

Drum-Buffer-Rope and rolling horizon MILP for shifting bottleneck synchronization in sugar beet logistics

Corneliu Iațco

*Academy of Economic Studies of Moldova, Ph.D. Candidate, Chisinau, Republic of Moldova
corneliu.iatco@gmail.com*

Abstract

Objectives: This study develops and validates a cyber-physical mathematical framework for synchronizing stochastic agricultural harvesting with deterministic industrial processing in regional sugar beet logistics. The framework addresses real-time coordination conflicts between distributed field supply nodes, limited transport capacity, dynamic field accessibility, and continuous factory demand. **Prior work:** Traditional static logistics models treat agro-industrial cargo as inert mass and primarily optimize spatial distance. Such approaches are poorly suited to environments characterized by high-frequency schedule nervousness, weather-driven field accessibility changes, shifting bottlenecks, and non-linear raw material degradation during temporary field storage. **Approach:** The proposed Digital Twin architecture combines dynamic Drum-Buffer-Rope (DBR) pull logic for shifting bottlenecks with rolling Model Predictive Control (MPC) and a multi-phase Mixed-Integer Linear Programming (MILP) heuristic. Spatial transport work is handled during seasonal stratification, feasible route assignment, and topological filtering; within the active rolling MILP core it is treated as a quasi-invariant structural term. This allows the optimization layer to focus on time-sensitive trade-offs between sucrose preservation, factory starvation avoidance, and logistical feasibility. **Results:** Empirical validation using an 11-year industrial dataset comprising more than 180,000 unique transactional cycles shows that the multi-phase heuristic achieves sub-second solver response times in the tested industrial scenarios, with the active MILP core converging within 0.45 seconds. Operational deployment is associated with compression of factory reception-yard turnaround time from 24.6 to 13.8 minutes and a 41% reduction in active fleet dispersion. **Implications:** The adaptive sigmoid switching mechanism dynamically shifts objective priorities between quality preservation and plant starvation avoidance during environmental shocks. For practitioners, the framework provides a real-time execution tool for synchronizing harvesting, loading, transportation, and factory reception under stochastic disturbances. For researchers, it demonstrates a scalable approach to time-based buffer management in temporal agro-industrial networks. **Value:** The scientific novelty lies in integrating adaptive DBR pull constraints, rolling MILP optimization, graph-based topological filtering, and physics-informed quality inputs into a unified cyber-physical synchronization framework. The proposed architecture transforms agricultural logistics from a rigid cost center into a resilience-oriented decision system without compromising grower equity or systemic transparency.

Keywords: Smart Regional Development; Heuristic Optimization; Real-Time Control; Supply Chain Resilience; Fleet Coordination; Cyber-Physical Systems.

1. Introduction

The transition to Industry 4.0 paradigms in agro-industrial operations demands a fundamental shift from reactive dispatching to cyber-physically synchronized supply chain management [1]. Within the smart regional development discourse, this imperative represents a core component of institutional resilience and governance quality [2]. Nowhere is this transition more critical than in the sugar beet complex, where structural friction between stochastic agricultural harvesting and rigid processing cycles generates persistent operational inefficiencies across regional economies [3].

Conventional logistics planning relies on classical "push" models designed to minimize travel distance [4]. Under real-world agricultural conditions, this approach systematically fails, producing two macro-systemic dysfunctions. First, chaotic vehicle arrivals at

factories create structural bottlenecks, with queue times ranging from 20 to 55 minutes, artificially bloating fleet size and operational costs [5]. Second, biomass degradation during open field storage is ignored, causing irreversible sucrose loss from respiration and microbial activity before evacuation [6].

Recent operational research covers vehicle routing with time windows [4], heterogeneous fleets [4], and producer fairness [7]. However, these frameworks assume deterministic data or require computationally expensive global replanning impractical under erratic weather, soil, and machinery conditions [3]. A critical research gap remains: logistics models lack a mechanism for dynamic throughput identification that adapts to shifting bottlenecks across the entire supply chain, from field harvesting to factory processing.

To address this gap, this study introduces a multi-layer cyber-physical Digital Twin architecture [8, 9] integrating three interconnected methodological pillars that collectively replace spatial distance minimization with temporal degradation-driven synchronization:

- *Dynamic Drum-Buffer-Rope (DBR) logic* adapted for shifting agricultural bottlenecks, continuously subordinating the logistical transit rhythm to the network's true constraint [10, 11, 12];
- *Model Predictive Control (MPC)* operating on an iterative, daily rolling horizon simulation paradigm to eliminate operational schedule nervousness via synchronized graph data assimilation [1, 13, 14];
- *Predictive biochemical quality streams* ingested from an external companion physics-informed neural network (PINN) architecture [15, 16] to execute quality-aware temporal triage of decaying field piles [6,7].

The scientific contribution and novelty of this research are structured around three core pillars:

- *Algorithmic Formalization of a Dynamic Drum.* Unlike classical DBR architectures where the system constraint remains a static mechanical asset [11, 12], the proposed framework identifies the binding bottleneck at each computational interval as a multi-echelon minimum function balancing remaining field inventory, real-time harvesting capabilities, rolling factory demand, and active fleet carrying capacities [18].
- *Mathematical Uncoupling via Multi-Phase Temporal Decomposition.* The proposed rigid multi-phase heuristic shifts computational complexity from the spatial domain (global vehicle routing optimization) to the temporal domain (priority queue management). The centralized MILP core [19, 20] is utilized exclusively to validate aggregate daily capacity targets and structural boundary conditions, while high-frequency micro-routing execution is transferred to a Sequential Forward-Backward Scheduling heuristic, achieving solver convergence in under one second.
- *Bounded Risk-Driven Override Supervision.* An adaptive sigmoidal objective switching mechanism was developed, operating without numerical tariff competition. By treating spatial transport work as a quasi-invariant structural term after seasonal stratification,

feasible route assignment, and topological filtering, the supervisor handles the operational trade-off between biomass degradation risk and processing starvation penalties. This allows the system to deliberately accumulate inventory up to the factory's absolute physical ceiling (B_{max}) during emergency quality triage, a behavior that constitutes a risk-driven override of the standard pull signal rather than an optimization artifact.

In a broader smart governance context, this framework moves beyond technological novelty toward systemic resilience and public value [21]. By eliminating information asymmetry between independent agricultural producers and centralized industrial processors, the architecture establishes a transparent execution environment [22], strengthening institutional trust and mitigating Principal-Agent conflicts within regional agro-industrial clusters [2].

The remainder of this paper is organized as follows. Section 2 provides a complete formalization of the cyber-physical graph topology, the multi-layer data assimilation pipeline [22], and the mathematical MILP core [19, 20] coupled with the multi-phase heuristic. Section 3 validates the framework empirically using the 11-year industrial dataset, evaluating computational performance, link formation, and turnaround dynamics. Section 4 discusses the implications for schedule nervousness and paradigm shifts, and Section 5 concludes with key findings and future research directions.

2. Materials and Methods

This section presents the comprehensive methodological framework and architectural blueprint of the Digital Twin designed to synchronize the stochastic processes of agricultural harvesting with the deterministic cycles of industrial processing [8, 9]. The proposed architecture integrates cyber-physical graph topology, a multi-layer data ontology, and a hybrid optimization core combining MILP [19, 20], DBR pull logic [10, 11, 12], and MPC [13, 14], providing a rigorous mathematical foundation for the flow synchronization algorithms detailed below [1, 3, 15].

2.1. Cyber-Physical Topology and Graph Model

The spatially distributed structure of the regional supply chain is formalized as a time-varying directed graph [23]:

$$G(\tau) = (V(\tau), E(\tau)) \quad (1)$$

where the active set of nodes and arcs changes [23] at each discrete step τ according to availability calendars (see Section 2.2). For notational simplicity, the time index (τ) is dropped in the following subsections where the context is clear.

This graph-based representation is essential because it captures the spatial distribution of fields, factories, and buffers, while the temporal dimension reflects the dynamic availability of nodes and arcs due to weather, machinery status, and operational decisions [18].

The node set is decomposed into five non-overlapping functional echelons:

$$V = V_S \cup V_C \cup V_B \cup V_F \cup V_{dummy} \quad (2)$$

Where:

V_S – raw material generation nodes (fields)

V_C – field buffer zones (field piles)

V_B – factory accumulation nodes (factory piles)

V_F – processing nodes (factories)

V_{dummy} – virtual sink with infinite capacity, introduced to ensure numerical feasibility of the optimization problem under force majeure conditions (see Section 2.7.1).

This separation is not merely formal: it allows the model to distinguish between stochastic generation (fields), temporary storage (buffers), deterministic consumption (factories), and a mathematical safety valve (virtual sink) [11], each with its own dynamics and constraints.

To ensure strictly directed material flow, the arc set is restricted to:

$$E \subseteq (V_S \times V_C) \cup (V_C \times V_B) \cup (V_C \times V_F) \cup (V_B \times V_F) \quad (3)$$

The virtual sink is not part of the physical transportation network; it is activated only algebraically through slack variables in the MILP formulation when force majeure constraints make full evacuation infeasible within the current planning horizon.

This restriction eliminates physically unfeasible logistical loops [11] and enforces a pull flow from fields through buffers to factories [10, 12]. The multi-echelon architecture of this directed network is illustrated in Fig. 1.




Fig.1 – Directed multigraph architecture of the logistical network
Source: Compiled by the author.

Cyber-physical integration is achieved by mapping a vector of static parameters and dynamic state variables to each graph element [8, 9]. State updates occur discretely via

transactional data transmissions [1, 15] from physical checkpoints (weighbridges, laboratories, etc.) or manual operator inputs [22].

The upper layer represents the physical material flow from stochastic resource generation (fields) through temporary buffer zones and transportation to deterministic industrial consumption illustrated in Fig. 2. Transactional measurements continuously synchronize the digital state with the physical system. The Digital Twin integrates graph-state estimation, MILP optimization, Drum–Buffer–Rope pull logic, and Model Predictive Control (MPC), producing coordinated harvesting, loading, transportation, and processing schedules.




Fig.2 – High-level architecture of the proposed cyber-physical Digital Twin.
Source: Compiled by the author.

2.2. Multi-Factor Attributive Composition of Nodes

Each network node is described by a dynamic multi-factor state matrix integrating endogenous characteristics (volume, quality) and exogenous environmental variables (temperature, soil moisture). This richness is necessary because agricultural logistics cannot be reduced to distance alone – the state of each node determines its availability, throughput, and degradation rate [6, 7].

2.2.1. Field Nodes (V_S, V_C)

The logistical availability status of a field pile is binarized through deterministic algorithmic rules:

$$Status_j(t) = \begin{cases} 0, & B_{dir} + B_{hum} + B_{frost} \geq 1 \\ 1, & \text{otherwise} \end{cases} \quad (4)$$

Where:

B_{dir} – directive blocking,

B_{hum} – critical soil moisture,

B_{frost} – critical frost.

At each step τ , nodes with $Status_j(\tau) = 0$ are removed from the active set $V(\tau)$, and arcs incident to blocked nodes are excluded from $E(\tau)$. Availability is governed by machine-specific calendars that determine whether a node is accessible for operations at a given time. Thus, the mathematical core operates on a dynamic graph $G(\tau)$, which is reconstructed daily through the MPC loop [1, 13].

This dynamic pruning is not a simplification – it reflects physical reality: fields become inaccessible after heavy rain, frost damages crops, and operators may block nodes for contractual reasons [5]. The model must adapt to these changes without human intervention.

2.2.2. Machinery Fleet

Mobile assets (harvesters, loaders, trucks, etc.) are constrained by binary availability calendars [4]. Each asset operates within a nominal active window (10–12 hours), with the remaining time serving as a guaranteed buffer for absorbing intra-day stochastic disturbances [16].

The key insight here is that machines are not infinitely flexible. Their availability is limited by shift schedules, maintenance, and travel time between locations [27]. The model accounts for this through calendar-based constraints, preventing unrealistic schedules.

2.2.3. Factory Buffer Node (V_B)

The storage yard functions as a damping buffer between stochastic logistics and deterministic processing [1, 11]. Target buffer capacity $B_{target}(t)$ is governed by seasonal risk modes: Just-in-Time (early phase), stepwise increase (transit phase), and safety stock maximization (late phase), based on objective meteorological criteria [3].

The buffer is the critical link between the stochastic world of fields and the deterministic world of the factory. Without it, any disruption in supply would immediately stop production [14]. The model dynamically adjusts buffer targets to hedge against weather risks while minimizing degradation losses [6, 17].

2.2.4. Processing Node (V_F)

Characterized by binary operational status $z_b(t) \in \{0,1\}$ and current capacity $C_{available}(t)$. During emergency shutdown ($z_b(t) = 0$), incoming flows are not blocked, allowing accumulation up to $B_{max,b}$.

This design choice is deliberate: even when the factory is down, field evacuation must continue to prevent crop loss [6]. The buffer serves as a shock absorber, allowing harvest operations to proceed independently of factory status.

2.3. Multi-Layer Data System

The Digital Twin operates with three isolated, hierarchically subordinated data layers [1, 8]. This separation is architectural, not merely organizational – it prevents data corruption and ensures that planning never overwrites physical truth [22].

- *N-layer (Nominal Reference)*: Static baseline – equipment specifications, technological standards, historical yield profiles, retrospective meteorological data. Remains unchanged during the active season [3].
- *P-layer (Planning)*: Dynamic virtual snapshot of future system state, generated by optimization algorithms over planning horizon H [1]. Contains synchronized schedules of loading complexes and transport links [4]. Data in this layer represent operational intentions and cannot overwrite physical state.
- *E-layer (Empirical Transactional)*: Absolute physical truth – discrete transactional logs from weighbridges, GPS telemetry, and quality laboratories [15, 22]. State variables are updated only upon physical execution and confirmation of transactions [1].

The E-layer is the source of truth. Plans are continuously reconciled against it, ensuring that the Digital Twin never drifts away from reality [1, 13].

2.4. Fundamental Equation of Material Balance

The global logistical cycle is governed by a chain of non-strict inequalities representing cumulative material balance:

$$Q_{total}^{biolog} \geq Q_{total}^{harvested} \geq Q_{total}^{loaded} \geq Q_{total}^{transit} \geq Q_{total}^{factory_buffer} \geq Q_{total}^{processed} \quad (5)$$

This chain reflects the physics of mass conservation [18]: each step loses some mass due to handling, storage, or degradation [6]. The inequalities are non-strict because losses are inevitable – but the model must account for them explicitly [17].

This guarantees that routing algorithms cannot allocate volumes exceeding actual physical availability at any preceding technological stage [11]. It is a hard constraint that prevents the optimizer from generating infeasible plans [19, 20].

2.5. State Update Cycle and Predictive Control

The integration of data layers is executed through a Model Predictive Control (MPC) mechanism with a recursive Sequential Forward Scheduling loop [1, 13, 14]:

1. At day t , the system retrieves physical state from E-layer for period $t - 1$ [1]. If validated data are absent, P-layer values are used as baseline [22].
2. Based on N-layer constants and E-layer state, the optimization core computes operational plans strictly for a single day t [1, 3].
3. Upon transitioning to $t + 1$, the timeline shifts: planned metrics for t are reconciled with available data, exhausted nodes are removed from the graph [23].
4. At the end of the operational day, E-layer overwrites initial conditions for the next iteration, annulling accumulated virtual projections [1, 13].

This daily recalibration is essential because agricultural forecasts are inherently uncertain [3, 5]. Yield estimates can deviate by $\pm 50\%$, weather changes unpredictably, and machinery breaks down [27]. The MPC loop ensures that the model always starts from validated facts, not from stale assumptions [1, 14].

Discrete states are managed by a two-level trigger architecture [1]:

- *Strategic (Daily) Trigger*: Global recalculation of the logistical flow matrix
- *Tactical (Operational) Trigger*: Event-driven response to localized disturbances [4].

The behavioral logic of the rolling horizon optimization cycle and the interaction of the structural layers during state updates are presented schematically in Fig. 3.




Fig. 3 – Optimization cycle (MPC loop)
Source: Compiled by the author.

2.6. Computational Pipeline and Flow Synchronization

This subsection details the sequential computational steps that translate the high-level pull signal into actionable daily quotas and synchronized transport links. The pipeline is structured as a cascading filter: starting from the raw topological graph [23], it successively prunes infeasible nodes, computes physical generation limits [27], reconciles factory demand [1], and finally identifies the dynamic bottleneck [11, 12] that determines the day's evacuation rhythm.

2.6.1. Topological Graph Pruning

Prior to solver execution, algorithmic graph filtering eliminates physically unfeasible routes [23]:

- *Global node blocking*: If $Status_j(t) = 0$, node and incident edges are removed from active set for the current day [5]. Blocked nodes are never permanently purged – they are restored when environmental constraints permit [1].
- *Selective soil bearing capacity filter*: Heavy vehicles are excluded from nodes with borderline moisture ($W_j(t) > W_{boundary}$), restricting throughput to light echelon only [5, 27].

This pruning is not optional – it prevents the solver from wasting computational resources on routes that are physically impossible due to weather or soil conditions [1].

2.6.2. Physical Supply Cap (Raw Material Generation Limit)

The global daily physical throughput capacity $\Theta_{supply,b}(t)$ for factory zone b is the strict minimum of active field echelons:

$$\Theta_{supply,b}(t) = \min(Q_{rem,b}^{biolog}(t), Cap_{harvest,b}(t), Cap_{MAUS,b}(t)) \quad (6)$$

Where:

$Cap_{harvest,b}(t)$ – is the integrated harvesting capacity of combines [27];

$Cap_{MAUS,b}(t)$ – is the total productivity of loading complexes [16].

This formula identifies the objective bottleneck at the generation stage [11] – before transport and factory constraints are even considered. If combines or loaders fail, this limit drops immediately, constraining everything downstream [1].

2.6.3. Logistical Demand Computation

The daily logistical demand is computed factoring target inventory, carryover balance, and current factory operational mode:

$$Demand_b(t) = \max(0, B_{target,b}(t) - B_b(t-1) + C_{available,b}(t) \cdot z_b(t)) \quad (7)$$

Where:

$B_{target,b}(t)$ – is dynamically adjusted based on meteorological risk hedging (3–5 day forecast horizon) [1, 3].

This is the pull signal: the factory "pulls" material from the buffers based on its current consumption and desired safety stock [11, 12]. The model anticipates weather risks and adjusts the pull intensity accordingly [14].

2.6.4. Identification of the Dynamic Drum

The final evacuation quota $\Theta_{final,b}(t)$ is established as the minimum across overlapping constraints:

$$\Theta_{final,b}(t) = \min(\Theta_{demand,b}(t), \Theta_{supply,b}(t), Cap_{Fleet,b}(t)) \quad (8)$$

Where:

$Cap_{Fleet,b}(t)$ – is the aggregate carrying capacity of the transport fleet (detailed calculation provided in Appendix C) [4, 27].

The structural logic of the Plan-Fact Substitution mechanism, which enables seamless integration between the predictive P-layer and the empirical E-layer, is presented in Fig. 4. This identifies the Dynamic Drum – the floating bottleneck that subordinates the entire logistical rhythm to the absolute weakest link at each computational interval [10, 11]. Unlike classical DBR, where the drum is a fixed machine [11, 12], here it moves between generation, transport, and demand depending on current conditions [1, 18].




Fig. 4 – Plan-Fact Substitution
Source: Compiled by the author.

2.7. Mathematical Formulation of the Logistical Flow Optimization Problem (MILP)

This subsection formalizes the core mathematical engine of the Digital Twin – the MILP model [19, 20] that translates the aggregated capacity limits and demand signals from the computational pipeline into an optimal daily evacuation plan [1]. Unlike classical VRP formulations that minimize travel distance [4], the proposed MILP core operates exclusively on aggregated macro-flows [1, 18], balancing the operational trade-off between raw material preservation [6, 17] and factory continuity [11, 14], while embedding biochemical degradation penalties from the companion PINN model [15, 16]. The deliberate exclusion of spatial transport costs preserves the compact linear structure of the active MILP core and enables deterministic operational response times in the tested industrial scenarios [19, 20].

2.7.1. Objective Function

All quality penalties and operational risks are monetized to a uniform dimension [19]. After normalizing all penalty terms to a common monetary dimension, the weighting coefficients can be set to unity under standard operating conditions $\omega_1 = \omega_2 = \omega_3 = 1$. However, these coefficients are retained as tunable control parameters: they can be amplified, diminished, or completely deactivated ($\omega_k = 0$) to reflect strategic priorities. For example, during scenario modeling, the slack penalty ω_3 may be disabled to evaluate worst-case performance, or the degradation weight ω_1 can be increased to enforce stricter quality preservation under heatwave forecasts. The starvation penalty weight ω_2 is particularly illustrative: it is dynamically adjusted in real time based on the factory buffer supply coefficient $\rho_b(t)$, automatically increasing when inventories drop below critical thresholds (see Section 2.9, Eq. 14–15). This adaptability transforms the objective function from a static cost metric into an active decision-support tool.

The global optimization criterion J over rolling horizon H minimizes [1, 19]:

$$\min_{x \in X} J = \sum_{t=1}^H (\omega_1 \cdot \text{Loss}_{\text{biochem}}(t) + \omega_2 \cdot \text{Penalty}_{\text{starve}}(t) + \omega_3 \cdot \text{Penalty}_{\text{slack}}(t)) \quad (9)$$

Where:

- x – vector of routing and allocation decision variables;
- X – feasible set defined by capacity, availability, and mass-balance constraints;
- $\text{Loss}_{\text{biochem}}$ – represents monetized sucrose degradation in field storage nodes [6, 17]. The highest evacuation priority (Smart Triage) is assigned to piles with extreme predicted degradation rates [15, 16].

$Penalty_{starve}$ – protects the continuous processing cycle against factory buffer depletion below the operational minimum threshold [1, 14].

$Penalty_{slack}$ – functions as a mathematical pressure relief valve [19], preventing algorithmic infeasibility during force majeure disruptions [1].

$\omega_1, \omega_2, \omega_3$ – weighting coefficients

Explicit mathematical formulations for these three penalty components are provided in Appendix C (Sections C.9.1–C.9.3). The slack variables (Appendix C, Sections C.9.2–C.9.3; Eqs. C10–C11) represent the volume of unmet supply or artificial mass deficit introduced when physical constraints (e.g., complete field blockage by heavy precipitation) make full evacuation impossible within the current planning horizon [19, 20]. This mechanism prevents the solver from terminating with an infeasibility error and allows it to return the maximum achievable plan.

Spatial transportation costs are excluded from the active MILP objective only after seasonal stratification and feasible route assignment have been fixed [18] (see Section 4.1). Under this assumption, each evacuated ton is assigned to a predetermined absorption zone, and the corresponding transport work becomes a quasi-invariant structural term rather than an active optimization gradient. Therefore, spatial costs are not ignored physically; they are handled during the stratification and topological filtering stages, while the rolling MILP focuses on time-sensitive trade-offs between sucrose degradation and factory starvation [6, 17].

2.7.2. Structural Constraints

Mathematical flow allocation within the MILP core is constrained by the physical throughput capacity of the cyber-physical graph elements [1, 18]. To prevent the loosening of the linear programming relaxation (LP-relaxation) and exponential growth of branch-and-bound execution times characteristic of Big-M formulations [19, 20], the architecture applies a topological pre-filtering algorithm to the logistical network [23].

2.7.2.1. Flow Conservation at Generation Nodes:

$$B_j(t) = B_j(t-1) - \sum_{b \in V_B} X_{j,b}(t) \forall j \in V_C, t \in [1, H] \quad (10)$$

This equation establishes a rigid mathematical coupling between macro-logistical flows and the sequential transition algorithm [1, 18]. The remaining biomass within a field pile at the end of the current day is formed as the carryover balance minus the aggregate volume of evacuated raw material [6, 17].

2.7.2.2. Factory Buffer Balance:

$$B_b(t) = B_b(t-1) + \sum_{j \in V_C} X_{j,b}(t) - D_b(t) \forall b \in V_B, t \in [1, H] \quad (11)$$

The remaining biomass at the factory storage yard is formed as the carryover balance from the preceding day plus all incoming transit deliveries via road transport, minus the raw material volume extracted by the pull signal onto the processing line [1, 11].

2.7.2.3. Capacity Constraints:

$$B_b(t) \leq B_{max,b} \forall b \in V_B, t \in [1, H] \tag{12}$$

$$D_b(t) \leq C_{available,b}(t) \cdot z_b(t) \forall b \in V_B, t \in [1, H] \tag{13}$$

Raw material storage is strictly limited by the capacity of the factory pile, and the processing volume is restricted by the current operating capacity of the factory, factoring in its operational status [1, 14].

These constraints (2.7.2.1 – 2.7.2.3) enforce mass conservation and physical limits [18] – you cannot store more than the buffer holds, and you cannot process more than the factory can handle [6]. They are the hard boundaries within which the optimizer operates [19, 20].

2.8. Multi-Phase Routing and Predictive Control Heuristic

The global vehicle routing problem (VRP) for a heterogeneous fleet within a stochastic agricultural environment belongs to the class of NP-hard problems [4, 18]. The architecture decouples spatial routing from temporal queue management through a two-level decomposition [1, 19] (Table 1).

Table 1. Methodological decomposition of the multi-horizon optimization architecture

Level	Operational Focus	Decision Variables	Core Computational Method
Macro (Strategic)	Field evacuation ordering, daily quotas	$X_{j,b}(t)$ – aggregate flows	MILP (Formula 9)
Micro (Tactical)	Fleet link sync, shuttle dispatching	$N_{actual,m}, T_{cycle}$	Predictive Rolling Heuristic (Appendix C)

Source: Compiled by the author

This decomposition is possible because, after topological filtering by vehicle class, road-bearing capacity, and node accessibility, the remaining fleet is treated as functionally homogeneous within each feasible shuttle link [4, 27]. Trucks are not assumed to be physically identical; rather, they are interchangeable with respect to the business function of transporting beet mass along a validated link. The macro-level decides how much biomass must move; the micro-level synchronizes the available vehicles within each validated shuttle link [1, 14].

- Phase 0: Temporal Quota Allocation (Growers' Equity).** Fields are stratified across early, normal, and late periods using a three-level hierarchy: Directive (manual), Historical Compensation, and Proportional Allocation [7]. This ensures that no farmer is systematically disadvantaged by the scheduling algorithm [2].

- **Phase I: Spatial Decomposition.** Fields are grouped into macro-clusters based on geographic proximity and MAUS assignment [4]. Daily gross plans are generated for each loader [1, 16].
- **Phase II: Micro-Routing and MPC Conflict Resolution.** Local disturbances (vehicle failure, impurity changes) trigger local recalculation of T_{cycle} without global MILP restart [1, 14]. Compensatory maneuvers include calendar extension, intra-cluster redeployment, and proportional sequestration [4, 27].

A critical operational principle embedded in this heuristic is the prohibition of service fragmentation where equipment would otherwise be inefficiently redeployed between partially cleared field piles [1, 11]. The sequential list execution (Section 2.7.2.1) ensures that each field pile is fully evacuated before the loader is reassigned, eliminating hidden switching costs that arise from repeated equipment redeployment. Furthermore, the micro-routing heuristic resolves scheduling conflicts through iterative time-window compression (Forward-Backward scheduling), where evacuation deadlines $\tau_{evac,i}$ and harvest start times $\tau_{start,i}$ are monotonically adjusted until a feasible and conflict-free schedule is achieved [1, 10].

2.9. Intelligent Control Mechanisms and Adaptive Objective Switching

The adaptive weight switching mechanism dynamically re-prioritizes optimization criteria based on factory inventory thresholds [1, 13].

The controlling parameter is the supply coefficient:

$$\rho_b(t) = \frac{B_b(t)}{C_{available,b}(t)}, \text{ if } C_{available,b}(t) > 0; \infty \text{ otherwise} \quad (14)$$

This coefficient represents how many days of operation the factory buffer can sustain at current consumption rates [1, 14]. When it drops below critical, the system must shift priorities.

The starvation penalty weight ω_2 from the objective function (Eq. 9) is dynamically adjusted as follows:

$$\omega_2(t) = \omega_2^{base} + \frac{\Delta\omega_{max}}{1 + \exp(\gamma \cdot (\rho_b(t) - \rho_{critical}))} \quad (15)$$

The sigmoid provides a continuous, differentiable transition between modes, preventing numerical oscillations that would destabilize the solver [1, 19].

Two operational states:

- **Quality-Preservation Mode** ($\rho_b(t) > \rho_{critical}$): $\omega_2(t) \approx \omega_2^{base}$, focus on minimizing $Loss_{biochem}$ (Smart Triage via PINN [15, 16]). The system evacuates the most vulnerable piles first, regardless of distance [6, 17].
- **Emergency Pull Mode** ($\rho_b(t) < \rho_{critical}$): $\omega_2(t) \rightarrow \omega_2^{base} + \Delta\omega_{max}$, overriding quality preservation to prevent factory starvation [11, 14]. Every available vehicle is redirected to ensure the factory does not stop [1].

The cyber-physical resilience mechanism, which coordinates exogenous data acquisition, the predictive PINN, and the MPC loop, is illustrated in Fig. 5.




Fig. 5 – Cross-functional architecture of cyber-physical resilience
Source: Compiled by the author.

This mechanism transforms the model from a passive optimizer into an active risk manager, capable of balancing long-term quality goals against short-term survival constraints [1, 12, 21].

2.10. Transactional Telemetry and Microservice Architecture (Concise)

The Digital Twin is deployed as an event-driven microservice architecture (EDA) [1, 22]. The backend is implemented in Python with FastAPI, PuLP/HiGHS for MILP [19, 20], and containerized deployment on Debian. Data streams from weighbridges, GPS telemetry, and laboratories are aggregated through a message bus (RabbitMQ) into an Azure MS SQL Data Lake [1, 22]. A Power BI layer provides real-time visualization and operational analytics [26, 29]. The operator interface (Control Room UI) implements a Human-in-the-Loop paradigm with directive override capabilities [1, 32].

This architecture ensures scalability, resilience, and real-time responsiveness – critical properties for a system that must operate under peak seasonal loads without downtime [8, 9, 15].

3. Results

This section presents the empirical validation results of the developed cyber-physical system and its predictive control algorithms, specifically focusing on the flow synchronization and forward-backward scheduling frameworks. Within the contemporary smart regional development paradigm, the efficacy of digital platforms is evaluated by their capacity to ensure systemic resilience, high-quality operational governance, and structural optimization under stochastic disturbances [2]. The evaluation of the developed architecture was conducted using real-world agro-industrial data from an extended business case study [15]. The empirical results demonstrate that integrating Digital Twin technologies into regional logistical operations effectively prevents systemic capacity blockages during operational shocks [3].

Since the computational architecture is explicitly decoupled, this section focuses on the empirical validation of the macro-logistical routing and fleet synchronization metrics. Detailed empirical results demonstrating the operation of the adaptive objective function weights specifically the prevention of factory starvation and the minimization of biochemical degradation via PINN are comprehensively analyzed in the second part of this study [15].

3.1. Empirical Research Base and Macro-Logistical Network Topology

To verify the proposed multi-layer Digital Twin architecture, a representative dataset was compiled from the transactional logs of the empirical layer (E-layer) of an agro-industrial enterprise (Südzucker Moldova SRL) in the Republic of Moldova [22]. This regional context aligns with contemporary smart regional development discourse regarding the minimization of digital divides between urbanized industrial centers and agrarian territories [2], alongside the implementation of resource allocation optimization within developing institutional environments [18]. The retrospective analysis spans an 11-year operational period (harvesting campaigns from 2015 to 2025) and comprises an array of more than 180,000 unique heavy-duty logistical cycles [9].

The physical topology of the macro-logistical network is characterized by high combinatorial dimensionality and inter-seasonal spatial volatility [23]. The logistical supply graph undergoes reconfiguration annually. During the simultaneous operation of two processing facilities (2015–2017), the number of active raw material generation nodes reached peak values of up to 337 unique fields per season, requiring a heterogeneous fleet of up to 252 vehicles algorithmically allocated based on road bearing capacities [4]. The aggregated characteristics of the logistical problem scales solved annually by the optimization core are presented in Table B1 (Appendix B).

Starting in 2020, the enterprise's processing operations were completely consolidated into a single absorption node (the Drochia factory, DR), resulting in the decommissioning of the second processing plant. Analysis of transport flow transactional density refutes the

hypothesis that the recorded contraction in fleet size (Table B1 Appendix B) is a consequence of declining production volumes.

Historical logs show that during the peak year of 2016, an extensive fleet of 252 vehicles served two factories that collectively processed 23,003 deliveries (averaging 91 trips per vehicle per season). In 2021, following capacity consolidation, the logistical network contracted spatially, yet the total volume of transport remained elevated, with the single DR facility processing 19,323 deliveries. Consequently, while the total transactional workload decreased by 16%, the inbound traffic density focusing on the sole remaining facility increased by 68% compared to the historical average workload per factory during peak seasons (11,501 deliveries per plant).

Under hyper-concentrated traffic conditions, classical push logistics would be expected to generate queue growth and operational paralysis. After plant consolidation and the transition to the pull-based Digital Twin architecture, the deployed fleet decreased from 252 vehicles in 2016 to 88 vehicles in 2021, while the remaining factory absorbed substantially higher inbound traffic density. Although this reduction cannot be attributed to the algorithm alone, the simultaneous compression of queue times, fleet dispersion, and vehicle turnover indicates that the proposed Forward-Backward Scheduling mechanism materially contributed to resource consolidation and sustainable regional operations [1, 3, 7, 27].

The statistics of the terminal 2025 macro-season present a specific operational scenario where, despite an all-time low of active loading complexes (5 MAUS units), the predictive algorithm deployed an expanded fleet of 131 vehicles [16]. The nominal capacity of a single MAUS loader is approximately 550 t/h. To cover the targeted daily factory demand (4,000–4,200 t/day), the operation of a single unit is sufficient, whereas historically the loader fleet operated under capacity underutilization.

The driver behind the fleet expansion in 2025 was the onset of a prolonged period of adverse precipitation. In response to the impending meteorological shock and the risk of losing soil bearing capacity on unpaved roads, the predictive control loop initiated a forced raw material accumulation mode [1, 13] (see Section 2.6.3). To maximize the utilization of narrow favorable weather windows, the system operated the 5 available MAUS complexes at their peak design capacities [27]. To prevent downtime for these loaders, the carrying capacity of the shuttle transport links was increased, leading to the recruitment of 131 heavy-duty trucks [4, 16]. This case illustrates the flexibility of the Digital Twin framework in eliminating equipment underutilization and dynamically scaling the transport echelon to adapt to stochastic weather disturbances, thereby supporting macro-logistical network resilience [3, 8].

3.2. Computational Performance and the Forward-Backward Scheduling Algorithm

Deploying Model Predictive Control (MPC) systems within a real-time management loop requires immediate computational responses to exogenous disturbances [1, 13]. Because the routing mathematical model incorporating biochemical degradation belongs to the class of NP-hard problems, exact solvers encounter severe combinatorial scaling barriers. When applied to industrial parameters, commercial MILP solvers (e.g., IBM CPLEX) exhaust default computational limits (Time Limit = 3,600 s) while exhibiting optimality gaps of up to 4.53% [1]. To overcome this barrier, the proposed multi-phase heuristic (see Section 2.8) bypasses global, simultaneous matrix recalculations. Instead, the computational core relies on Sequential Forward-Backward Scheduling integrated with localized constraint reduction rules [1, 10] (see Sections 2.6.1 and 2.7.2).

To evaluate the computational efficiency and scalability of this architecture, a stress-test scenario was constructed to simulate a peak operational day within the regional supply chain [4]. Without preliminary topological filtering, the active search space comprises 10 mobile loaders (MAUS), 350 harvest-ready fields, 2 processing factories, and a heterogeneous fleet of 250 vehicles [27]. Because the allocation of the transport fleet into dedicated links is computed heuristically prior to MILP execution (see Section 2.6.4), vehicle assignment variables are algorithmically decoupled from the solver [4]. This configuration yields a compact mathematical matrix. Due to the multi-phase heuristic, which decouples vehicle routing from flow balancing, the MILP core operates exclusively on aggregated daily macro-flows [1, 19]. For a peak operational day with 350 active fields, the matrix contains approximately 350 continuous flow variables ($X_{j,b}(t)$) and 360 structural constraints [25]. This dimensionality enables the open-source HiGHS solver, integrated via the PuLP modeling library, to converge to an optimal feasible solution in 0.45 seconds [19]. The combinatorial complexity of micro-routing is entirely absorbed by the Sequential Forward-Backward Scheduling heuristic (Section 2.8), which operates on dedicated shuttle links without burdening the MILP core [1, 10].

When executed within a single-threaded Python environment on a standard laptop (Intel Core i7, 11th Gen, 4 cores, 8 Logical Processors @ 3.30GHz, 32 GB DDR4 RAM), the open-source HiGHS solver integrated via the PuLP modeling library converges rapidly. Matrix generation requires approximately 0.07 s, and complete solver execution (including presolve and branch-and-bound phases) successfully converges to an optimal feasible solution in 0.45 s [19]. Consequently, tactical optimization for a single operational day requires roughly 0.5 s. Extrapolating this performance to a full-season strategic projection (100 days) yields a convergence time of under 60 s. To ensure methodological transparency, the single-threaded benchmark script and synthetic dataset are publicly accessible.

The computation of raw material biochemical degradation parameters ($\Delta S_j(t)$) via the external PINN architecture is detailed in the companion article [15]; these predictive quality metrics are ingested by the optimization core as exogenous inputs [16]. In the full

production pipeline, the execution phase encompassing neural inference, multi-horizon MILP simulation, and asynchronous transactional synchronization with a cloud-based Azure MS SQL Data Lake under peak seasonal workloads reliably terminates within a maximum 15-minute window [1, 22]. This benchmark confirms the technological resilience and reaction speed of the proposed architecture, supporting its practical readiness for integration into smart regional management loops under macro-systemic shocks [1, 3, 29].

3.3. Formation of Shuttle Transport Links and Fleet Consolidation

To provide an empirical validation of the developed multi-phase heuristic and the transport link allocation algorithm (see Section 2.6.4), a retrospective analysis was performed on transactional logs (exceeding 180,000 database records) for the period spanning 2015 to 2025 [5, 22]. The core evaluation metric chosen to track the migration from a chaotic push logistics structure to a synchronized pull framework was the daily fleet dispersion per loading asset, defined as the average number of unique cargo trucks serving a single mobile MAUS loading complex within a given shift [27]. The historical dataset is presented in Table B2 (Appendix B), and its corresponding graphical interpretation is shown in Fig. 6.




Fig. 6 – Boxplot analysis of daily fleet dispersion
Source: Compiled by the author based on empirical transactional logs of S dzucker Moldova SRL

During the 2015–2018 historical baseline period (prior to the integration of the specialized routing algorithms), the logistical data was characterized by a massive Interquartile Range (IQR) and a significant concentration of extreme outliers [5]. The maximum daily fleet dispersion regularly reached 23.4 to 25.1 unique transport units assigned to a single MAUS loader (as documented at the F le ti processing plant, FL). This equipment allocation and broad statistical variance highlighted the systemic inefficiencies of push dispatching: individual trucks completed a limited number of trips per day, causing structural queues at both the field pile loading areas and the factory buffer yards [4, 11].

The initial iterations of the mathematical core and its companion shuttle link algorithm were introduced into active operations between 2019 and 2020 [1]. The empirical findings captured in the boxplot visualization (Fig. 6) reflect this technological pivot: by the 2020 season, the median values and overall dispersion of even the most heavily taxed MAUS loader compressed sharply, capping the absolute maximum at 12.6 unique trucks per day.

Over the subsequent operational cycles (2021–2024), the DBR pull loop tightly constrained the size of the allocated transport link by anchoring it directly to the predictive logical cycle times [10, 11]. Consequently, during the 2024 harvesting campaign, the highest-intensity loader was sustained by a lean transport pool showing minimal statistical deviation (peaking at 13.6 trucks per day). This contraction of fleet dispersion by nearly 41%, accomplished without sacrificing targeted delivery volumes, provides strong empirical evidence of the conversion of a fragmented transport fleet into closed, high-density shuttle links with optimized round-trip times [1, 27]. This algorithmic resource consolidation directly mitigates the carbon footprint associated with idle vehicle delays, advancing the efficiency of regional logistics networks [2, 26].

3.4. Logistical Cycle Time Optimization and Fleet Minimization

The operational efficiency of shuttle transport links depends directly on the stability of the total transactional logistical cycle time (T_{cycle}) [4, 27]. In classical push-based agro-logistics management systems, the most volatile and stochastic component of this cycle is the vehicle waiting time in the queue at the factory reception yard (t_{queue}). Under chaotic raw material arrival patterns, the duration of a heavy-duty truck's stay at the reception yard varies on average from 20 to 55 minutes, which artificially bloats the required fleet size [5, 25]. The empirical dynamics of the average truck processing time at reception yards over an 11-year period are presented in Table B3 (Appendix B) and illustrated in Fig. 7.




Fig. 7 – Dynamics of average truck processing time
Source: Compiled by the author based on empirical transactional logs of Süducker Moldova SRL

The historical dataset covering more than 180,000 unique heavy-duty truck transactions demonstrates the efficiency of implementing a Digital Twin-driven pull architecture based on the Drum-Buffer-Rope framework [6, 10, 11]. From 2020 onward, following the complete consolidation of raw material processing into a single absorption node (the Drochia factory – DR), classical queuing theory would predict an exponential growth of queues due to a two-fold concentration of the inbound vehicle stream [5].

However, the synergy of predictive arrival slot management and extensive hardware and software modernization of the throughput gate layout prevented a logistical collapse [1, 22]. Consequently, the average vehicle turnaround time at the reception yard dropped to an all-time low of 13.8 minutes during the 2024 season.

This technological integration eliminated manual buffers at key checkpoints of the logistical graph [22]. The gross (M_{brutto}) and tare (M_{tara}) check-weighing procedures were fully automated: contactless reading of RFID transponders was integrated directly with the weighbridge terminals, instantly generating deterministic JSON transactions into the Digital Twin's data layer [24, 30]. The registration time at these checkpoints was reduced to a physical floor of 40–60 seconds.

Following the transition from barcodes to RFID in 2015, the registration time at the weighbridge terminals stabilized near the physical minimum and could no longer serve as the primary source of further cycle-time reduction [22]. The subsequent decline in average reception-yard turnaround time from 24.6 minutes in 2015 to 13.8 minutes in 2024 (Table B3, Appendix B) is therefore primarily attributable to algorithmic synchronization of arrivals and queue compression, rather than to further reductions in checkpoint processing time [1, 10, 11].

4. Discussion

The empirical deployment of the proposed Digital Twin architecture shifts regional agro-industrial supply chain management from a reactive, distance-centric dispatching paradigm to a continuous, predictive cyber-physical synchronization loop [1, 5]. The resulting operational insights provide alternative perspectives on fundamental logistical challenges traditionally explored in operations research, while also revealing the practical boundaries of the proposed approach.

4.1. Paradigm Shift: Temporal Prioritization vs. Spatial Distance Minimization

A core theoretical insight of this study is the deliberate rejection of classical Vehicle Routing Problem (VRP) frameworks for biomass evacuation scheduling [4, 18]. Traditional VRP models assume static customer demand and instant product availability, treating fields merely as coordinate points on a geographic map where spatial distance minimization dictates optimality [4]. As demonstrated in Higgins [25] for sugarcane transport, classical VRP formulations encounter severe operational limits because they treat distributed agricultural hubs as static entities. In real-world sugar beet logistics, this assumption completely collapses due to severe technological asymmetries [6, 17].

The duration of the in-field harvesting operation exceeds the active transport loading duration by a significant factor, transforming the distributed fields into dynamic, time-degrading inventory buffers [6, 27]. Under such conditions, minimizing spatial travel distance becomes secondary to managing temporal buffer degradation [17]. By utilizing an elastic intra-day vehicle time budget instead of continuous spatial graph restructuring, the proposed multi-phase heuristic treats each field clamp not as a point on a map, but as a perishable resource with non-linear biochemical decay [1, 16].

The computational core successfully prioritizes temporal synchronization matching the factory's continuous pull signal with the fields' degradation rate (Smart Triage) allowing the system to algorithmically absorb localized equipment breakdowns or transit delays without triggering resource-intensive global matrix recalculations [1, 10].

Comparison with existing approaches. The most comprehensive prior work on sugar beet harvest planning is presented by Baunach, Spinler, and Birge [31], which proposes an integrated multi-commodity flow model on a time-expanded network [23]. While their approach is commendable for its system-wide scope, it rests on two assumptions that fundamentally limit its applicability to real-time operations: (i) sugar beets are treated as non-perishable, eliminating the core degradation-driven synchronization problem [6]; and (ii) the model is deterministic and integrated, which makes it brittle under the high-frequency stochastic shocks (weather, machinery failures, yield variability) that characterize actual campaigns [1, 3].

Distinctive contributions of this work. In contrast, the presented framework:

- Explicitly models non-linear biochemical degradation via a companion PINN architecture [15, 16].
- Introduces a dynamic DBR mechanism for shifting bottlenecks, where the "drum" floats between generation, transport, and demand [10, 11].
- Decouples macro-level quota validation from micro-level execution through a two-level heuristic, enabling real-time cyber-physical synchronization under uncertainty [1, 19].
- Treats spatial transport work as a quasi-invariant structural term after seasonal stratification and feasible route assignment, allowing the active MILP objective to focus on time-sensitive degradation and starvation risks [18].

This combination of features is, to the best of author knowledge, absent from existing agricultural logistics models, which typically treat transport as a primary cost driver rather than a structural constant [4, 25].

4.2. Overcoming "Schedule Nervousness" via Rolling Horizon Simulation

Operational testing confirms that under heavy agro-industrial stochasticity (continuous yield refinements, sudden precipitation, machinery downtime), rigid long-term plans lose operational validity within a 24-hour window [3, 5]. Applying classical global optimization in such unstable environments triggers severe schedule nervousness (Scheduling Nervousness): minor deviations force the mathematical core to completely reconstruct the entire equipment timeline, causing vehicle arrival windows at the fields to shift erratically along the time axis [1].

To resolve this issue, the Digital Twin architecture replaces single-pass global planning with an iterative, daily rolling horizon computation [1, 13]. The core innovation lies in applying an identical algorithm for both generating predictive plans and recording historical transactions, converting the framework core into a step-by-step timeline simulator [1].

This methodological approach creates a single mathematical continuum, processing both historical actuals and forward-looking plans on an identical graph topology [23, 25]. This algorithmic homogenization allows operators to leverage a unified analytical toolkit: the exact same matrix is utilized for constructing retrospective plan-to-actual reports and for executing multi-scenario stress testing of the logistics network [1, 9]. The system's capacity for predictive shock modeling establishes a vital infrastructural foundation for ensuring cyber-physical resilience and executing smart regional management strategies (Smart Governance) [1, 3, 29].

Practical implications. The rolling horizon paradigm proved critical during the 2025 season, when a sudden precipitation event triggered a forced accumulation mode [1, 13]. Rather than recalculating the entire seasonal plan, the system locally adjusted quotas and fleet allocations within a single daily cycle [5, 27], demonstrating that the architecture's resilience is not a theoretical property but a demonstrated operational capability.

4.3. Cyber-Physical Resilience and the Role of Disintermediation

A fundamental, yet rarely discussed, vulnerability of decentralized supply chains is the opportunistic behavior of independent contractors (the Principal-Agent Problem) [2, 21]. Analysis of the baseline scenario revealed that prior to the integration of the Digital Twin model, managing a single logistical cycle required the synchronous participation of up to 11 distinct actors (7 internal factory operators and 4 external network representatives). This excessive concentration of the human element (Human-in-the-Loop) exponentially slowed down checkpoint throughput and created a breeding ground for opportunistic collusion [3, 22].

The Digital Twin addresses this through two complementary mechanisms:

- **Hardware-level disintermediation.** The architecture eliminates manual buffers at key checkpoints through RFID integration and automated weighbridges [24, 30]. Any

transaction lacking a continuous, hardware-validated timestamp chain (harvesting start → MAUS loading → gross weight terminal → laboratory analysis → factory buffer yard → tare weight terminal) is automatically discarded [3, 22]. This reduces administration overhead by more than 80% and ensures that contract closures and transport billing are executed automatically based on validated E-layer logs [1, 15].

- **Data-level triangulation.** The system continuously evaluates the variance of an individual vehicle's empty weight, enforcing a rigid physical boundary where tare fluctuations cannot exceed 1,000 kg [5, 22]. Transactions exceeding this empirical corridor are flagged as weighbridge positioning anomalies, securing absolute integrity for the net weight (M_{neto}) parameter used in the MPC solver [15]. Furthermore, cross-referencing departure and arrival timestamps with declared distances allows the system to detect fraudulent source-field declarations. Any attempt to declare a more distant field forces the truck's required average speed past physical limits (exceeding 60–100 km/h), triggering an automatic audit flag [4, 27].
- **Institutional implications.** By eliminating information asymmetry, the Digital Twin transforms institutional relationships into a substantially more transparent, paperless cyber-physical space of electronic governance (E-Governance) [2, 21, 22]. This establishes a robust foundation for institutional trust, which is particularly critical in developing regional contexts where contract enforcement mechanisms may be weak [2].

4.4. *Balancing Economic Efficiency and Growers' Equity*

The implementation of multi-year quota balancing (see Section 2.8.1, Phase 0) provides empirical evidence that algorithmically safeguarding growers' equity (Growers' Equity) does not conflict with global supply chain efficiency [7]. Historical compensation algorithmically shields smallholder farms from logistical discrimination during high-risk early and late harvesting windows [2]. This fully aligns with the principles of inclusive "smart governance" (Smart Governance) and sustainable regional development actively championed within contemporary research published in Smart Cities and Regional Development (SCRD) [2, 21, 29].

A critical nuance. The model does not pursue the unachievable goal of absolute equality among suppliers [7]. Instead, it provides a formal mechanism for preventing systematic discrimination: no single agricultural producer is repeatedly forced to bear disproportionate losses due to the calendar allocation of harvesting windows across consecutive seasons [2]. Crucially, the algorithm does not sacrifice global supply chain efficiency, as social balancing is isolated within the preliminary quota allocation phase and does not interfere with real-time operational route optimization [1, 4]. Consequently, the system minimizes aggregate logistical expenditures rather than merely redistributing them from one group of stakeholders to another.

4.5. *Research Limitations*

Despite the demonstrated efficiency and high performance of the proposed Digital Twin architecture, this study is subject to several objective limitations:

- **Meteorological Dependency.** The predictive accuracy of the adaptive weight switching mechanism (Objective Switching) and preventive quota sequestration is highly dependent on the reliability of exogenous weather forecasts over the planning horizon $H_{forecast}$ [1, 3]. While rolling horizon recalibration mitigates this dependency, extreme weather events with lead times shorter than the forecast window remain a source of residual uncertainty [13].
- **Telemetry Ingestion Limits.** A fundamental barrier to absolute algorithmic autonomy is the current limitation in collecting real-time field telemetry [5, 22]. As noted in the companion study [15], predictive biochemistry modeling depends on continuous field data stream inputs; transitioning to automated digital control is constrained by regional infrastructure and deployment capitalization costs across remote locations [2, 32]. The current implementation relies on manual temperature probing, which introduces both latency and sampling bias [6].
- **PINN Integration Status.** The Physics-Informed Neural Network (PINN) module is currently in a pilot validation phase (Shadow Mode) and has been tested predominantly on retrospective historical datasets [15, 16]. Full-scale deployment within a production environment requires further calibration of the network's hyperparameters using expanded multi-season datasets [15]. This is essential to improve operational robustness of predictions under extreme climatic deviations and to achieve seamless integration of the neural network outputs into the transactional loop of the corporate ERP system [22].
- **Computational Scaling Gaps.** The computational benchmark presented in Section 3.2 reflects an isolated simulation environment [19]. While the open-source HiGHS solver converged in 0.45 s on the synthetic stress-test scenario, full-scale industrial deployment with a complete transactional dataset may introduce latency due to database I/O and network communication overhead [1, 22]. This constraint is partially addressed by isolation protocols between the optimization core and the transactional layer, but remains a subject of ongoing performance tuning.
- **Generalizability.** The empirical validation is based on a single regional context (Republic of Moldova) and a specific crop (sugar beet) [2, 9]. While the mathematical framework is designed to be universal [18, 23], the quantitative performance metrics (e.g., queue time reduction, fleet consolidation) may vary across different agricultural systems, transportation infrastructures, and regulatory environments [2, 27]. Further validation across diverse contexts is required to establish the full generalizability of the approach.

4.6. Implications for Future Research and Practice

The findings of this study open several avenues for future investigation:

- **Integration with Reinforcement Learning.** The current MILP core relies on manually calibrated weights for the adaptive objective function [1, 19]. Future work will explore coupling the developed framework with deep Reinforcement Learning algorithms to automate hyperparameter tuning under escalating regional climatic instability, enabling the system to learn optimal switching policies from historical data [15].

- **Full-scale PINN Deployment.** The next phase of this research involves deploying the PINN module in a live production environment, replacing the current shadow-mode validation with real-time integration [15, 16]. This will require addressing the telemetry infrastructure gap through the deployment of IoT temperature sensors across field piles [5, 32].
- **Multi-Product and Multi-Site Extension.** While the current model focuses on a single product (sugar beet) within a regional cluster, the mathematical framework is agnostic to the specific commodity [6, 18]. Extending the architecture to manage multiple products, multiple processing sites, and inter-regional coordination represents a natural next step [28].
- **Policy and Governance Implications.** Beyond technical refinement, the Digital Twin architecture offers a governance tool for regional authorities [21, 22]. By providing transparent, auditable data on supply chain performance and stakeholder equity, the framework can support evidence-based policymaking in agricultural development and food security [2, 29].

5. Conclusion

This study presents the primary architectural and mathematical formalization of a comprehensive Digital Twin framework designed to ensure the cyber-physical resilience of regional agro-industrial supply chains [3, 5]. The developed optimization framework translates extensive empirical transactional records into a scientifically grounded model of smart operational governance and regional resource allocation [1, 17].

The research yields the following key conclusions:

5.1. *Logistical Pull Synchronization*

Replacing continuous spatial micromanagement of heterogeneous fleets with discrete transactional state updates (E-layer) and stabilized shuttle transport links significantly enhances logistical network resilience [9, 16]. Integrating a predictive pull framework based on Drum-Buffer-Rope principles [10, 11] with a compact Mixed-Integer Linear Programming (MILP) core [19] enables the system to algorithmically smooth traffic density peaks and minimize vehicle processing queues at factory reception terminals without reducing total throughput volumes [4, 27].

Empirical validation: The 11-year industrial dataset (180,000+ transactional cycles) demonstrates a 41% reduction in active fleet dispersion and a substantial consolidation of deployed fleet resources after plant consolidation and the transition to the pull-based Digital Twin architecture. The fleet decreased from 252 vehicles in 2016 to 88 vehicles in 2021 while the remaining factory absorbed substantially higher inbound traffic density, indicating that the synchronization mechanism materially contributed to resource consolidation and queue compression [5, 27].

5.2. Rolling Horizon Control and Priority Shifts

Empirical validation confirms that due to high stochastic variance in agricultural logistics, static seasonal schedules lose operational validity within short intervals [3, 5]. The deployed Model Predictive Control (MPC) framework addresses this via recursive, daily recalculations over a rolling horizon, maintaining identical graph topologies for historical actuals and predictive plans [1, 13].

Furthermore, integrating a piece-wise supply coefficient with a smooth sigmoidal weight switching mechanism enables the solver to autonomously shift objective functions between quality preservation and factory starvation avoidance, safeguarding industrial process continuity under severe environmental disruptions [1, 9]. This adaptive mechanism was successfully stress-tested during the 2025 season, when a prolonged precipitation event triggered a forced accumulation mode, and the system locally adjusted quotas and fleet allocations without recalculating the entire seasonal plan [1, 13].

5.3. Dynamic Bottleneck Identification

The introduction of the Dynamic Bottleneck Principle formalizes the floating constraint that subordinates the entire logistical rhythm to the weakest link at each computational interval:

$$\Theta(\tau) = \min(A(\tau), B(\tau), D(\tau), \Gamma_{load}(\tau)) \quad (16)$$

where:

- $A(\tau)$ – available resource (supply layer) [6, 17],
- $B(\tau)$ – transport capacity (infrastructure layer) [4, 27],
- $D(\tau)$ – unsatisfied demand (consumption layer) [14, 28],
- $\Gamma_{load}(\tau)$ – loading agent capacity (equipment layer) [10, 16].

This constitutes a significant extension of classical Theory of Constraints (TOC) thinking, where the "drum" is typically a fixed mechanical asset [11, 12]. In the proposed framework, the bottleneck dynamically shifts between four independent layers depending on current conditions [1, 18].

5.4. Biochemical Degradation Management

While the full-scale deployment of Physics-Informed Neural Networks (PINN) for real-time degradation prediction remains a subject of ongoing work [15], the conceptual integration of thermodynamic laws governing sucrose decay into a monetized logistical penalty establishes a rigorous foundation for intelligent prioritization (Smart Triage) [16, 17]. The retrospective analysis reveals that up to 15–18% of total seasonal biomass is traditionally processed during the final 4–5 weeks of the campaign, within which average unmonitored loss of digestion reaches 11–15% of baseline potential. The proposed architecture is designed to algorithmically shave off this "tail" of the degradation curve, transforming the logistics framework from a cost center into an instrument for safeguarding enterprise profitability [1, 17].

5.5. Socio-Economic Justice (Growers' Equity)

The implementation of multi-year quota balancing (Phase 0) provides empirical evidence that institutional fairness can be protected without conflicting with global supply chain efficiency [7]. Integrating a temporal quota allocation module with historical compensation provides a formal mechanism for preventing repeated logistical discrimination against independent agricultural producers, supporting the principles of inclusive sustainable regional development [2, 21]. The three-level hierarchy (Directive, Historical Compensation, Proportional Allocation) reduces the risk that any single producer repeatedly bears disproportionate losses due to calendar allocation of harvesting windows, while the social balancing mechanism remains isolated from real-time operational optimization [1, 4].

5.6. Cyber-Physical Disintermediation

The Digital Twin architecture eliminates manual buffers at key checkpoints through RFID integration and automated weighbridges, reducing administrative overhead by more than 80% [24, 30]. Hardware-validated timestamp chains ensure that transactions lacking continuous validation are automatically discarded. Data triangulation mechanisms—including empty weight variance analysis and speed-to-distance consistency checks—detect fraudulent activities [27, 30], establishing a transparent, paperless cyber-physical governance space [21, 22, 23].

5.7. Research Implications and Future Directions

The proposed cyber-physical architecture serves as a scalable methodological baseline for managing decentralized logistical networks handling perishable biomass [3, 8]. Future research trajectories will focus on:

- **Reinforcement Learning Integration.** Coupling the developed MILP core with deep Reinforcement Learning algorithms to automate objective function hyperparameter tuning under escalating regional climatic instability [1, 19].
- **Full-Scale PINN Deployment.** Transitioning the pilot PINN module from shadow-mode validation to live production integration, requiring IoT sensor deployment across field piles for real-time temperature profiling [15, 32].
- **Multi-Product and Multi-Site Extension.** Extending the architecture to manage multiple commodities, multiple processing sites, and inter-regional coordination, leveraging the framework's commodity-agnostic mathematical core [18, 28].
- **Policy and Governance Applications.** Beyond technical refinement, the Digital Twin architecture offers a governance tool for regional authorities, providing transparent, auditable data on supply chain performance and stakeholder equity to support evidence-based policymaking in agricultural development and food security [2, 21, 29].

5.8. Closing Reflection

The scientific contribution of this work lies not in any single algorithmic component, but in the original integration of:

- Hierarchical graph stratification [23],
- Temporal decoupling of spatial and temporal domains [1, 18],
- Four-layer Dynamic Bottleneck Principle [10, 11],
- Quasi-invariance of spatial work with explicit loss sensitivity [6, 18],
- Forward-Backward scheduling as a formal Rope mechanism [1, 10],
- Sigmoid-based adaptive objective switching [1, 13],
- MPC with granular resilience and practical stability [1, 14],

into a unified cyber-physical framework for Total Evacuation Problems [1, 8].

To the best of the author's knowledge, this is one of the first formalizations that transforms agricultural logistics from a rigid cost center into a resilient cyber-physical synchronization system without compromising stakeholder equity or systemic transparency [2, 21].

The framework is potentially transferable to a broader class of problems where stochastic generation must be synchronized with deterministic consumption under limited throughput capacity and changing network topology [18, 23], including bulk-material logistics, emergency evacuation planning, and distributed resource coordination [1, 28]. Further validation is required before claiming full generalizability across these domains.

Acknowledgements.

The author expresses profound gratitude to the operational, logistics, and agronomic departments of Sdzucker Moldova SRL and Agro-SZM SRL for their long-term cooperation in testing and successfully implementing the developed Digital Twin architecture into a real production environment.

Conflict of Interest / Author Declaration. The author of this study, Corneliu Iatco, is currently employed as the Head of the IT Department at Sdzucker Moldova SRL. This professional affiliation may be perceived as a potential conflict of interest and is therefore explicitly disclosed. The author declares that this research is an applied industrial case study based on mathematical models and algorithms originally developed by the author and integrated into daily operations at Sdzucker Moldova SRL. The company provided access to infrastructure and historical data for retrospective analysis but had no role in the theoretical formulation, algorithmic design, interpretation of the published results, or preparation of the manuscript.

Data Availability & Intellectual Property. The mathematical models, logic, and predictive control algorithms (including the MPC frameworks and the PINN architecture) described in this study are the sole intellectual property and independent development of the author. The 11-year empirical dataset and the specific corporate hardware-software implementation of the system represent the proprietary commercial assets of Sdzucker Moldova SRL. Due to strict corporate confidentiality policies (NDA), raw historical datasets are not publicly available.

However, to ensure methodological reproducibility and transparency of the computational logic, the core computational algorithms (the MILP core and the companion PINN training module) have been isolated into a standalone Python Proof-of-Concept (PoC). These scripts, along with synthetic datasets that reproduce the mathematical topology and statistical structure of the real-world regional supply chain, are made publicly accessible via the following GitHub repository: <https://github.com/iatco13/sbdt>.

Funding. The theoretical foundation of this research was carried out on a personal initiative and received no external grant funding. The practical implementation of the author's algorithms into the production environment was funded by the internal budgets of Südzucker Moldova SRL as part of its planned operational digitalization strategy.

AI Use Disclosure. The original manuscript was drafted in Russian. In preparing the final English version of the manuscript, the author utilized Artificial Intelligence tools (specifically, Google Gemini Advanced and Google NotebookLM) solely as auxiliary instruments for translation, academic language editing, and formatting. The author exercised full human oversight over this process and assumes sole responsibility for the accuracy, originality, and scientific integrity of the final text. In accordance with current publication ethics, AI tools are not listed as co-authors.

Appendix A. Mathematical Notation and Nomenclature

Table A1. Sets and Indices

Symbol	Definition
V	Complete set of heterogeneous network nodes ($V = V_S \cup V_C \cup V_B \cup V_F \cup V_{dummy}$)
E	Set of directed transportation arcs (edges)
	V_S Subset of raw material generation nodes (fields)
	V_C Subset of field buffer zones (field piles)
	V_B Subset of factory accumulation nodes (factory piles)
	V_F Subset of processing nodes (factories)
	V_{dummy} Virtual sink node introduced for numerical feasibility under force majeure conditions
j	Index for a specific field pile node ($j \in V_C$)
b	Index for a specific factory buffer node ($b \in V_B$)
m	Index for a mobile cleaner-loader complex (MAUS)
t	Discrete time step index within the rolling planning horizon ($t \in [1, H]$)

Source: Compiled by the author

Table A2. Parameters (Constants & Exogenous Inputs)

Symbol	Definition
$B_{max,b}$	Maximum physical capacity of the factory buffer b (t)
$B_{target,b}(t)$	Dynamic target climatic inventory in the factory buffer (t)
$C_{available,b}(t)$	Current operating processing capacity of factory b (t/day)
$\Delta S_j(t)$	Predictive sucrose loss coefficient in field pile j, computed via PINN (fraction)
η	Standardized technological sugar extraction efficiency coefficient (fraction)
P_{sugar}	Market value of one ton of white sugar (monetary units)
P_{starve}	Specific penalty for raw material undersupply (monetary units/t)
P_{slack}	Prohibitive penalty for utilizing artificial slack variables
$\omega_1, \omega_2(t), \omega_3$	Structural weights in the objective function (static and dynamically adaptive)
T_{cycle}	Total transactional round-trip time of a transport unit (hours)
$z_b(t)$	Binary operational status indicator of factory b (1 = normal, 0 = shutdown)

Source: Compiled by the author

Table A3. Decision and State Variables (MILP Core)

Symbol	Definition
$X_{j,b}(t)$	Continuous macro-flow variable: volume dispatched from j to b (t)
$B_j(t)$	Physical volume of biomass remaining in field pile j at time t (t)
$B_b(t)$	Physical volume of biomass in the factory buffer b at time t (t)
$D_b(t)$	Actual volume of raw material pulled into processing at factory b (t)
$X_{short_fac,b}(t)$	Continuous slack variable representing raw material deficit for factory b (t)
$X_{short,j}(t)$	Continuous slack variable representing artificial mass deficit at node j (t)
$N_{actual,m}(t)$	Integer variable: actual number of vehicles allocated to the MAUS link m

Source: Compiled by the author

Appendix B: Empirical Data and Macro-Logistical Statistics

Table B1. Macro-statistics of supply chain topology and logistical resource allocation (2015–2025)

Macro-Season	Factories No	MAUS No	Generation Nodes (Unique Fields)	Total Fleet	Harvesters No	Transaction No
2015	2	8	256	151	26	16044
2016	2	9	320	252	26	23003
2017	2	9	337	229	23	28236
2018	2	9	247	170	24	20802
2019	2	10	184	129	23	17690
2020	1	7	134	92	26	10135
2021	1	6	189	88	24	19323
2022	1	7	145	112	21	12537
2023	1	8	154	78	19	13249
2024	1	7	167	99	20	7355
2025	1	5	120	131	16	11795

Source: Compiled by the author based on empirical transactional logs of Südzucker Moldova SRL.

Table B2. Analysis of daily fleet dispersion across mobile loading complexes (2015–2025)

Season	Factory ID Type	MAUSID	Avg Cars per Day	Operating Days
2015	DR MIN (Lowest)	1	11.2	4
2015	DR MAX (Best)	2	18.7	19
2015	FL MIN (Lowest)	3	6.3	3
2015	FL MAX (Best)	4	23	37
2016	DR MIN (Lowest)	5	11.9	34
2016	DR MAX (Best)	6	16.6	95
2016	FL MIN (Lowest)	7	13.8	49
2016	FL MAX (Best)	8	19.9	50
2017	DR MIN (Lowest)	9	6.8	4
2017	DR MAX (Best)	10	21.2	80
2017	FL MIN (Lowest)	11	10.8	4
2017	FL MAX (Best)	12	23.4	75
2018	DR MIN (Lowest)	13	8.2	4
2018	DR MAX (Best)	14	19	82
2018	FL MIN (Lowest)	15	6	1
2018	FL MAX (Best)	16	25.1	53
2019	DR MIN (Lowest)	17	4.5	17
2019	DR MAX (Best)	18	17.7	41
2019	FL MIN (Lowest)	19	6.7	3
2019	FL MAX (Best)	20	18.1	21
2020	DR MIN (Lowest)	21	5	2
2020	DR MAX (Best)	22	12.6	58
2021	DR MIN (Lowest)	23	8.7	32
2021	DR MAX (Best)	24	13.1	108
2022	DR MIN (Lowest)	25	10.1	53
2022	DR MAX (Best)	26	17.8	69
2023	DR MIN (Lowest)	27	7.9	15
2023	DR MAX (Best)	28	12.7	75
2024	DR MIN (Lowest)	29	6.5	49
2024	DR MAX (Best)	30	13.6	50
2025	DR MIN (Lowest)	31	10.2	75
2025	DR MAX (Best)	32	18.3	97

Source: Compiled by the author based on empirical transactional logs of Südzucker Moldova SRL.

Table B3. Dynamics of average truck processing time at reception yards (2015–2025)

Season	Factory ID	Avg_Time_Minutes	Total_Trucks_Measured
2015	DR	24.6	10145
2015	FL	16.7	5899
2016	DR	20.6	15885
2016	FL	18.3	7118
2017	DR	16.6	16680
2017	FL	20.1	11556

Season	Factory ID	Avg_Time_Minutes	Total_Trucks_Measured
2018	DR	17.2	12026
2018	FL	19.5	8770
2019	DR	14.6	14314
2019	FL	16.9	3376
2020	DR	14.4	10135
2021	DR	15.6	19323
2022	DR	14.4	12537
2023	DR	14.5	13249
2024	DR	13.8	7355
2025	DR	14.2	11795

Source: Compiled by the author based on empirical transactional logs of Südzucker Moldova SRL.

Appendix C: Transport Link Capacity Calculation Details

This appendix provides the detailed derivations for the transport link capacity calculations referenced in Section 2.6.4 (Equation 8).

C.1. Logistics Cycle Time

The total round-trip time of a single transport unit is defined as:

$$T_{cycle} = t_{load} + t_{transit_to} + t_{queue} + t_{unload} + t_{transit_back} \quad (C1)$$

Where:

- t_{load} – loading time at the field pile,
- $t_{transit_to}$ – loaded transit time to the factory,
- t_{queue} – waiting time at the factory reception,
- t_{unload} – unloading time,
- $t_{transit_back}$ – empty return time to the field.

C.2. Ideal Fleet Size per Loader

To prevent loader starvation (i.e., to keep the MAUS continuously operational), the theoretical number of trucks required is:

$$N_{ideal,m} = \left\lceil \frac{T_{cycle}}{t_{load}} \right\rceil \quad (C2)$$

C.3. Required Trips to Fulfill the Daily Plan

The number of trips needed to execute the gross daily loading plan for loader m is:

$$R_{required,m}(t) = \left\lceil \frac{Plan_{brutto,m}(t)}{Cap_a} \right\rceil \quad (C3)$$

Where:

- $Plan_{brutto,m}(t)$ – is the gross daily loading plan (tons/day);
- Cap_a – is the nominal payload capacity of a single truck (tons).

C.4. Maximum Trips per Vehicle per Day

Given the available daily time fund of a vehicle $T_{calendar,a}(t)$ (hours/day), the maximum number of trips one truck can complete is:

$$N_{trips_per_vehicle_per_day} = \left\lfloor \frac{T_{calendar,a}(t)}{T_{cycle}} \right\rfloor \quad (C4)$$

C.5. Minimum Required Fleet Size

The minimum number of vehicles required to execute the daily plan is:

$$N_{required,m}(t) = \left\lceil \frac{R_{required,m}(t)}{N_{trips_per_vehicle_per_day}} \right\rceil \quad (C5)$$

C.6. Actual Fleet Allocation

The actual number of vehicles assigned to the link is determined by the tighter of the two physical and economic constraints:

$$N_{actual,m}(t) = \min \left(N_{ideal,m}, \max(N_{required,m}(t), 1) \right) \quad (C6)$$

This formulation guarantees:

- at least one vehicle is always assigned to each active loader,
- no more than the ideal fleet is allocated, preventing unnecessary idle capacity.

C.7. Link Throughput Capacity

The cumulative daily throughput of the transport link assigned to loader m is:

$$Cap_{link,m}(t) = \sum_{a=1}^{N_{actual,m}(t)} \left(\frac{T_{calendar,a}(t)}{T_{cycle}} \cdot Cap_a \right) \quad (C7)$$

C.8. Aggregate Fleet Capacity for Factory Zone

The total fleet capacity for factory zone b is obtained by aggregating the throughput of all active links:

$$Cap_{Fleet,b}(t) = \sum_{m \in MAUS_b} Cap_{link,m}(t) \quad (C8)$$

This aggregated value is used in Equation (8) of the main text as the transport capacity constraint in the Dynamic Drum identification.

C.9. Objective Function Components Definition

This subsection provides the explicit formulations for the three penalty components used in the objective function (Equation 9 of the main text).

C.9.1. Biochemical Degradation Loss

The loss from sucrose degradation in field storage nodes is computed as:

$$Loss_{biochem}(t) = \sum_{j \in V_C} (B_j(t) \cdot \Delta S_j(t) \cdot \eta \cdot P_{sugar}) \quad (C9)$$

Where:

$B_j(t)$ – remaining biomass volume in field pile j at day t (t),

$\Delta S_j(t)$ – predictive sucrose loss coefficient for field pile j on day t , provided by the PINN model (fraction),

η – standardized technological sugar extraction efficiency coefficient (fraction),

P_{sugar} – current market value of one ton of white sugar (monetary units/t).

Interpretation: This term penalizes piles with high predicted degradation, forcing the optimizer to prioritize their evacuation (Smart Triage).

C.9.2. Factory Starvation Penalty

The penalty for insufficient supply to the processing plant is:

$$Penalty_{starve}(t) = P_{starve} \cdot \sum_{b \in V_B} X_{short_fac,b}(t) \quad (C10)$$

Where:

P_{starve} – specific economic penalty for raw material undersupply (monetary units/t),

$X_{short_fac,b}(t)$ – continuous slack variable representing the deficit in factory buffer b at day t (t).

Interpretation: This term activates when the factory buffer cannot satisfy current processing demand, preventing production stoppages.

C.9.3. Logistical Slack Penalty

The penalty for resorting to artificial mass deficit (pressure relief) is:

$$Penalty_{slack}(t) = P_{slack} \cdot \sum_{j \in V_C} X_{short,j}(t) \quad (C11)$$

Where:

P_{slack} – prohibitive penalty for using slack variables (monetary units/t),

$X_{short,j}(t)$ – continuous slack variable representing artificial mass deficit at generation node j on day t (t), introduced to maintain numerical feasibility under force majeure.

Interpretation: This term is a mathematical safeguard. Under normal conditions $X_{short,j}(t) = 0$; it only becomes active when physical constraints prevent full evacuation, allowing the solver to return a feasible (though suboptimal) plan.

References

- [1] D. Ivanov and A. Dolgui, "A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0," *Production Planning & Control*, vol. 32, no. 9, pp. 775-788, 2021.
- [2] M. Rosca and R. Ciobanu, "Smart Cities and the Rural-Urban Divide in Moldova: Balancing Innovation and Tradition," *Smart Cities and Regional Development (SCRD) Journal*, vol. 9, no. 3, pp. 21-38, 2025.
- [3] E. V. Khudyakova, M. N. Stepansevich, and K. S. Muzalev, "Sovershenstvovanie upravleniya biznes-processami tsekha sakharnogo zavoda na osnove tekhnologii tsifrovyykh dvoynikov (Improving the business process management of a sugar factory shop based on digital twin technology)," *Izvestiya Timiryazevskoy sel'skokhozyaystvennoy akademii*, no. 1, pp. 15-25, 2025. (in Russian).
- [4] C. Koç, T. Bektas, O. Jabali, and G. Laporte, "An adaptive large neighborhood search for heterogeneous vehicle routing problem with time windows," *Computers & Operations Research*, vol. 64, pp. 105-115, 2016.
- [5] A. Balashov, "Justification of the logistics of transporting root crops to the beet point of a sugar factory," *Science in the Central Russia*, no. 1, pp. 54-60, 2026.
- [6] M. Asadi, *Beet-sugar handbook*. Hoboken, New Jersey: John Wiley and Sons, Inc., 2007.
- [7] D. Drenovac, Đ. Stakić, and M. Milanović, "Sugar beet transportation problem under growers' equity regulations: Metaheuristic approach," *International Journal of Industrial Engineering Computations*, vol. 16, no. 4, pp. 1123-1142, 2025.
- [8] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in *Transdisciplinary Perspectives on Complex Systems*, Springer, 2017, pp. 85-113.
- [9] F. Tao, Q. Qi, A. Liu, and A. Kusiak, "Digital Twins and Cyber-Physical Systems toward Smart Manufacturing and Industry 4.0: Correlation and Comparison," *Engineering*, vol. 5, no. 4, pp. 653-661, 2019.
- [10] Y. Siradeghyan and N. Shirvani, "Order Releasing and Scheduling for a Multi-Item MTO Industry: An Efficient Heuristic Based on Drum Buffer Rope," *Applied Sciences*, vol. 12, no. 4, p. 1925, 2022.
- [11] E. M. Goldratt, *The Goal: A Process of Ongoing Improvement*, 3rd ed. North River Press, 2004.
- [12] E. M. Goldratt, *Critical Chain*. North River Press, 1997.
- [13] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789-814, 2000.
- [14] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Publishing, 2009.
- [15] C. Iațco, "Physics-Informed Digital Twin for Sucrose Loss Prevention and Transparent Governance in Regional Sugar Beet Supply Chains," *Smart Cities and Regional Development (SCRD) Journal*, (under review). (companion article).
- [16] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686-707, 2019.
- [17] A. Rong, R. Akkerman, and M. Grunow, "An optimization approach for managing fresh food quality throughout the supply chain," *International Journal of Production Economics*, vol. 131, no. 1, pp. 421-429, 2011.
- [18] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [19] H. P. Williams, *Model Building in Mathematical Programming*, 5th ed. John Wiley & Sons, 2013.
- [20] G. L. Nemhauser and L. A. Wolsey, *Integer and Combinatorial Optimization*. Wiley, 1988.
- [21] C. Vrabie, "Editorial: From Smart Systems to Smart Governance," *Smart Cities and Regional Development (SCRD) Journal*, vol. 10, no. 2, p. C2, 2026.
- [22] ISA-95 (IEC 62264), *Enterprise-Control System Integration*. International Society of Automation, 2000-2015.
- [23] P. Holme and J. Saramäki, "Temporal Networks," *Physics Reports*, vol. 519, no. 3, pp. 97-125, 2012.
- [24] O. Ahumada and J. R. Villalobos, "Application of planning models in the agri-food supply chain: A review," *European Journal of Operational Research*, vol. 196, no. 1, pp. 1-20, 2009.
- [25] A. Higgins, "Scheduling of road vehicles in sugarcane transport: A case study at an Australian sugar mill," *European Journal of Operational Research*, vol. 170, no. 3, pp. 987-1000, 2006.

- [26] J. A. Marmolejo-Saucedo, "Utilising Digital Twins to Bolster the Sustainability of Logistics Processes in Industry 4.0," *Sustainability*, vol. 16, no. 6, p. 2575, 2024.
- [27] S. Fryshev, M. Volokha, and M. Ikalchik, "Analysis of throughput rate of harvesting and transportation complex of machines for sugar beet production," *Naukovij žurnal «Tehnika ta energetika»*, vol. 12, no. 1, pp. 84-93, 2022.
- [28] I. Fikry, M. Gheith, and A. Eltawil, "An integrated production-logistics-crop rotation planning model for sugar beet supply chains," *Computers & Industrial Engineering*, vol. 157, p. 107300, 2021.
- [29] K. A. Bakarr, J. I. Chernor, A. S. Marrah, et al., "Smart city technologies, resource allocation and sustainable urban planning - A critical review," *Smart Cities and Regional Development (SCRD) Journal*, vol. 10, no. 1, pp. 33-54, 2026.
- [30] D. Brcanov, S. Dakić, D. Đokić, N. Gvozdenović, and S. Zekić, "Optimization of transport activities in the sugar beet harvesting campaign," *Economics of Agriculture (Ekonomika poljoprivrede)*, vol. 72, no. 1, pp. 155-169, 2025.
- [31] Baunach, Ludwig and Spinler, Stefan and Birge, John R., "Integrating Harvest Planning Along the Supply Chain: A Case Study From the European Sugar Beet Industry," November 26, 2025. Available: <http://dx.doi.org/10.2139/ssrn.5341899>
- [32] J. Lee, "Cyber Physical Systems for Predictive Manufacturing," in *Proceedings of the 2015 IEEE International Conference on Automation Science and Engineering (CASE)*, 2015, pp. 1–6.