

Disruption-aware reserve-constrained dispatching of battery-electric AGVs for brownfield manufacturing intralogistics

Muhammad Rafay Ikram

muhammadrarafay.ikram1@gmail.com

Researchers <https://orcid.org/0009-0004-3000-9038>

Research Article

Keywords: automated guided vehicles, manufacturing intralogistics, real-time dispatching, battery management, discrete-event simulation, brownfield manufacturing

Posted Date: August 7th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10560768/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at The International Journal of Advanced Manufacturing Technology on October 7th, 2026. See the published version at <https://doi.org/10.1007/s00170-026-19264-2>.

Abstract

Battery-electric automated guided vehicles (AGVs) in brownfield factories must serve time-critical material movements while preserving enough energy to reach a charger and adapting to temporary route restrictions. This study develops disruption-aware reserve-constrained dispatching (RCRD), an event-driven rule for manufacturing intralogistics. Candidate assignments are first screened for post-task charger reachability and then ranked using normalized empty travel, route inflation, reserve headroom, projected tardiness, and task priority. The manufacturing application is evaluated on a reference production network with ten material-transfer stations, eight AGVs, two chargers, stochastic task arrivals, and four operating profiles. A matched energy-aware comparator uses identical coefficients, charging logic, candidate sets, and random scenarios but omits route inflation. Coefficients were locked before confirmatory testing. Across 240 matched seed-profile observations and 432,000 completed tasks, RCRD reduced mean tardiness by 6.9%; the paired reduction was 3.37 time units (95% confidence interval 1.00-5.78). The reduction reached 11.2% under simultaneous demand and aisle stress and increased under a one-charger bottleneck. Mean dispatch computation was 58.8 microseconds. The results provide an implementable decision rule for fleet managers and manufacturing-system integrators, while showing that route-state information is most valuable when disturbances interact with limited operational slack.

1. Introduction

Automated guided vehicles (AGVs) and autonomous mobile robots are established resources for flexible material flow in manufacturing. They decouple workstations from fixed conveyors and allow production systems to respond to changes in product mix, routing, and demand. Their performance, however, depends on interrelated scheduling, routing, parking, and charging decisions. Reviews of AGV and mobile-robot control show that local assignment choices can transfer congestion, starvation, or charging delay to another part of the production system [1, 2, 14, 15].

These interactions are especially important in brownfield factories. Shared aisles may be restricted by maintenance, forklift traffic, quality holds, or temporary work zones; chargers are finite; and transport demand can change sharply during line recovery or rush production. The closest vehicle may be unable to finish an urgent task and retain enough energy to reach a charger, while a vehicle with ample energy may incur a substantial detour around a blocked aisle. An industrially useful controller must combine urgency, travel, route state, and post-task energy without creating a computational burden at each dispatch event.

Existing methods range from transparent multi-attribute rules [3, 4] to integrated machine-AGV scheduling, multi-objective optimization, and battery-constrained matheuristics [5–8]. Recent benchmarks and adaptive dispatching studies also show that policy choice depends on the operating instance and available information [16, 17]. These developments improve decision quality, but they do not isolate a practical question faced by manufacturing engineers: after an interpretable policy already

considers empty travel, urgency, priority, and battery reserve, what additional value is created by current route-disruption information?

Application-oriented manufacturing research increasingly combines analytical rules with simulation to evaluate intralogistics decisions before deployment. Recent work has used discrete-event simulation to assess AGV implementation in automotive manufacturing [18], dynamic dispatching and lot-splitting in a large assembly job shop [19], and physically informed AGV energy models for production-system planning [20]. The present study follows this applications-based perspective. It evaluates a controller that can be implemented at the fleet-management layer and uses a matched ablation in which the proposed rule and the energy-only comparator differ only in the route-inflation term.

Three questions guide the study. First, does disruption-aware reserve-constrained dispatching improve transport service relative to FIFO-nearest, EDD-nearest, and a matched energy-aware core? Second, under which combinations of demand, aisle disruption, and resource capacity does route-state information provide measurable benefit? Third, can the rule retain sub-millisecond decision latency and a stable directional effect under moderate coefficient changes?

The contribution is an application-ready decision rule and a controlled evaluation of its information value. The method formalizes a post-task reserve that covers both the assigned transport and travel to a reachable charger, embeds that condition in a normalized dispatch score, and evaluates the route term with policy-independent scenarios, separate development and confirmatory seeds, paired inference, capacity tests, and coefficient sensitivity analysis. The intended users are manufacturing-system engineers, AGV fleet managers, and warehouse-control or manufacturing-execution-system integrators responsible for brownfield intralogistics.

2. Related work and research positioning

2.1 Real-time AGV dispatching in manufacturing

AGV control includes guide-path design, fleet sizing, routing, deadlock avoidance, task dispatching, and charging [1, 2, 14, 15]. FIFO-nearest and EDD-nearest remain attractive because their decisions are easy to audit and execute. Their limitation is myopia: age, due date, or proximity alone does not represent downstream energy, route, and workstation consequences. Multi-attribute rules combine several operating signals [3], while dispatching-rule analysis can support system design and policy selection [4]. More recent adaptive and benchmarking research confirms that no single scheduling architecture is best for every intralogistics instance [16, 17].

Integrated scheduling extends the decision boundary. Sivarami Reddy et al. [5] coordinated machines, AGVs, and tools in a flexible manufacturing system, Liu et al. [6] formulated dynamic multi-AGV assignment as a knowledge-guided multi-objective problem, and Singh et al. [7] and Boccia et al. [8] developed battery-constrained formulations and matheuristics. Such approaches can produce high-

quality schedules but require a formal optimization model and repeated solution effort. RCRD addresses the lower-level operating layer in which decisions must be recomputed frequently from live fleet state.

2.2 Battery-aware assignment and charging

Battery limits turn AGV dispatching into a coupled assignment-and-charging problem. Kabir and Suzuki [9] showed that charging duration and policy alter productive fleet hours and manufacturing flexibility. Later studies used feature-based reinforcement learning for charging [10], deep reinforcement learning for combined dispatching, parking, and charging [11], and mixed-integer programming for recharge sequencing [12]. Physically informed AGV energy models further show why production planning should represent movement and payload effects rather than treating energy as a binary constraint [20].

Energy can be represented as a soft score, a hard feasibility condition, or both. A purely soft penalty may still select a task that leaves insufficient energy for a safe charger transfer. A conservative fixed threshold can avoid depletion but remove useful capacity too early. The proposed method therefore combines a hard expected post-task charger reserve with a soft reserve-headroom term. This separates a safety-oriented feasibility decision from the performance preference applied among feasible assignments.

2.3 Disruption response and application-oriented simulation

Resilience-oriented fleet control adapts when route or resource conditions depart from the nominal state. Goutham and Stockar [13] update fleet allocation under blocked passages and failures through an offline-online search architecture. RCRD considers a narrower industrial information setting: the fleet controller receives current detour costs from the route planner but does not maintain a search tree or learned value function. This is consistent with brownfield installations in which route-state data are available but the plant requires transparent assignment logic.

Discrete-event simulation is appropriate when policy effects emerge through stochastic arrivals, queues, charging, and shared paths. Stadtler et al. [16] demonstrate the importance of instance-specific benchmarks for real-time AGV scheduling. Within advanced manufacturing, Papa and Panico [18] use simulation to validate AGV-supported automotive logistics decisions, while Sudevan et al. [19] combine priority dispatching with simulation in a large assembly job shop. These studies establish an applications-based standard: the model should represent a concrete manufacturing decision, expose relevant resource interactions, and translate results into implementable engineering guidance.

2.4 Research gap and positioning

The unresolved question is the incremental value of current route-disruption information after energy reserve, urgency, empty travel, and priority have already been represented. A defensible test requires a comparator that changes no other policy feature, identical exogenous scenarios, and confirmatory seeds that were not used during coefficient selection. RCRD is evaluated under those conditions in a reference

brownfield manufacturing application. The method is not presented as globally optimal; its purpose is to determine whether a low-computation route signal is operationally useful and to identify the manufacturing conditions in which that value is most likely to appear.

3. Dynamic transport problem

3.1 System definition and assumptions

Let V denote a fleet of homogeneous battery-electric AGVs and $J(t)$ the set of transport jobs waiting at decision time t . Job j has release time r_j , pickup node o_j , delivery node d_j , due time q_j , and priority $p_j \in \{0,1,2\}$. Vehicle v has current node x_v , availability time a_v , and state of charge s_v . A finite set C of non-preemptive chargers has time-varying availability. Jobs are single-load pickup-and-delivery requests. The route controller supplies a current effective distance $D_t(a,b)$ and nominal distance $D_0(a,b)$.

The operating objectives are to reduce mean and tail tardiness, protect on-time service, limit empty travel, and shorten the completion horizon without assigning a vehicle that lacks expected energy to finish the task and reach a charger. The study treats safety interlocks, collision avoidance, and vendor path execution as lower-level hard controls outside the dispatch score.

3.2 Distance, route inflation, and reserve feasibility

$$D_{vj}(t) = D_t(x_v, o_j) + D_t(o_j, d_j) \quad (1)$$

The effective distance for candidate pair (v,j) is the sum of empty travel to pickup and loaded travel to delivery. Route inflation is the excess over the corresponding nominal path:

$$I_{vj}(t) = D_{vj}(t) - [D_0(x_v, o_j) + D_0(o_j, d_j)]. \quad (2)$$

Expected post-task state of charge is $s_{vj}^{post} = s_v - eD_{vj}(t)$, where e is expected energy use per distance unit. Let $c^*(d_j, t)$ be the nearest feasible charger after delivery. A pair is reserve-feasible when:

$$s_{vj}^{post} \geq s_{safe} + eD_t(d_j, c^*(d_j, t)). \quad (3)$$

The safety reserve s_{safe} is eight state-of-charge percentage points. The soft reserve feature is $R_{vj} = \max\{0, 12 - h_{vj}\}$, where h_{vj} is the headroom above the hard requirement. Thus, a pair can be feasible yet receive a penalty when its remaining margin is small.

4. Disruption-aware reserve-constrained dispatching

4.1 Normalized vehicle–job score

At a job-release, vehicle-completion, or charge-completion event, the controller orders waiting jobs by due date, priority, release time, and identifier and retains at most $K = 18$ urgent candidates. For every available vehicle and candidate job, it applies the hard reserve screen and computes normalized empty distance $E_{vj}/10$, route inflation $I_{vj}/5$, soft reserve penalty $R_{vj}/20$, projected tardiness $T_{vj}/50$, and priority $p_j/2$. The locked RCRD score is:

$$C_{vj} = 1.5625(E_{vj}/10) + 3.75(I_{vj}/5) + 3.125(R_{vj}/20) + 1.5(T_{vj}/50) - 0.6(p_j/2). \quad (4)$$

The feasible pair with the lowest score is assigned. The coefficients express operating trade-offs; they are not physical constants. Selection used only development seeds and a pre-specified composite of mean tardiness, tail tardiness, on-time service, empty travel, and completion horizon. The chosen values were then fixed before the confirmatory experiments in Section 5. The matched energy-core comparator uses Eq. (4) with the route-inflation coefficient set to zero; every other policy element remains unchanged.

4.2 Charging and conventional policies

All policies share the same expected reserve-feasibility screen and a fixed charging target of 85% state of charge. When no reserve-feasible vehicle–job pair exists, the available vehicle with the earliest achievable charge completion is routed to a charger. This common charging controller prevents the proposed method from gaining an artificial advantage through different charging thresholds or targets.

FIFO-nearest evaluates the full waiting queue, gives precedence to the oldest job for which at least one vehicle is reserve-feasible, and selects the nearest feasible vehicle. EDD-nearest follows the same logic but gives precedence to the earliest due date. The matched energy-core and RCRD share the same 18-job urgency shortlist. These definitions preserve the standard meaning of the conventional rules while ensuring an exact information ablation for the principal comparison.

4.3 Computational complexity

Let W be the number of waiting jobs, F the number of available vehicles, and $K = \min(W, 18)$. Queue ordering requires $O(W \log W)$, and feature evaluation requires $O(FK)$ for the score-based policies. Because K is capped at 18, the pairwise calculation remains bounded in practice. In the confirmatory runs, the mean RCRD decision took 58.8 microseconds, and the mean run-level 95th percentile was 81.6 microseconds on the stated test environment. These times are well below a sub-second fleet-control cycle and require neither offline training nor a repeated global optimization solve.

5. Manufacturing application and experimental design

5.1 Reference brownfield manufacturing system

The application represents a brownfield production area in which battery-electric AGVs connect ten material-transfer interfaces distributed among machining, subassembly, inspection, final assembly, and

supermarket functions. The 7 x 6 rectilinear network contains four parking positions, eight AGVs, and two chargers. Jobs represent unit-load transfers between production interfaces and arrive dynamically with an origin, destination, priority, and delivery allowance. Each run contains 450 completed tasks. The model is vendor-neutral and uses normalized time, distance, and state-of-charge units so that the dispatching logic can be calibrated to a plant without disclosing proprietary layout or production data.

Travel time includes stochastic speed and handling variation, and energy use is proportional to effective distance with a job-level perturbation. Temporary restrictions are generated independently of policy decisions in 15-time-unit intervals and add detour penalties to routes crossing affected horizontal or vertical zones. Job attributes, disruption schedules, initial battery states, and stochastic multipliers are generated before policy execution, ensuring that every policy receives the same exogenous manufacturing scenario.

For deployment, the model assumes a standard information interface: the manufacturing execution or warehouse-control system supplies task release, due time, and priority; the fleet manager supplies vehicle location and state of charge; the route planner supplies effective and nominal path lengths; and the charger controller supplies reachable charger state. RCRD returns a vehicle-job assignment or a charge command. No plant-identifying data are required by the score itself.

5.2 Operating profiles and manufacturing stress cases

Table 1
Operating profiles used in the locked confirmatory experiment

Profile	Arrival rate	Block probability	Blocked zones	Travel factor	Due factor
Nominal	0.34	0.015	1	1.00	2.40
Demand surge	0.46	0.020	1	1.04	2.15
Aisle disruption	0.35	0.120	3	1.10	2.30
Combined stress	0.45	0.100	3	1.14	2.05

Four profiles separate routine production from operational stress. Demand surge increases arrival intensity and tightens delivery allowances while maintaining low blockage probability. Aisle disruption retains near-nominal demand but increases the frequency and spatial extent of route restrictions. Combined stress applies high arrival intensity, multiple blocked zones, the largest travel factor, and the tightest due-date factor. The design therefore distinguishes the effect of route-state information under ordinary operation, demand pressure, route disturbance, and interacting disturbances.

5.3 Development-confirmation separation and experiment size

Coefficient screening used only development seed blocks (10–29, 100–139, and 900–919). After selecting Eq. (4), the coefficients and charging policy were locked. No result from the confirmatory seed blocks was used to revise the method. The main confirmatory experiment used seeds 1300–1359: 60 seeds × four profiles × four policies × 450 jobs = 432,000 completed jobs.

Resource-capacity robustness used seeds 1400–1439 and compared RCRD with the matched energy-core under three configurations—six AGVs/two chargers, eight AGVs/one charger, and ten AGVs/two chargers—under demand-surge and combined-stress profiles. This design produced 216,000 jobs. One-factor-at-a-time sensitivity used seeds 1500–1519, $\pm 25\%$ perturbations of five coefficients, three stress profiles, and the matched core on the same seeds, producing 324,000 jobs. The complete locked evidence base therefore contains 972,000 completed jobs.

5.4 Performance measures and statistical analysis

The primary endpoint is run-level mean tardiness, where task tardiness is $\max\{0, c_j - q_j\}$. Secondary measures are 95th-percentile tardiness, on-time rate, empty-travel share, and completion horizon. Mean wait, tail cycle time, utilization, charge service, and charger queueing are retained in the reproducibility files. Empty-travel share is empty distance divided by total vehicle distance.

The principal comparison is paired by seed and profile. Improvement is defined as matched-core minus RCRD for delay, movement, and horizon measures, and as RCRD minus matched-core for on-time rate. The paper reports 20,000-resample paired bootstrap confidence intervals. A two-sided Wilcoxon signed-rank test is reported as a distribution-robust sensitivity analysis. Holm adjustment is applied within each five-outcome family for secondary confirmatory interpretation. Standardized paired effects are reported as d_z . Conventional-rule comparisons are descriptive paired benchmarks rather than the basis of the novelty claim.

5.5 Model verification and reproducibility

Verification included deterministic reruns from fixed seeds, policy-independent scenario generation, checks that the matched core and RCRD differ only in route-inflation weight, exact reproduction of aggregate tables from run-level files, and a smoke test covering release, dispatch, charging, and completion events. The reproducibility package contains executable source code, locked coefficients, development rankings, raw confirmatory results, analysis scripts, figures, package versions, and checksums.

5.6 Use of generative AI tools

OpenAI ChatGPT was used to support code review, experimental-design checks, literature cross-checking, language editing, and document preparation. All numerical results were generated by the supplied deterministic simulator and verified against the run-level output files. The author reviewed the code, references, analyses, and final manuscript and retains full responsibility for the work.

6. Results

6.1 Overall policy performance

Table 2
Overall means across 240 runs per policy in the main confirmatory experiment

Policy	Mean tard.	P95 tard.	On time (%)	Empty share (%)	Horizon	Decision μ s
FIFO-nearest	113.78	236.07	16.9	45.7	1481.7	149.9
EDD-nearest	112.13	234.54	16.6	45.5	1475.1	146.0
Matched energy-core	49.01	126.61	58.2	38.1	1347.0	56.4
Proposed RCRD	45.64	122.50	58.7	38.1	1337.8	58.8

RCRD had the lowest mean tardiness, 95th-percentile tardiness, and completion horizon among the four policies. Compared with FIFO-nearest, mean tardiness fell by 59.9%, tail tardiness by 48.1%, empty-travel share by 16.6%, and completion horizon by 9.7%; on-time service increased by 41.8 percentage points. The corresponding reductions relative to EDD-nearest were 59.3%, 47.8%, 16.4%, and 9.3%. These comparisons show the value of an energy-aware multi-attribute policy, but they do not isolate the contribution of route information. That question is addressed by the matched comparison below.

6.2 Matched incremental effect of route-disruption information

Table 3
RCRD improvement over the matched energy-aware core across 240 paired observations; positive values favor RCRD

Outcome	Mean improvement	Bootstrap 95% CI	Wilcoxon p	Holm p	d_z	Unit
Mean tardiness	3.3683	[1.0018, 5.7760]	0.0540	0.2245	0.177	time units
P95 tardiness	4.1056	[-2.1603, 10.3127]	0.2150	0.6449	0.083	time units
On-time rate	0.0054	[-0.0075, 0.0187]	0.8576	1.0000	0.052	proportion
Empty-travel share	-0.0003	[-0.0019, 0.0012]	0.6356	1.0000	-0.026	proportion
Completion horizon	9.2403	[1.0066, 17.5153]	0.0449	0.2245	0.143	time units

Mean tardiness fell from 49.01 to 45.64 time units, a reduction of 3.37 time units (6.9%). The paired bootstrap interval, [1.00, 5.78], did not include zero. The Wilcoxon test was close to the conventional threshold ($p = 0.054$), which is consistent with a small, skewed effect that was not present in every scenario. Completion horizon improved by 9.24 time units, with a bootstrap interval of [1.01, 17.52], although the multiplicity-adjusted rank test was inconclusive. Confidence intervals for tail tardiness, on-time rate, and empty-travel share included zero. The overall evidence therefore supports a reduction in average delay and a possible reduction in completion horizon, not across-the-board superiority on every measure.

6.3 Scenario dependence

Table 4
Matched energy-aware core and RCRD performance by operating profile

Profile	Core mean tard.	RCRD mean tard.	Reduction (%)	Core on time (%)	RCRD on time (%)	Horizon gain
Nominal	8.26	6.92	16.2	88.4	89.0	12.5
Demand surge	71.67	70.76	1.3	39.9	40.9	8.2
Aisle disruption	17.51	17.35	0.9	76.2	75.5	0.3
Combined stress	98.59	87.53	11.2	28.1	29.5	15.9

The clearest benefit appeared under combined stress. Mean tardiness fell from 98.59 to 87.53 time units, an 11.2% reduction. The paired improvement was 11.06 time units [3.33, 18.75], with a Holm-adjusted Wilcoxon p -value of 0.0488. Differences were smaller under nominal operation and demand surge. Under isolated aisle disruption, mean tardiness was nearly unchanged, while tail tardiness and empty travel were slightly worse with RCRD. Route information therefore helped most when detours coincided with high demand and limited slack; a blocked aisle by itself was not enough to guarantee an improvement.

6.4 Resource-capacity robustness

Table 5
Resource-capacity robustness of the RCRD mean-tardiness advantage

Capacity configuration	Profile	Mean-tardiness gain	Bootstrap 95% CI
6 AGVs / 2 chargers	Demand surge	-0.90	[-6.40, 4.41]
6 AGVs / 2 chargers	Combined stress	7.98	[-2.91, 18.54]
8 AGVs / 1 charger	Demand surge	1.46	[-1.74, 4.73]
8 AGVs / 1 charger	Combined stress	14.00	[6.33, 21.74]
10 AGVs / 2 chargers	Demand surge	0.27	[-4.27, 4.19]
10 AGVs / 2 chargers	Combined stress	6.98	[-0.81, 14.80]

The strongest capacity result occurred with one charger under combined stress. RCRD improved mean tardiness by 14.00 time units [6.33, 21.74], tail tardiness by 31.40 [15.62, 47.07], and completion horizon by 34.16 [15.88, 52.04]. With six AGVs under combined stress, completion horizon improved by 28.34 [2.55, 54.07], although the mean-tardiness interval included zero. Most other conditions were neutral or directionally favorable. The exception was the fleet-constrained demand-surge case, where movement and horizon slightly favored the matched core. These results point to charger scarcity and combined stress—not resource scarcity in general—as the conditions in which route information is most helpful.

6.5 Coefficient sensitivity and computation

On the independent sensitivity seeds, the fixed coefficient set retained a 5.23-time-unit mean-tardiness advantage over the matched core. Nine of the ten one-at-a-time $\pm 25\%$ perturbations also produced a positive advantage. The direction reversed only when the empty-distance coefficient was reduced by 25%, yielding a -7.56-time-unit difference; increasing the same coefficient by 25% yielded a 16.00-time-unit advantage. Empty-distance weighting was therefore the most influential tuning choice in this experiment, and the reported coefficients should not be transferred to another plant without calibration.

The mean computation time per RCRD decision was 58.8 microseconds, compared with 56.4 microseconds for the matched core. Across runs, the mean 95th-percentile latency was 81.6 microseconds. Computing route inflation therefore added little absolute overhead in this benchmark.

7. Discussion

7.1 Contribution to advanced manufacturing

The principal contribution is not a new global optimizer but an implementable manufacturing-control rule whose added information can be isolated. The matched design holds the candidate set, coefficients, charging policy, reserve condition, and exogenous scenario constant. This allows the route-inflation term to be interpreted as an engineering feature rather than as one element of a broadly different policy

architecture. The result is directly relevant to plants that can obtain detour estimates from a route planner but cannot support repeated optimization or retraining.

The experiments also show that additional state information is not automatically beneficial. Route inflation produced little change during demand-only or isolated-disruption profiles and occasionally traded small increases in travel for lower delay. Its value became clearer when disruption interacted with high demand or limited charging capacity. For advanced manufacturing systems, the implication is that controller complexity should be justified by the operating regime rather than by the availability of data alone.

7.2 Relationship to manufacturing intralogistics research

RCRD extends multi-attribute dispatching [3] with an explicit post-task charger reserve and a live route-disruption signal. It occupies a different implementation layer from integrated scheduling and matheuristics [5–8], and a less data-intensive layer than adaptive or learning-based dispatching [11, 17]. The application logic complements recent advanced-manufacturing studies that use simulation to support AGV implementation and production scheduling [18, 19]. The reserve formulation is also consistent with evidence that energy consumption and charging policy change usable fleet capacity [9–12, 20].

7.3 Scope and interpretation of the observed benefit

RCRD improved the pre-specified primary endpoint overall and produced a clear gain under combined stress, but secondary outcomes were mixed and the standardized paired effect was small. The bootstrap interval for mean tardiness excluded zero, whereas the multiplicity-adjusted rank test was not significant. The result should therefore be interpreted as conditional engineering evidence rather than universal policy dominance. The manufacturing value lies in identifying a low-cost control feature that becomes useful when route restrictions interact with limited slack.

8. Manufacturing application and deployment guidance

Implementation requires five signals that are normally available across a manufacturing execution system, warehouse-control system, or AGV fleet manager: released transport requests with due time and priority; vehicle position and state of charge; current path distance or route inflation; charger location and availability; and calibrated travel-energy factors. The reserve screen should remain a hard rule outside the weighted performance score, while vendor safety limits and traffic-control interlocks remain authoritative.

A practical deployment should proceed in four stages. First, export historical tasks, route states, battery traces, and charger events to calibrate distance and energy scales. Second, replay the current policy and RCRD in a digital twin or offline event log. Third, operate RCRD in shadow mode and compare recommended assignments with the live fleet manager. Fourth, activate the rule on a limited route set or

shift while retaining the incumbent controller as a fallback. This sequence separates algorithm evaluation from safety certification and production authorization.

The results also provide a selection rule. Plants with stable routes and ample charging capacity are unlikely to justify a large route-inflation weight. Plants experiencing frequent temporary closures, high transport backlog, or charger bottlenecks are stronger candidates. The one-charger combined-stress experiment illustrates this boundary: the same route signal that was nearly neutral in several ordinary settings produced a substantial delay reduction when detours and charging contention interacted.

The likely users are manufacturing and intralogistics engineers responsible for AGV commissioning, fleet-management configuration, and material-flow performance. Because the score uses arithmetic features and evaluates only the most urgent waiting tasks, it can be implemented in a fleet manager, warehouse-control layer, or manufacturing-execution interface without a separate training infrastructure. Plant-specific coefficient governance remains necessary after layout, battery, product-mix, or routing changes.

9. Limitations and future research

The study uses a representative manufacturing network rather than a named factory, and its parameters are not calibrated to a specific AGV vendor. The network uses rectilinear distance and zone-level detour penalties instead of continuous traffic simulation. Vehicle interactions, deadlocks, pedestrian yielding, payload-dependent energy, battery aging, temperature, nonlinear charging, communication latency, maintenance failure, and charger failure are outside the model. These choices support reproducible comparison but limit plant-level external validity.

The hard reserve is based on expected rather than worst-case energy use. A production deployment should use vendor battery estimation, route-specific safety margins, and conservative fallback logic. The score also assumes that route inflation is current and sufficiently accurate; delayed or noisy detour estimates could reduce the observed benefit. The application evidence therefore supports pre-deployment evaluation, not direct use without calibration and safety validation.

Future research should validate the controller with timestamped task, route-controller, state-of-charge, and charger records from multiple plants. A digital twin could replace zone penalties with predictive travel-time distributions and explicitly model collision-free routing. Further work should compare rolling-horizon mixed-integer assignment, auction mechanisms, and learning-based policies on a shared benchmark, and examine heterogeneous fleets, opportunity charging, battery-aging cost, inaccurate route information, and online coefficient adaptation.

10. Conclusion

This study developed a disruption-aware reserve-constrained dispatching rule for battery-electric AGVs in brownfield manufacturing intralogistics. Candidate assignments are screened for expected post-task

charger reachability and ranked using normalized empty distance, route inflation, reserve headroom, projected tardiness, and task priority. The matched comparator removes only route inflation, allowing the added value of current route state to be estimated without changing the remainder of the controller.

Across the locked confirmatory experiment, RCRD reduced mean tardiness by 6.9% relative to the matched energy-aware core. The paired reduction was 3.37 time units with a 95% bootstrap interval of [1.00, 5.78]. Mean tardiness fell by 11.2% under combined demand and aisle stress, and the advantage increased under a one-charger bottleneck. Results in several low-stress settings were similar, while mean decision time remained 58.8 microseconds.

The method is therefore most relevant to manufacturing systems in which temporary route restrictions interact with demand pressure or charging contention. It provides a transparent control option for fleet managers and system integrators, but it is not a universal replacement for an incumbent policy. Factory calibration, shadow-mode testing, and vendor-specific safety validation are required before deployment.

Declarations

Competing interests

The author has no relevant financial or non-financial interests to disclose.

Funding

The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Author contribution

Muhammad Rafay Ikram performed the conceptualization, methodology, software development, validation, formal analysis, investigation, data curation, visualization, writing of the original draft, review and editing, and project administration. The author read and approved the final manuscript.

Data availability

The complete reproducibility package accompanying this submission contains the simulator, locked coefficients, development rankings, raw run-level data, analysis scripts, figures, software-environment record, and checksums. The package is supplied as Supplementary Information with the submission. A public archival DOI will be added before publication if the manuscript is accepted.

References

1. Le-Anh T, de Koster MBM (2006) A review of design and control of automated guided vehicle systems. *Eur J Oper Res* 171:1–23. <https://doi.org/10.1016/j.ejor.2005.01.036>
2. Vis IFA (2006) Survey of research in the design and control of automated guided vehicle systems. *Eur J Oper Res* 170:677–709. <https://doi.org/10.1016/j.ejor.2004.09.020>
3. Jeong BH, Randhawa SU (2001) A multi-attribute dispatching rule for automated guided vehicle systems. *Int J Prod Res* 39:2817–2832. <https://doi.org/10.1080/00207540110051860>
4. Shen YC, Kobza JE (1998) A dispatching-rule-based algorithm for automated guided vehicle systems design. *Prod Plan Control* 9:47–59. <https://doi.org/10.1080/095372898234514>
5. Sivarami Reddy N, Ramamurthy DV, Padma Lalitha M, Prahlada Rao K (2022) Integrated simultaneous scheduling of machines, automated guided vehicles and tools in multi machine flexible manufacturing system using symbiotic organisms search algorithm. *J Ind Prod Eng* 39:317–339. <https://doi.org/10.1080/21681015.2021.1991014>
6. Liu L, Qu T, Thürer M, Ma L, Zhang Z, Yuan M (2023) A new knowledge-guided multi-objective optimisation for the multi-AGV dispatching problem in dynamic production environments. *Int J Prod Res* 61:6030–6051. <https://doi.org/10.1080/00207543.2022.2122619>
7. Singh N, Dang QV, Akcay A, Adan I, Martagan T (2022) A matheuristic for AGV scheduling with battery constraints. *Eur J Oper Res* 298:855–873. <https://doi.org/10.1016/j.ejor.2021.08.008>
8. Boccia M, Masone A, Sterle C, Murino T (2023) The parallel AGV scheduling problem with battery constraints: a new formulation and a matheuristic approach. *Eur J Oper Res* 307:590–603. <https://doi.org/10.1016/j.ejor.2022.10.023>
9. Kabir QS, Suzuki Y (2018) Increasing manufacturing flexibility through battery management of automated guided vehicles. *Comput Ind Eng* 117:225–236. <https://doi.org/10.1016/j.cie.2018.01.026>
10. Lin CC, Chen KY, Hsieh LT (2023) Real-time charging scheduling of automated guided vehicles in cyber-physical smart factories using feature-based reinforcement learning. *IEEE Trans Intell Transp Syst* 24:4016–4026. <https://doi.org/10.1109/TITS.2023.3234010>
11. Singh N, Akcay A, Dang QV, Martagan TG, Adan IJBF (2024) Dispatching AGVs with battery constraints using deep reinforcement learning. *Comput Ind Eng* 187:109678. <https://doi.org/10.1016/j.cie.2023.109678>
12. Lo Iudice F, Olm JM, Suárez R (2024) Optimal recharge sequencing in multi-AGV systems: a mixed ILP approach. *Comput Ind Eng* 196:110510. <https://doi.org/10.1016/j.cie.2024.110510>
13. Goutham M, Stockar S (2024) Resilient fleet management for energy-aware intra-factory logistics. In: 2024 American Control Conference, pp 839–844. <https://doi.org/10.23919/ACC60939.2024.10644599>
14. Fragapane G, de Koster R, Sgarbossa F, Strandhagen JO (2021) Planning and control of autonomous mobile robots for intralogistics: literature review and research agenda. *Eur J Oper Res* 294:405–426.

<https://doi.org/10.1016/j.ejor.2021.01.019>

15. De Ryck M, Versteyhe M, Debrouwere F (2020) Automated guided vehicle systems, state-of-the-art control algorithms and techniques. *J Manuf Syst* 54:152–173.
<https://doi.org/10.1016/j.jmsy.2019.12.002>
16. Stadtler J, Stahmann P, Janiesch C (2026) A benchmark for selecting real-time scheduling approaches for fleets of automated guided vehicles in intralogistics. *Bus Inf Syst Eng*.
<https://doi.org/10.1007/s12599-026-01003-2>
17. Zou G, Liu H (2025) An adaptive dispatching approach for automated guided vehicles system. *Discov Appl Sci* 7:1332. <https://doi.org/10.1007/s42452-025-07893-3>
18. Papa I, Panico M (2026) Simulation-based logistics optimisation in automotive manufacturing: from route rationalisation to smart AGV implementation. *Int J Adv Manuf Technol* 144:3001–3015.
<https://doi.org/10.1007/s00170-026-18039-z>
19. Sudevan S, Krishnaswamy R, Mohanavelu T et al (2026) An integrated MCDM-simulation approach to study the effect of lot-splitting in a large-scale assembly job shop: a case study. *Int J Adv Manuf Technol* 144:3449–3472. <https://doi.org/10.1007/s00170-026-18015-7>
20. Meißner M, Massalski L (2020) Modeling the electrical power and energy consumption of automated guided vehicles to improve the energy efficiency of production systems. *Int J Adv Manuf Technol* 110:481–498. <https://doi.org/10.1007/s00170-020-05796-8>

Figures




Figure 1




Figure 2




Figure 3




Figure 4




Figure 5

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- MANIFEST.csv
- OnlineResource1SupplementaryInformation.pdf
- README.md
- SHA256SUMS.txt
- analyzerresults.py
- candidatelerativeranking.csv
- conventionalcomparisons.csv
- coreweights.csv
- environment.txt
- mainbyprofile.csv
- mainbyrun.csv
- mainoverall.csv
- pairedinference.csv
- ranking.csv
- rcrdweights.csv
- requirements.txt
- robustnessbyrun.csv
- robustnessinference.csv
- robustnesssummary.csv
- runlockedexperiment.py
- sensitivitybyrun.csv
- sensitivitycorebyrun.csv
- sensitivitycoresummary.csv
- sensitivitysummary.csv
- sensitivityvscore.csv
- simulator.py
- smoketest.py