 ± 0.0
TO_BE_7AMR_Aux	100	92.1 ± 1.2	42.3 ± 2.4	0.0 ± 0.0
TO_BE_8AMR_Aux	100	94.7 ± 0.8	48.4 ± 2.3	0.0 ± 0.0
TO_BE_9AMR_Aux	100	96.1 ± 0.4	53.6 ± 2.1	0.0 ± 0.0
TO_BE_6AMR_Aux	120	75.1 ± 3.5	24.5 ± 2.2	0.1 ± 0.1
TO_BE_7AMR_Aux	120	85.7 ± 2.0	33.0 ± 2.2	0.0 ± 0.0
TO_BE_8AMR_Aux	120	91.4 ± 1.3	39.6 ± 2.2	0.0 ± 0.0
TO_BE_9AMR_Aux	120	94.0 ± 1.1	45.3 ± 2.0	0.0 ± 0.0
TO_BE_6AMR_Aux	140	62.3 ± 2.4	17.4 ± 1.1	0.9 ± 0.5
TO_BE_7AMR_Aux	140	77.2 ± 2.1	23.8 ± 1.7	0.0 ± 0.0
TO_BE_8AMR_Aux	140	84.7 ± 1.7	30.5 ± 2.0	0.0 ± 0.0
TO_BE_9AMR_Aux	140	89.9 ± 1.5	36.6 ± 2.0	0.0 ± 0.0

C.3.4 Replication-Level Completion Time Outputs

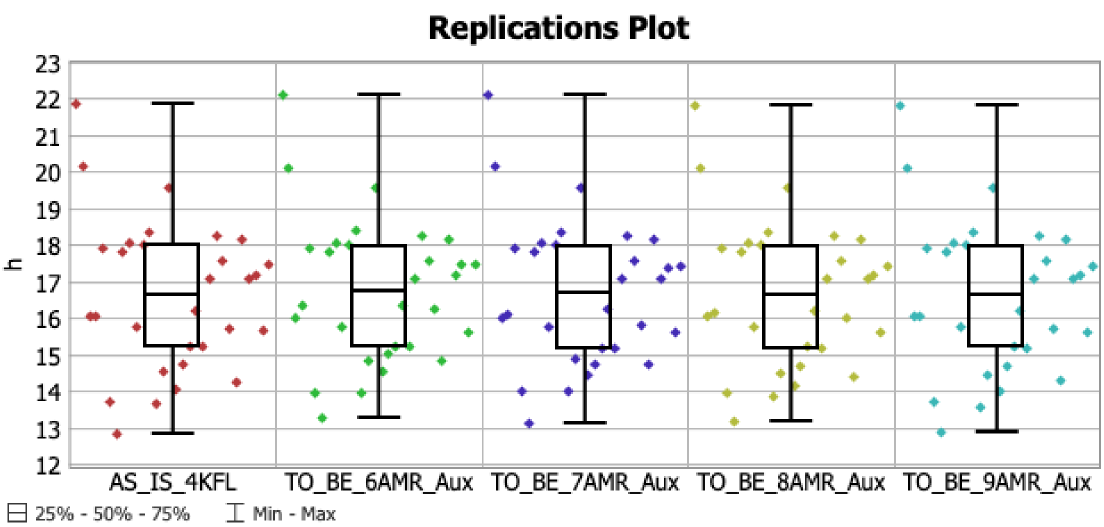

Figure C.7: Replication-level completion time results for Block C under the 100-truck workload.

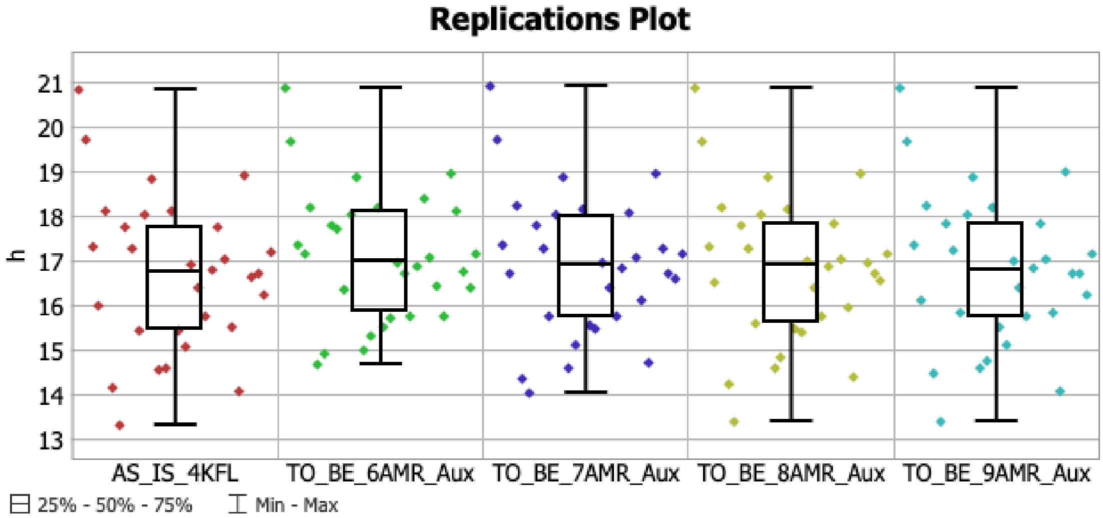

Figure C.8: Replication-level completion time results for Block C under the 120-truck workload.

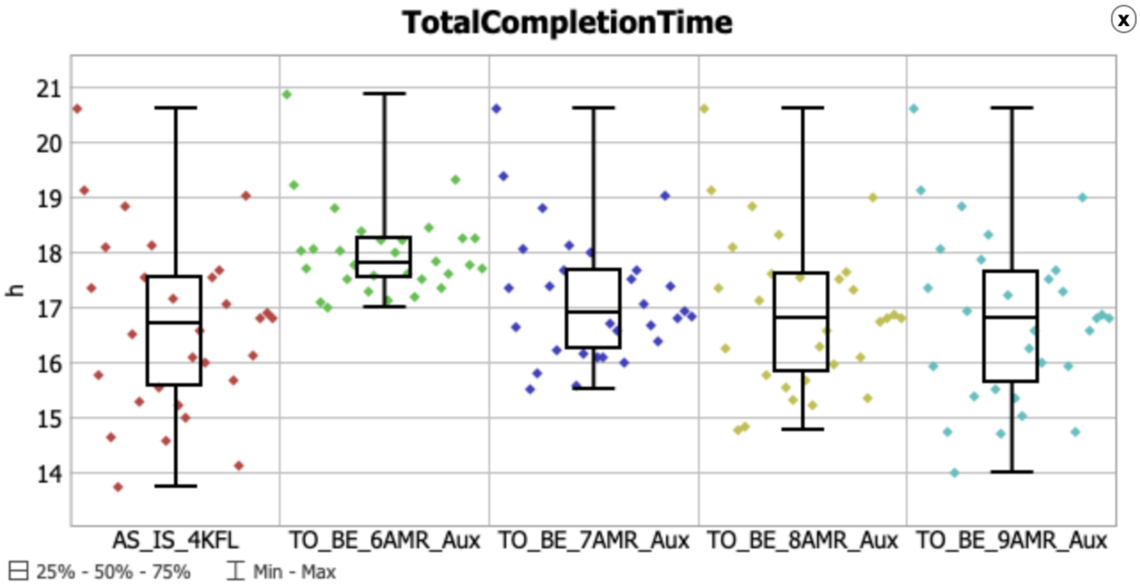

Figure C.9: Replication-level completion time results for Block C under the 140-truck workload.

C.4 Block D: Sensitivity to Rotation-Heavy Load Composition

C.4.1 Loading-Time and Shared-Load Participation Outputs

Table C.10: Loading-time and shared-load participation outputs for Block D, reported as mean ± 95% confidence interval (minutes) across 30 replications. Shared-load participation is reported in per cent.

Scenario	W	Loading time	Non-rotated cycle	Rotated cycle	Shared loads (%)
TO_BE_7AMR_Aux	120	31.54 ± 0.31	1.32 ± 0.01	2.01 ± 0.01	59.3
TO_BE_8AMR_Aux	120	28.74 ± 0.30	1.38 ± 0.01	2.06 ± 0.01	76.9
TO_BE_9AMR_Aux	120	27.34 ± 0.20	1.41 ± 0.01	2.08 ± 0.00	85.7
TO_BE_10AMR_Aux	120	25.77 ± 0.14	1.47 ± 0.00	2.11 ± 0.00	97.5
TO_BE_11AMR_Aux	120	26.03 ± 0.11	1.46 ± 0.00	2.11 ± 0.00	96.0
TO_BE_7AMR_Aux	140	32.19 ± 0.20	1.31 ± 0.01	2.00 ± 0.01	54.9
TO_BE_8AMR_Aux	140	29.55 ± 0.15	1.36 ± 0.01	2.04 ± 0.00	71.7
TO_BE_9AMR_Aux	140	27.70 ± 0.16	1.40 ± 0.01	2.07 ± 0.00	83.1
TO_BE_10AMR_Aux	140	26.12 ± 0.20	1.43 ± 0.01	2.10 ± 0.00	93.1
TO_BE_11AMR_Aux	140	26.17 ± 0.09	1.44 ± 0.00	2.10 ± 0.00	92.9

C.4.2 Battery and Charging Outputs

Table C.11: Battery and charging outputs for Block D, reported as mean ± 95% confidence interval across 30 replications.

Scenario	W	Avg battery level (%)	Avg charging state (%)	Mandatory charging events
TO_BE_7AMR_Aux	120	69.3 ± 3.0	17.9 ± 1.6	0.2 ± 0.2
TO_BE_8AMR_Aux	120	80.2 ± 2.2	24.0 ± 1.9	0.0 ± 0.0
TO_BE_9AMR_Aux	120	87.0 ± 1.7	30.3 ± 2.2	0.0 ± 0.0
TO_BE_10AMR_Aux	120	95.3 ± 0.7	45.1 ± 2.3	0.0 ± 0.0
TO_BE_11AMR_Aux	120	96.6 ± 0.3	50.0 ± 2.1	0.0 ± 0.0
TO_BE_7AMR_Aux	140	60.6 ± 1.4	14.2 ± 0.5	1.2 ± 0.4
TO_BE_8AMR_Aux	140	71.3 ± 1.8	17.5 ± 1.1	0.0 ± 0.0
TO_BE_9AMR_Aux	140	79.9 ± 1.4	22.6 ± 1.5	0.0 ± 0.0
TO_BE_10AMR_Aux	140	84.9 ± 1.4	27.6 ± 1.7	0.0 ± 0.0
TO_BE_11AMR_Aux	140	91.2 ± 1.4	33.7 ± 1.6	0.0 ± 0.0

C.4.3 Replication-Level Completion Time Outputs

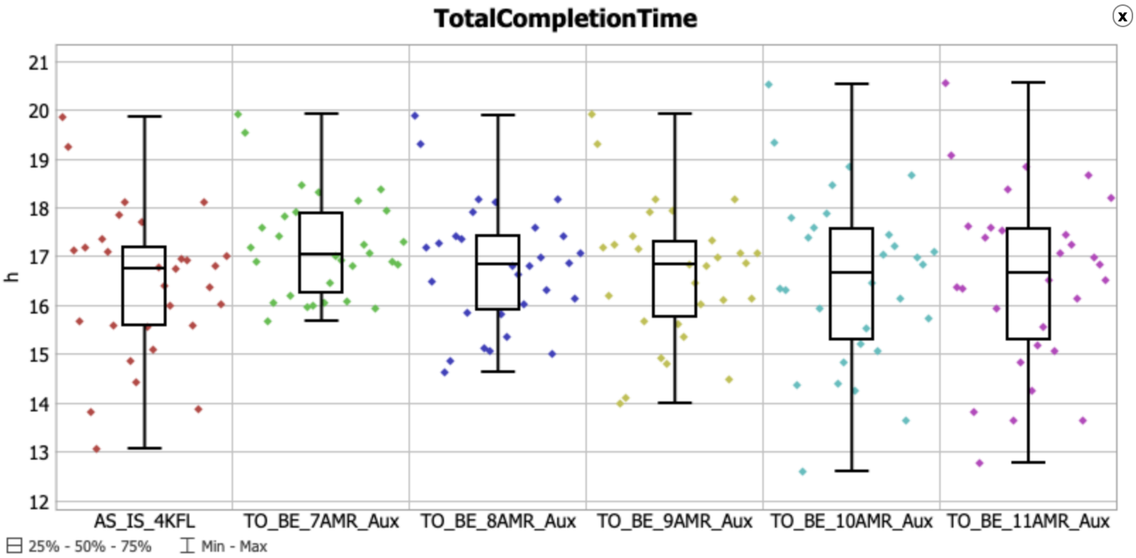

Figure C.10: Replication-level completion time results for Block D under the 120-truck rotation-heavy workload.

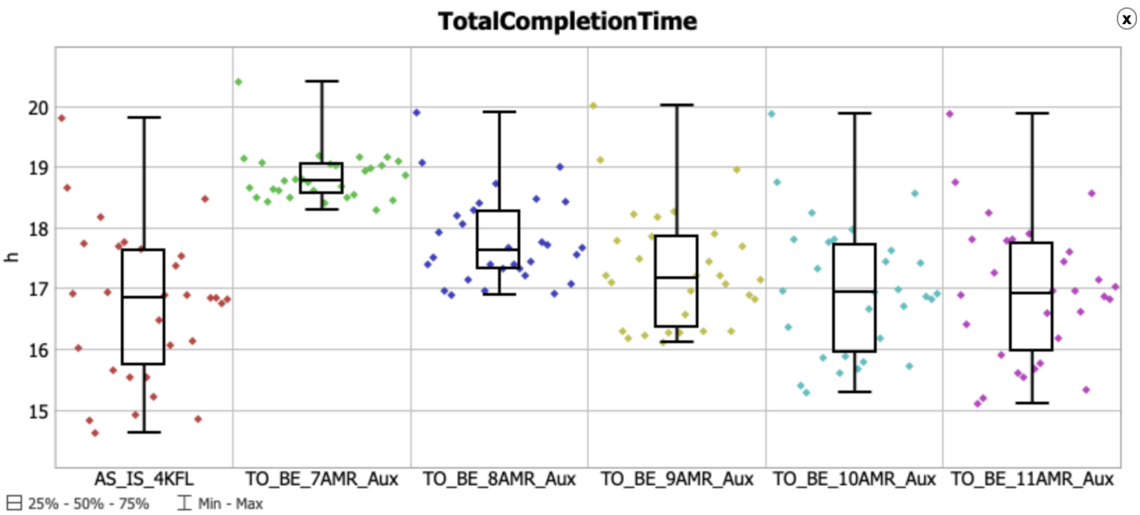

Figure C.11: Replication-level completion time results for Block D under the 140-truck rotation-heavy workload.

C.5 Block E: High-Demand Sensitivity Scenario

C.5.1 Replication-Level Completion Time Outputs

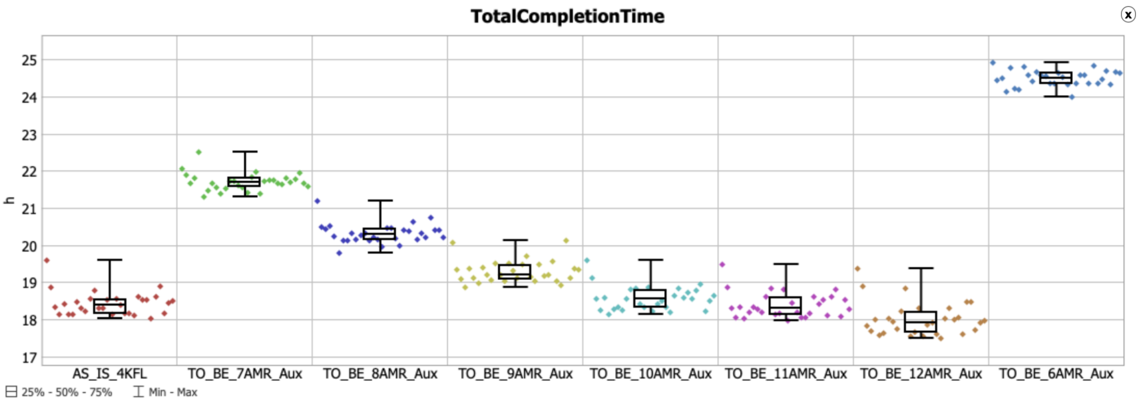

Figure C.12: Replication-level completion time results for High Demand Sensitivity Scenario under the 200-truck workload

C.5.2 AMR Per-State Utilisation Outputs

Table C.12: High-demand scenario (W200) – AMR fleet utilisation across 30 replications.

Fleet	Idle + charging (%)[a]	Working state (%)	Min / Max per AMR (%)[b]	AMRs below 20%[c]
6 AMR	19.07	80.93	18.61/19.36	6/6
7 AMR	21.14	78.86	20.64/21.61	0/7
8 AMR	25.73	74.27	25.01/26.76	0/8

[a]Mean combined idle-and-charging share across all AMRs and replications.

[b]Lowest and highest per-AMR means (averaged over replications), indicating workload evenness.

[c]Number of AMRs whose mean idle-and-charging share falls below the 20% threshold.

Table C.13: High-demand scenario (W200) - Battery and charging outputs, reported as mean ± 95% confidence interval across 30 replications.

Measure	6 AMR	7 AMR	8 AMR	9 AMR	10 AMR	11 AMR	12 AMR
Avg battery level (%)	57.6 ± 0.4	65.9 ± 1.0	76.5 ± 0.5	81.2 ± 0.3	84.4 ± 0.6	91.6 ± 0.6	96.9 ± 0.1
Avg charging state (%)	15.1 ± 0.2	15.4 ± 0.3	19.1 ± 0.3	23.2 ± 0.4	27.2 ± 0.5	32.5 ± 0.5	36.1 ± 0.6
Mandatory charging events	5.90 ± 0.15	0.47 ± 0.23	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00

C.6 Stress-Test Scenario

C.6.1 Battery and Charging Outputs

Table C.14: Battery and charging outputs for the stress-test scenario, reported as mean ± 95% confidence interval across 30 replications.

Scenario	Avg battery level (%)	Avg charging state (%)	Mandatory charging events
TO_BE_8AMR_Aux	70.3 ± 1.4	14.0 ± 0.3	0.0 ± 0.0
TO_BE_9AMR_Aux	75.4 ± 0.8	16.6 ± 0.5	0.0 ± 0.0
TO_BE_10AMR_Aux	78.3 ± 0.7	19.4 ± 0.7	0.0 ± 0.0
TO_BE_11AMR_Aux	84.3 ± 0.3	25.8 ± 0.2	0.0 ± 0.0
TO_BE_12AMR_Aux	89.2 ± 0.2	28.1 ± 0.1	0.0 ± 0.0

C.6.2 Replication-Level Completion Time Outputs

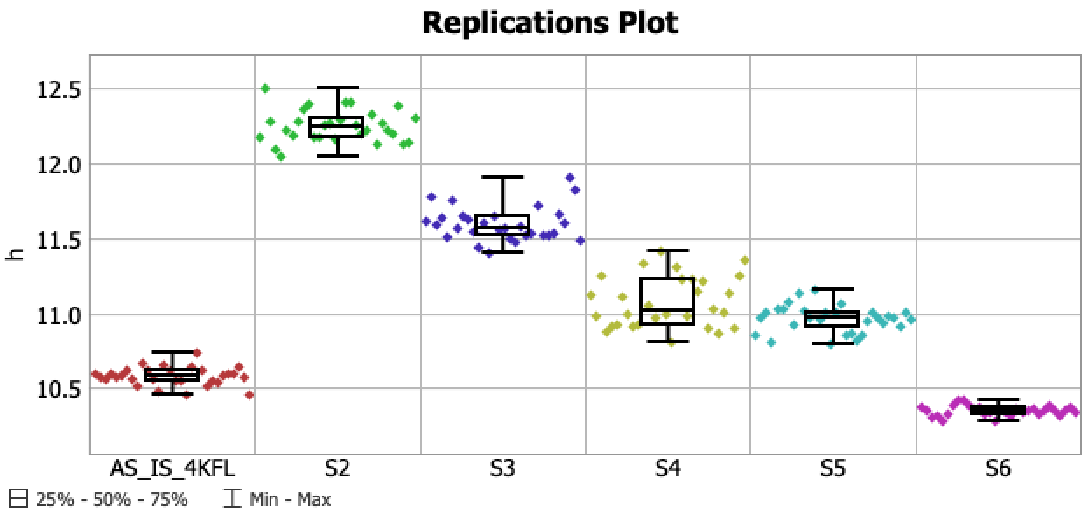

Figure C.13: Replication-level completion time results for the stress-test scenario.