Agent-Based Simulation and Digital Twin Modeling for Dynamic Storage Strategies in Dark Warehouses

Simon Holzhauer and Philipp Nenninger
Institute of Energy Efficient Mobility
Karlsruhe University of Applied Sciences
Karlsruhe, Germany
e-mail: {simon.holzhauer | philipp.nenninger}@h-ka.de

***Abstract*—The increasing agility and dynamics of global e-commerce, coupled with steadily rising operational costs and labor shortages, necessitate the transition toward fully automated storage facilities known as dark warehouses. Within these environments, optimizing the Storage Location Assignment Problem (SLAP) is critical to maintaining high-throughput logistics. However, evaluating complex, dynamic storage strategies in active facilities is computationally demanding, requiring robust predictive models rather than static equations. To address this challenge, this paper presents an agent-based discrete-event simulation environment developed as the computational core of an industrial digital twin. Implemented in Python, the model accurately maps the physical kinematics and asymmetric interactions between Automated Guided Vehicles (AGVs) and Automated Storage and Retrieval Systems (AS/RS) within a pallet high-bay warehouse. The simulation systematically evaluates four storage strategies: Random Fixed, Shared Random, Access-frequency-based Fixed Location (FSN), and a Dynamic Predictive Reslotting strategy under varying stochastic load scenarios and sudden seasonal demand shifts. Quantitative analysis reveals a non-linear relationship between hardware resource utilization and systemic performance. The study demonstrates that while static models severely degrade during demand volatility, the dynamic, predictive reslotting algorithm is superior. By leveraging idle time to proactively reorganize inventory, the dynamic strategy minimizes AGV traffic congestion, significantly reduces average retrieval times and enhances overall throughput predictability.**

Keywords-Digital Twins; Agent-Based Simulation; Dark Warehouse; Storage Location Assignment Problem; Intralogistics; Automated Guided Vehicles.

I. INTRODUCTION

The contemporary supply chain is under unprecedented pressure, characterized by the demand for maximum agility and the mitigation of steadily rising costs for storage space, energy, and human capital. This operational pressure is further amplified by a growing global shortage of skilled labor, forcing industrial enterprises to design resilient, highly efficient, and fail-safe automated processes. In this context, *dark warehouses*, which are fully automated storage facilities operating without direct human intervention or environmental conditioning, have emerged as a critical evolutionary step in intralogistics [1][2]. Understanding the economic drivers behind this paradigm shift requires a closer look at the traditional cost structures governing modern warehousing. Evaluating these baseline expenditures helps contextualize why rigid, legacy control frameworks are increasingly being replaced by cyber-physical optimization strategies.

Warehousing and inventory management represent a substantial financial burden within global supply chains, conventionally accounting for roughly 20% to 30% of total logistics costs [3]. Within these facilities, the physical retrieval and order-picking processes are highly labor- and capital-intensive, historically consuming an estimated 55% to 65% of the total warehouse operating budget [4][5]. The optimization of storage and retrieval processes represents a pivotal lever for reducing overall supply chain expenditure. By deploying automated hardware alongside intelligent, data-driven storage strategies, dark warehouses mitigate workforce shortages while simultaneously enhancing throughput velocity, operational flexibility, and order fulfillment precision [6][7].

Evaluating the performance of these complex systems under fully automated conditions requires sophisticated simulative approaches. While industrial software exists for warehouse design, it is frequently proprietary, vendor-locked, and unsuitable for independent, strategic algorithmic evaluation [8][9]. Additionally, scientific literature often isolates specific subproblems of warehousing, failing to capture the systemic interdependencies, race conditions, and spatial congestion inherent in a fully automated facility [4][7].

To address this gap, this investigation focuses on the development of an agent-based simulation environment that acts as the analytical engine for an industrial digital twin [10]–[12]. This framework enables the holistic, quantitative analysis of various storage strategies, spanning static fixed-location models to dynamic, predictive reslotting algorithms. By subjecting these cyber-physical models to varying stochastic demand scenarios, including severe utilization loads and sudden seasonal structure shifts, the research quantifies the limits of hardware scaling versus software orchestration.

The paper is organized as follows: Section II discusses related work surrounding digital twins and simulation methodologies. Section III details the methodology, specifically the storage location assignment problem formulation and the cyber-physical environment architecture. Section IV outlines the mathematical kinematic equations governing the agent movement profiles. Section V presents the algorithmic frameworks for the evaluated storage strategies. Section VI describes the experimental setup and scenario generation. Section VII presents the quantitative results of the simulation. Section VIII provides a comprehensive discussion of these findings. Finally, the conclusion and future work are given in Section X.

 ISBN: 978-1-68558-400-9

II. Related Work in Industrial Twins and Automated Warehousing

The concept of the digital twin has evolved from passive geometric modeling into a predictive, analytical engine essential for complex manufacturing and logistics networks [13]. This evolution is formalized by the five-dimensional digital twin model, comprising the physical entity, virtual entity, services, data, and connections [12]. In dark warehouses, the virtual entity must accurately mirror physical kinematics and spatial constraints; however, the ultimate value lies in the "services" dimension: utilizing the digital replica to run predictive simulations and proactively optimize system parameters before physical execution [12].

The current technological frontier integrates DTs with Agent-Based Simulation (ABS) to create autonomous, self-optimizing entities [14]. Unlike top-down system dynamics or rigid Discrete-Event Simulation (DES), ABS models systems from the bottom up. By defining independent decision-making agents, such as AGVs and AS/RS, ABS captures emergent system dynamics, spatial congestion, and routing conflicts that linear models systematically fail to represent [15][16]. When embedded within a digital twin architecture, ABS acts as the critical bridge between theoretical Warehouse Management System (WMS) optimization and real-time, collision-free execution in automated environments [17][18].

III. Formulating the Storage Location Assignment Problem

The core operational challenge addressed by the simulated digital twin is the Storage Location Assignment Problem (SLAP). SLAP is formally defined as the optimal mapping of a set of products, or Stock Keeping Units (SKUs), to a set of available storage locations to maximize warehousing efficiency [19][20]. Efficiency in this context encompasses the minimization of transport distances, the reduction of retrieval times, and the optimal utilization of spatial capacity.

In its most fundamental form, often referred to as the *Greenfield* scenario where the warehouse is assumed to be entirely empty, SLAP can be mathematically formulated as a classical binary integer programming problem. The theoretical objective function Z aims to minimize the expected retrieval costs based purely on historical access frequencies and spatial distances [21]. This objective function is expressed as:

$$\min Z = \sum_{i\in I}\sum_{j\in J} f_i \cdot d_j \cdot y_{ij} \tag{1}$$

In this formulation, i serves as the index for a product from the set of all products I, while j is the index for a storage location from the set of all locations J. The parameter f_i denotes the access frequency of product i, and d_j represents the physical distance of storage location j to the target input/output (I/O) point. The binary decision variable y_{ij} indicates whether product i is assigned to location j (taking a value of 1) or not (taking a value of 0).

This optimal assignment is subject to strict capacity constraints, ensuring that each product is assigned to exactly one location ($\sum_{j\in J} y_{ij} = 1$) and that each storage location holds no more than one product ($\sum_{i\in I} y_{ij} \leq 1$) [20]. If the minimum of Z is achieved under these conditions, the storage allocation is considered optimal for a static environment.

However, the Greenfield assumption is rarely applicable in continuous dark warehouse operations. A functional warehouse represents a *Brownfield* scenario, containing existing, dynamically fluctuating stock. Any algorithmic attempt to optimize the layout during active operations (reslotting) must mathematically account for the immediate operational penalty of moving an item. Therefore, an extended, practice-oriented objective function must be formulated. This function weighs the long-term retrieval efficiency against the short-term kinematic effort required for relocation:

$$\min Z = \sum_{i\in I}\sum_{j\in J} c_{ij} \cdot x_{ij} + \sum_{i\in I}\sum_{j\in J} u_{ij} \cdot x_{ij} \tag{2}$$

In (2), the binary decision variable x_{ij} dictates whether product i is assigned to or relocated to location j during the active operational phase. Here, c_{ij} represents the expected retrieval cost, analogous to $f_i \cdot d_j$ from the Greenfield model. The critical addition is u_{ij}, which represents the precise cost of physically relocating product i from its current position to the new optimal position j [22].

The inclusion of these relocation costs, combined with the scale of modern distribution centers, transforms the SLAP into a highly complex, NP-hard problem [19]. Because exact mathematical optimization models, such as Mixed-Integer Linear Programming (MILP), suffer from exponential calculation times and are unsuitable for real-time digital twin synchronization, the framework relies on highly efficient heuristic approaches. Specifically, greedy algorithms and metaheuristics are utilized, making locally optimal choices at discrete intervals to approximate a global optimum without violating the real-time computational constraints of the simulation engine [19].

These computational bottlenecks and the inherent spatial complexities of the SLAP can be effectively overcome by embedding this heuristic logic within a dynamic, agent-based simulation. By acting as the virtual entity of the digital twin, this environment evaluates the storage algorithms not in a static mathematical way, but against the emergent traffic and kinematic constraints of a real system.

IV. Architecting the Digital Twin and Cyber-Physical Environment

To ensure that this simulative evaluation maintains empirical validity and industrial applicability, the digital twin must be constructed based on exact physical parameters representative of modern industrial pallet high-bay warehouses. The architecture establishes strict rules regarding spatial topology, capacity, and the interaction interfaces between different classes of autonomous agents.

The reference layout simulates a freestanding high-bay warehouse. The structural geometry consists of 12 distinct racks, each containing 200 compartments arranged in a grid

 ISBN: 978-1-68558-400-9

of 20 vertical columns and 10 horizontal levels, yielding a total systemic capacity of 2,400 storage locations. The storage methodology is strictly single-depth. This architectural decision ensures direct, unhindered access to every individual pallet, allowing the simulation to enforce strict first-in-first-out (FIFO) retrieval principles, which are often compromised in double-deep or block storage configurations. The simulation assumes standard European pallets, leading to a parameterized compartment width of 1000 mm, depth of 1500 mm, and a height of 2000 mm to accommodate both the pallet and the stored goods [23].

The spatial dimension of the digital twin is modeled using the Python library NetworkX, which maps the physical layout into a rigid, undirected graph structure as seen in Figure 1[24]. Within this graph, nodes represent physical positions, encompassing intersections, I/O points, and racking interfaces, while weighted edges represent navigable paths based on exact physical distances. This graph serves as the absolute coordinate system for all agent routing.

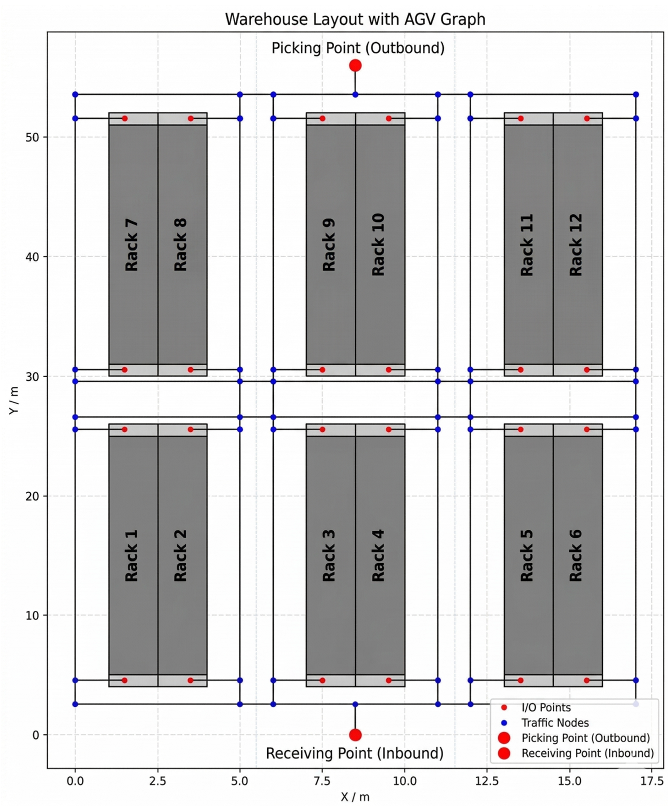

Figure 1. Warehouse Layout with Graph

To simulate the traffic constraints and emergent congestion characteristic of dark warehouses, strict capacity limits are enforced on the graph nodes via the SimPy discrete-event framework. The central goods receipt and goods issue points act as the simulation boundaries and possess high-capacity buffers to accommodate simultaneous staging by multiple agents. The nodes representing the physical handover points between AGVs and SRUs are strictly limited to a capacity of 1, as mechanical constraints dictate that an SRU can only interface with a single AGV at any given moment. Furthermore, intersection and transit routing nodes are limited to a capacity of 2. This prevents deadlocks while enabling fluid bidirectional traffic within the aisles.

Pathfinding within this constrained graph relies on the A^* search algorithm, which guarantees shortest-path optimality across weighted networks [25]. Because the central WMS logic generates thousands of routing requests per hour, the A^* algorithm is optimized via caching mechanisms. Calculated paths between static nodes are retained in memory, drastically reducing computational overhead and allowing the simulation to process extended time horizons rapidly.

V. Kinematic Equations and Agent Movement Profiles

A critical component of a detailed digital twin is the precise mathematical representation of cyber-physical kinematics [26]. The simulation advances beyond static topological routing by calculating the exact temporal duration of every mechanical action, utilizing distinct velocity profiles that account for the acceleration and deceleration phases of AS/RS.

A. AS/RS Kinematic Modeling

An AS/RS operates on a two-dimensional vertical plane within the storage aisles. Because the horizontal drive motor along the X-axis and the hoisting mechanism along the Z-axis operate concurrently and independently, the total travel time T_{travel} is dictated by the slower of the two axes. This relationship is mathematically represented by applying the Chebyshev distance metric to the temporal domain [27]:

$$T_{total} = \max(T_x, T_z) + T_{cycle} \tag{3}$$

In (3), T_{cycle} represents the constant duration for the mechanical fork cycle, encompassing the time required to physically extract or deposit a pallet. The simulation defines this interaction time as 5.0 seconds for pickup and 5.0 seconds for drop-off, yielding a 10-second T_{cycle}.

The calculation for the travel time of an individual axis (e.g., T_x) must differentiate between long and short travel distances. For long distances, the machine reaches its maximum velocity, forming a trapezoidal velocity profile (encompassing acceleration, constant velocity, and deceleration phases). For short distances, the deceleration phase must commence before maximum velocity is achieved, forming a triangular profile. The critical distance d_{crit} required to reach the maximum velocity v_{max} and immediately decelerate to zero, assuming a constant acceleration and deceleration rate a, is calculated as:

$$d_{crit} = \frac{v_{max}^2}{a} \tag{4}$$

Consequently, the precise time T_x required to cover a horizontal distance Δx is computed as a piecewise kinematic function [26]:

$$T_x = \begin{cases} \frac{v_{max,x}}{a_x} + \frac{\Delta x}{v_{max,x}} & \text{if } \Delta x \geq \frac{v_{max,x}^2}{a_x} \text{ (Trapezoidal)} \\ 2 \cdot \sqrt{\frac{\Delta x}{a_x}} & \text{if } \Delta x < \frac{v_{max,x}^2}{a_x} \text{ (Triangular)} \end{cases} \tag{5}$$

The calculation for the vertical axis time T_z follows identical logic, substituting the respective vertical velocity and acceleration parameters. In the simulation, the AS/RS parameters are calibrated to match modern industrial performance standards: the horizontal maximum velocity is configured as $v_{max,x} = 3.0$ m/s with an acceleration of $a_x = 1.0$ m/s^2, while the vertical hoisting velocity is set to $v_{max,z} = 1.5$ m/s with an acceleration of $a_z = 1.0$ m/s^2.

B. Automated Guided Vehicle Kinematics

In contrast to the continuous multi-axis movement of AS/RS, AGVs navigate sequentially along an orthogonal grid on the horizontal plane. Their theoretical travel distances conform strictly to Manhattan distance geometry. To maintain detailed physical representation, the simulation models AGVs with strict kinematic constraints: they must decelerate to a complete stop to execute 90-degree rotations at grid intersections.

The total travel time T_{travel} for an AGV is therefore the sum of all linear transit segments and rotational delays:

$$T_{travel} = \sum_{i=1}^{N} T_{seg,i} + \sum_{j=1}^{M} T_{turn,j} \tag{6}$$

In (6), N represents the total number of straight segments, M denotes the number of direction changes (where $M = N - 1$), $T_{seg,i}$ is the kinematic time required to traverse the i-th segment, and $T_{turn,j}$ applies a constant rotational penalty of 2.0 seconds per 90-degree turn.

The linear traversal time $T_{seg,i}$ is computed using the identical piecewise trapezoidal and triangular kinematic logic established for the AS/RS. To reflect standard industrial safety and performance limits, the baseline AGV kinematic parameters are defined with a maximum velocity of $v_{max} = 1.5$ m/s (equivalent to 5.4 km/h) and a constant acceleration and deceleration rate of $a = 1.0$ m/s^2, which strictly complies with the DIN EN ISO 3691-4 specifications [28].

C. Kinematic Validation

To ensure the integrity of the digital twin, the discrete-event simulation engine was validated against manual kinematic calculations. An edge-case scenario was tested: the retrieval of a pallet located in the compartment closest to the output node (rack 10, compartment 20/1). The manual calculation aggregated the triangular velocity profiles of the short SRU movement, the SRU interaction times, the piecewise AGV segment times across an 8.5-meter multi-segment path, and six total 90-degree AGV rotations. The theoretical manual calculation yielded a total execution time of 54.24 seconds. The Python-based digital twin simulated the identical event in 54.22 seconds, confirming a near-perfect mathematical alignment and validating the engine for large-scale scenario testing.

VI. Storage Strategies and Dynamic Reslotting Logic

The validated simulation serves as the computational testbed for evaluating four distinct storage assignment policies. These strategies advance in complexity from static, rule-based allocations to highly dynamic, predictive algorithms, allowing for a comprehensive benchmarking of intralogistics optimization techniques [19].

The first strategy evaluated is *Dedicated Random Storage*. This baseline strategy mirrors the most rudimentary form of warehouse management. Upon initial system setup, each SKU is assigned a permanent, randomly selected storage compartment. This assignment remains fixed regardless of fluctuating demand or access frequency [20]. While computationally trivial to implement, this approach severely restricts spatial efficiency and leads to suboptimal routing distances, serving primarily as a lower-bound benchmark for measuring subsequent optimizations.

The second strategy is *Shared Random Storage* (often colloquially referred to in industry as chaotic storage). In this approach, no dedicated locations exist. Any incoming pallet is routed to a randomly selected, currently available empty slot anywhere in the warehouse network [29]. While this strategy maximizes spatial density by preventing reserved-but-empty slots, it inherently lacks intelligent geographic positioning. Because retrieval typically follows strict First-In-First-Out (FIFO) rules, the system must locate the oldest instance of an SKU wherever it was randomly placed, heavily penalizing the retrieval times of fast-moving items [4].

The third strategy is *FSN-Based Dedicated Storage*. This methodology introduces statistical categorization based on the velocity of inventory turnover, explicitly dividing SKUs into Fast-moving, Slow-moving, and Non-moving (FSN) tiers [21]. By analyzing historical order frequencies rather than economic value, the WMS partitions the rack topologies. Fast-moving items are statically assigned to the most accessible zones nearest to the I/O points, while non-moving items are relegated to the deepest, most distant racks [29]. Though highly efficient in a perfectly stable environment, the rigid nature of this static zoning creates severe vulnerabilities when demand patterns undergo seasonal or sudden micro-structural shifts, leading to systemic degradation over time.

The fourth, and most advanced, framework implemented within the digital twin is the *Dynamic Reslotting Strategy*. This approach merges intelligent initial slotting with proactive, continuous background reorganization (interleaving) [22]. The WMS operates as an adaptive agent, continuously monitoring SKU access frequencies over a rolling 24-hour historical window. The algorithmic evaluation is triggered at discrete 5-minute intervals. To prevent interference with active fulfillment, the physical reorganization only executes during periods of low systemic load, specifically when AGV utilization drops below 50%.

A greedy heuristic identifies "fast-movers" currently stranded in low-quality, distant locations and "slow-movers" occupying high-value, accessible coordinates near the I/O points. The algorithmic decision to execute a background relocation is governed by a localized profit function G:

$$G = (t_{current} - t_{new}) \cdot f_{sku} \cdot \beta - t_{reloc} \quad (7)$$

This localized profit function represents a novel approach to task interleaving, actively evaluating immediate kinematic penalties against long-term routing efficiency. The economic viability weighting factor, β, was empirically calibrated to 0.70. This threshold ensures that background relocations are only executed when the anticipated long-term retrieval savings significantly outweigh the immediate energy consumption and mechanical wear incurred by the move.

In this profit function, the term $(t_{current} - t_{new})$ represents the expected temporal savings (in seconds) per future retrieval operation if the item is moved to the improved location. This delta is multiplied by f_{sku}, the projected access frequency of the item, and β, a weighting factor for economic viability (calibrated to 0.70). From this projected temporal profit, the immediate kinematic cost t_{reloc}, the exact time incurred by the AGVs and AS/RS to physically perform the move, is subtracted [22].

If G yields a positive value, the digital twin autonomously dispatches an AGV to reorganize the physical environment. To prevent spatial deadlocks and race conditions during these complex product swaps, the WMS utilizes dedicated mutex markers (defined as _SWAP_ states) that temporarily lock specific rack coordinates. This guarantees absolute data consistency and collision-free routing while the AGV acts as a mobile buffer [17].

VII. Experimental Setup and Scenario Generation

The quantitative evaluation of these complex assignment strategies requires robust, highly realistic data inputs. Because proprietary industrial order logs are typically confidential, the digital twin utilizes a stochastic scenario generator parameterized by an open-source, real-world eCommerce dataset from a major footwear manufacturer [30]. Analysis of this reference dataset reveals a pronounced turnover distribution, classic to warehousing logistics: merely 24.0% of the SKUs account for 79.8% of the total order frequency (constituting the high-velocity F-articles), while 23.6% of the SKUs represent 15.1% of the orders (S-articles), and the remaining 52.4% represent a mere 5.1% of the orders (the N-articles) [20][30].

Because the reference dataset originally consists of individual footwear SKUs, order lines were computationally aggregated into homogeneous unit loads (full-pallet equivalents). This transformation ensures the demand profiles accurately reflect the physical and volumetric constraints of a pallet high-bay warehouse.

To model natural logistical clustering, operational peaks, and variance, the simulation relies on an exponential distribution for order inter-arrival times. This ensures the system is subjected to realistic burst demands rather than a perfectly uniform, deterministic stream of tasks [31]. A 30-day continuous simulation horizon is utilized. This extended timeframe is critical to ensure statistical validity and to capture both micro-fluctuations (e.g., daily operational cycles) and macro-trends (e.g., seasonal catalog shifts) without being skewed by initial warm-up periods [31].

To rigorously test the resilience and adaptability of the storage algorithms, the digital twin is subjected to four distinct stress scenarios [22]:

- **Scenario 1 (Low Load):** The base arrival rate is tuned so the system operates at approximately 40% AGV utilization. Resources are abundant, queueing delays are minimal, and the environment represents optimal, non-congested operating conditions.
- **Scenario 2 (High Load):** The order arrival rate is drastically increased to push AGV utilization to a sustained 80%. The system operates near its mechanical breaking point, severely limiting robot idle time and inducing heavy spatial congestion at grid intersections.
- **Scenario 3 (Medium Load with Volatility):** Base utilization is maintained at 60%, but the scenario incorporates aggressive day/night demand shifts and a sudden "seasonal shift" at Day 15. At this juncture, the statistical access probabilities of the F-articles and S-articles are inverted to simulate a rapid product catalog transition.
- **Scenario 4 (High Load with Volatility):** This scenario combines the sustained 80% stress load with the day/night cycles and the Day 15 seasonal inversion. It represents the most constrained operating environment for the WMS, designed to test the limits of algorithmic adaptability under severe physical constraints.

VIII. Simulation Results

The execution of the continuous 30-day simulation experiments across the four scenarios generated extensive telemetry datasets, allowing for a deep quantitative analysis of both hardware resource dimensioning and strategic software efficacy.

Prior to comparing the storage strategies, the digital twin was utilized to determine the optimal AGV fleet size. Holding the baseline arrival rate constant, the fleet was scaled from 1 to 30 units. The analysis revealed a distinct non-linear relationship. At fleet sizes below 7 units, the AGVs represented a critical systemic bottleneck, with wait times growing exponentially. As the fleet expanded to 10 units, retrieval times plummeted and stabilized. Increasing the fleet from 10 to 30 AGVs yielded statistically negligible improvements in throughput. Therefore, a fleet of 10 AGVs was selected for all scenario testing, providing a baseline utilization of 60–70% that balances efficiency with realistic capital expenditure (CAPEX) constraints. Conversely, the telemetry data from the 12 AS/RS units revealed chronic underutilization across all scenarios, indicating an asymmetric hardware bottleneck within the reference design.

Table I presents the results for Scenario 1, operating under low load conditions.

 ISBN: 978-1-68558-400-9

TABLE I. SCENARIO 1: LOW LOAD CONDITIONS

Strategy	Storage [s]	Storage Δ[%]	Retrieval [s]	Retrieval Δ[%]	SD [s]
Random	80.48	0.00	89.00	0.00	19.6
Shared	78.86	-2.01	88.30	-0.79	19.8
FSN Fixed	84.99	+5.60	81.70	-8.20	16.4
Dynamic	76.02	-5.54	75.60	-15.06	15.6

Under these ideal conditions, the Dynamic Reslotting strategy significantly outperformed, reducing average retrieval times by 15.06% compared to the baseline (from 89.00 s to 75.6 s). Furthermore, it reduced operational variance, dropping the standard deviation to 15.6 seconds. The FSN strategy performed adequately in retrieval (-8.20%) but suffered during storage (+5.60%), as inbound AGVs frequently bypassed empty slots to reach reserved, deeper zones [29].

Table II presents the results for Scenario 2, operating under sustained high load.

TABLE II. SCENARIO 2: SUSTAINED HIGH LOAD

Strategy	Storage [s]	Storage Δ[%]	Retrieval [s]	Retrieval Δ[%]	SD [s]
Random	86.32	0.00	97.60	0.00	23.7
Shared	83.35	-3.44	94.20	-3.48	22.3
FSN Fixed	93.91	+8.79	94.40	-3.28	21.3
Dynamic	95.15	+10.23	88.70	-9.12	20.2

The high load fundamentally altered the efficiency matrix. The Dynamic strategy maintained its lead in retrieval (-9.12%), but resulted in a corresponding 10.23% increase in storage times (from 86.32 s to 95.15 s). Because the AGVs were constantly occupied with active order fulfillment, the WMS successfully suppressed background reslotting algorithms, temporarily preventing perfect layout optimization to prioritize immediate throughput [22]. Simultaneously, the FSN strategy experienced significant degradation. By statically clustering all fast-movers in the front racks, the FSN strategy created intense spatial hotspots, causing multiple AGVs to queue and gridlock at the same AS/RS transfer nodes [17].

Table III details the system's performance in Scenario 3, incorporating demand volatility and a seasonal shift under medium load.

TABLE III. SCENARIO 3: MEDIUM LOAD WITH VOLATILITY

Strategy	Storage [s]	Storage Δ[%]	Retrieval [s]	Retrieval Δ[%]	SD [s]
Random	80.07	0.00	90.47	0.00	20.4
Shared	78.96	-1.39	89.41	-1.17	20.5
FSN Fixed	85.20	+6.41	83.10	-8.25	18.6
Dynamic	80.20	+0.16	77.62	-14.21	18.1

In Scenario 3, the resilience of the Dynamic slotting strategy became evident. Despite the inversion of product demand profiles at Day 15, the algorithm utilized the night cycle lulls to completely restructure the physical layout. By swapping obsolete stock with newly trending SKUs, it maintained a 14.21% retrieval advantage over the baseline (from 90.47 s to 77.62 s).

Table IV presents the results for Scenario 4, combining high load, high volatility, and a seasonal shift.

TABLE IV. SCENARIO 4: HIGH LOAD WITH VOLATILITY

Strategy	Storage [s]	Storage Δ[%]	Retrieval [s]	Retrieval Δ[%]	SD [s]
Random	83.53	0.00	94.34	0.00	22.4
Shared	81.79	-2.08	93.12	-1.30	22.0
FSN Fixed	89.23	+6.83	89.61	-4.99	21.9
Dynamic	90.04	+7.80	82.65	-12.39	19.9

In the most demanding scenario, the systemic breakdown of static optimization was highly pronounced. A temporal analysis of the FSN strategy revealed that immediately following the Day 15 shift, retrieval times spiked drastically and remained elevated. The fixed assignments became severely suboptimal, with obsolete products blocking premium slots while newly trending products were stranded deep in the warehouse network [21]. In contrast, the Dynamic slotting strategy detected this statistical drift and continuously reallocated inventory during the brief night shifts to recover efficiency, ultimately achieving a sustained 12.39% retrieval advantage even under maximum load (from 94.34 s to 82.65 s). This advantage in scenario 4 is also shown in Figure 2, which compares the retrieval time of the different strategies.

IX. DISCUSSION

The vast telemetry datasets generated by the agent-based simulation yield several higher-order insights with profound implications for the design and deployment of industrial digital twins.

The most critical revelation is the inevitable decay of static optimization models. The failure of the FSN strategy in Scenarios 3 and 4 highlights a primary limitation of static facility design: static allocation models inevitably degrade over time due to stochastic shifts in consumer demand [20]. An industrial digital twin must not merely serve as a passive monitoring dashboard; it must feature continuous, autonomous realignment mechanisms [13]. The superiority of the Dynamic Reslotting strategy proves that optimization is not a singular project to be completed at facility commissioning, but a continuous computational cycle. The digital twin must act as a closed-loop control system, detecting statistical anomalies and autonomously dispatching agents to rebalance the logistical topology before overall throughput is compromised [12].

The success of the dynamic reslotting algorithm relies heavily on the concept of spatiotemporal arbitrage. The algorithm actively trades excess systemic capacity in one time period, such as the lower-demand night shifts, for enhanced kinetic performance during peak daytime operations. During complex

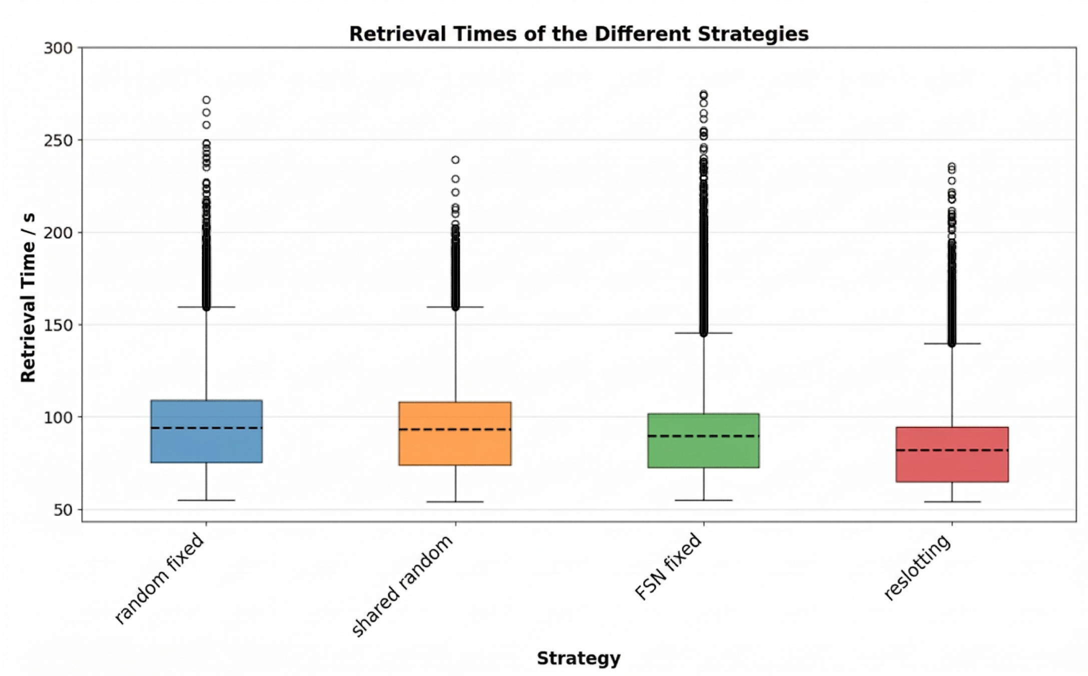

Figure 2. Scenario 4 - Retrieval Times of the Strategies

"swap" operations, the WMS utilizes the AGV's payload capacity as a temporary mobile buffer while the AS/RS maneuvers pallets into new positions [22]. This algorithmic behavior effectively integrates transport and storage capacities. In a sophisticated digital twin, AGVs are no longer viewed simply as delivery vectors moving from point A to point B, but as dynamic, floating nodes within the active memory architecture of the warehouse itself [1].

The observed increase in storage times under the Dynamic strategy is a calculated operational trade-off. By proactively reserving highly accessible slots near the I/O points for fast-moving retrievals, inbound storage tasks are frequently relegated to deeper racks. However, because outbound order fulfillment is typically the critical bottleneck in e-commerce intralogistics, sacrificing storage velocity to guarantee retrieval velocity yields a net-positive systemic throughput.

This leads to a counterintuitive insight regarding the total cumulative distance traveled by the AGV fleet. One might logically assume that the background relocations executed by the dynamic slotting strategy would vastly inflate the total mileage. However, our primary data from Scenario 3 reveals that the total fleet distance for the baseline random strategy was 4,647 km, while the dynamic slotting strategy, including all background relocation drives, totaled only 4,337 km. The optimization of the spatial layout compensates for the added mechanical effort of the reorganization itself. By proactively bringing the most frequently accessed goods closer to the I/O nodes, the marginal travel savings of every subsequent commercial pick accumulate into a net reduction in overall AGV travel distance.

Finally, the simulation exposes the strict limits of physical hardware scaling. The saturation curve of the AGV fleet, demonstrating diminishing returns beyond 10 units, proves that physical scaling cannot unilaterally solve logical bottlenecks. Introducing additional hardware into a highly utilized, constrained spatial graph simply generates exponential traffic conflicts, thereby lowering the marginal utility of each robot [17]. Maximum warehouse throughput is fundamentally bound by software orchestration. The finding that the 12 AS/RS units were chronically underutilized while AGVs were saturated indicates that future dark warehouse facility designs need to evaluate the number of AS/RS units as well.

X. Conclusion and Future Work

The transition to fully automated dark warehouses represents a massive capital investment for the logistics sector. To maximize the return on this infrastructure, enterprises must move beyond rigid control systems and embrace adaptive, AI-driven operations. This paper demonstrated the efficacy of an agent-based simulation acting as the core of an industrial digital twin, quantitatively proving that dynamic, predictive reslotting significantly outperformed static models. The dynamic strategy achieved the shortest retrieval times, reducing them by up to 15.06% and drastically collapsed operational variance, transforming temporal downtime into kinetic efficiency and enhancing overall throughput predictability. Most critically, it demonstrated the unique capacity to autonomously adapt to unforeseen seasonal demand shifts, leveraging off-peak hours to rebalance the physical layout.

Based on the insights gained from this validated architecture, future work must expand the capabilities of the digital twin. To thoroughly validate the proposed greedy heuristic, future investigations will benchmark it against exact Mixed-Integer Linear Programming (MILP) solvers on smaller-scale instances to quantify the optimality gap. Additionally, the evaluation framework will be expanded to financially quantify the added energy consumption, mechanical wear, and labor costs associated with continuous reslotting. Finally, to transition this architecture toward real-world deployment under highly congested multi-agent conditions, future models will evaluate the integration of Deep Reinforcement Learning (DRL) algorithms to replace the heuristic, allowing the warehouse to completely and autonomously self-govern its spatial topology.

ACKNOWLEDGMENTS

This work was developed as part of the move.mORe project. The move.mORe project funded by the Innovation University (IHS) promotes research-based regional transfer of ideas, knowledge and technology in the field of sustainable mobility and is jointly run by the universities of Karlsruhe and Offenburg

REFERENCES

[1] G. Fragapane, R. de Koster, F. Sgarbossa, and J. O. Strandhagen, “Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda,” *European Journal of Operational Research*, vol. 294, no. 2, pp. 405–426, 2021.

[2] S. Winkelhaus and E. H. Grosse, “Logistics 4.0: A systematic review towards a new logistics system,” *International Journal of Production Research*, vol. 58, no. 1, pp. 18–43, 2020.

[3] P. Baker and Z. Halim, “An exploration of warehouse automation implementations: Cost, service and flexibility issues,” *Supply Chain Management: An International Journal*, vol. 12, no. 2, pp. 129–138, 2007.

[4] R. De Koster, T. Le-Duc, and K. J. Roodbergen, “Design and control of warehouse order picking: A literature review,” *European Journal of Operational Research*, vol. 182, no. 2, pp. 481–501, 2007.

[5] J. A. Tompkins, J. A. White, Y. A. Bozer, and J. M. A. Tanchoco, *Facilities planning*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2010.

[6] N. Boysen, R. de Koster, and F. Weidinger, “Warehousing in the e-commerce era: A survey,” *European Journal of Operational Research*, vol. 277, no. 2, pp. 396–411, 2019.

[7] K. Azadeh, R. De Koster, and D. Roy, “Robotized and automated warehouse systems: Review and recent developments,” *Transportation Science*, vol. 53, no. 4, pp. 917–945, 2019.

[8] S. Weyer, M. Schmitt, M. Ohmer, and D. Gorecky, “Towards industry 4.0-standardization as the crucial challenge for highly modular, multi-vendor production systems,” *IFAC-PapersOnLine*, vol. 48, no. 3, pp. 579–584, 2015.

[9] E. Negri, L. Fumagalli, and M. Macchi, “A review of the roles of digital twin in cps-based production systems,” *Procedia Manufacturing*, vol. 11, pp. 939–948, 2017.

[10] G. Lugaresi and M. Macchi, “Digital manufacturing modelling: A review of digital twin applications in machining and intralogistics,” *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1143–1148, 2018.

[11] B. Rodič, “Industry 4.0 and the new simulation modelling paradigm,” *Organizacija*, vol. 50, no. 3, pp. 193–207, 2017.

[12] F. Tao et al., “Digital twin-driven product design framework,” *International Journal of Production Research*, vol. 57, no. 12, pp. 3935–3953, 2019.

[13] W. Kritzinger, M. Karner, G. Trampusch, and M. Sihn, “Digital twin in manufacturing: A categorical literature review and classification,” *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1016–1022, 2018.

[14] S. Wang, J. Wan, D. Li, and C. Zhang, “Towards smart factory for industry 4.0: A self-organized multi-agent system with big data based feedback and coordination,” *Computer Networks*, vol. 101, pp. 158–168, 2016.

[15] C. M. Macal and M. J. North, “Tutorial on agent-based modelling and simulation,” *Journal of Simulation*, vol. 4, no. 3, pp. 151–162, 2010.

[16] P. O. Siebers, C. M. Macal, J. Garnett, D. Buxton, and M. Pidd, “Discrete-event simulation is dead, long live agent-based simulation!” *Journal of Simulation*, vol. 4, no. 3, pp. 204–210, 2010.

[17] Y. A. Bozer and D. A. Aldarondo, “A simulation-based comparison of routing algorithms for multi-load agvs,” *European Journal of Operational Research*, vol. 268, no. 2, pp. 500–516, 2018.

[18] N. Kousi et al., “A digital twin for adapting robot behavior in flexible intralogistics,” *Procedia Manufacturing*, vol. 28, pp. 171–176, 2019.

[19] J. J. Reyes, E. L. Solano-Charris, and J. R. Montoya-Torres, “The storage location assignment problem: A literature review,” *International Journal of Industrial Engineering Computations*, vol. 10, no. 2, pp. 199–224, 2019.

[20] J. J. Bartholdi and S. T. Hackman, *Warehouse & Distribution Science*, Release 0.98. Supply Chain and Logistics Institute, Georgia Institute of Technology, 2019.

[21] W. H. Hausman, L. B. Schwarz, and S. C. Graves, “Optimal storage assignment in automatic warehousing systems,” *Management Science*, vol. 22, no. 6, pp. 629–638, 1976.

[22] J. A. Pazour and H. J. Carlo, “Reslotting strategy for a unit-load automated storage and retrieval system,” *IIE Transactions*, vol. 47, no. 8, pp. 823–839, 2015.

[23] *EN 13698-1: Pallet production specification - Part 1: Construction specification for 800 mm x 1200 mm flat wooden pallets*, European Committee for Standardization (CEN), Brussels, Belgium, 2003.

[24] NetworkX, “Networkx documentation,” 2025, Accessed: 2026-05-27. [Online]. Available: https://networkx.org/

[25] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[26] H. Hwang, W. Baek, and M.-K. Lee, “Travel-time models considering the operating characteristics of the storage and retrieval machine,” *International Journal of Production Research*, vol. 42, no. 14, pp. 2877–2889, 2004.

[27] Y. A. Bozer and J. A. White, “Travel-time models for automated storage/retrieval systems,” *IIE Transactions*, vol. 16, no. 4, pp. 329–338, 1984.

[28] ISO, *Industrial trucks — safety requirements and verification — part 4: Driverless industrial trucks and their systems*, Standard DIN EN ISO 3691-4, International Organization for Standardization, Geneva, CH, 2020.

[29] C. G. Petersen and G. Aase, “An evaluation of routing and volume-based storage policies in an order picking operation,” *Decision Sciences*, vol. 35, no. 2, pp. 269–289, 2004.

[30] R. Furlan de Assis, *Order picking dataset from a warehouse of a footwear manufacturing company*, version 1, Mendeley Data, 2024. Accessed: 2026-05-27. [Online]. Available: https://data.mendeley.com/datasets/pf2w725pw3/1

[31] A. M. Law, *Simulation Modeling and Analysis*, 5th ed. New York: McGraw-Hill Education, 2015.