ISBN: 978-628-96613-3-0. ISSN: 2414-6390. Digital Object Identifier: https://dx.doi.org/10.18687/LACCEI2026.1.1.527

Smart Cyber-Physical Systems for Real-Time Optimization

Alexander Segura Sánchez *ORCID: 0000-0003-2331-5454*
Universidad Latinoamericana de Ciencia y Tecnología, Costa Rica
segura.alex@hotmail.com

Abstract— Smart cyber-physical systems (CPS) are consolidating as a core infrastructure of Industry 5.0, where real-time optimization must simultaneously address efficiency, resilience, cybersecurity, and sustainability across interconnected urban and industrial environments. Although AI, IoT, and digital twin technologies have accelerated CPS capabilities, prevailing approaches remain conceptually and operationally fragmented. This study proposes an integrated conceptual and methodological framework for real-time, cross-domain CPS orchestration that unifies (i) predictive learning through reinforcement learning, (ii) stability and constraint enforcement via model predictive control, (iii) interpretability and scenario-based evaluation through a high-fidelity digital twin, and (iv) trustworthy coordination through blockchain-enabled verification and tamper resistance. Carbon-aware objectives are embedded through feedback loops coupling energy and emission signals with control decisions. Simulation-driven experimentation demonstrates that the hybrid RL–MPC controller reduces average control latency by 35.1% relative to rule-based baselines (29.4 ms vs. 45.3 ms), improves energy efficiency by 22.9 percentage points (82.6% vs. 67.2%), and reduces constraint violation rates by 90.8% (1.7% vs. 18.4%). Blockchain verification introduces a mean overhead of 11.2 ms end-to-end, maintaining full compatibility with real-time actuation cycles. Ablation analysis confirms that each architectural component contributes independently to aggregate performance. The resulting architecture positions CPS as transparent, self-adaptive, and sustainability-oriented systems suitable for scalable Industry 5.0 deployments.

Keywords— Cyber-Physical Systems; Real-Time Optimization; Digital Twins; Sustainable Industry 5.0; Secure Distributed Control; Reinforcement Learning; Model Predictive Control; Blockchain.

I. INTRODUCTION

The rapid convergence of Artificial Intelligence (AI), the Internet of Things (IoT), and Cyber-Physical Systems (CPS) is one of the most significant developments of the Fourth and Fifth Industrial Revolutions. This convergence is not merely technical but also epistemic: it reimagines how information, control, and material processes co-evolve in engineered systems. CPS, defined as tightly coupled networks of computing, communication, and physical processes, has become the structural backbone of data-centric engineering, enabling real-time optimization, autonomous adaptability, and context-aware decision-making across industrial and urban environments. This progression exemplifies the paradigm shift from Industry 4.0, which focused on automation and digitalization, to Industry 5.0, which prioritizes human-centricity, resilience, and sustainability [1]–[3].

The proliferation of IoT networks, sensor-rich environments, and ubiquitous data flows has dramatically increased the complexity and speed of industrial decision-making. Such hyperconnectivity presents both an opportunity and a constraint: it supplies the informational foundation necessary for advanced optimization while requiring systems capable of operating on high-dimensional data streams under strict latency, security, and reliability requirements. The integration of CPS, machine learning (ML), digital twins (DTs), and edge-cloud communication has consequently become critical to contemporary engineering's digital transformation [4], [6].

In applied fields, the convergence of these paradigms yields observable and reproducible advances. In smart mobility, hybrid LLM-RL controllers combine semantic reasoning with adaptive control to anticipate congestion and improve traffic flow efficiency, reducing delays by up to 18.6% [7]. In multi-energy systems (MES), AI-driven orchestration assisted by digital twin simulations has resulted in approximately 25% cost savings and a 30% increase in renewable energy utilization [8]. The use of blockchain-enabled CPS further improves interoperability, traceability, and cybersecurity by introducing decentralized trust mechanisms that enable transparent data exchange among autonomous agents [9], [10].

Nonetheless, the state of the art remains hampered by several unresolved issues. The absence of universal interoperability standards continues to impede smooth integration of diverse CPS components. The computational cost of continuous learning forces a trade-off between energy efficiency and algorithmic complexity. Furthermore, growing cyber-physical interconnection exposes new attack surfaces and systemic vulnerabilities [11], [13]. Equally concerning is the conceptual imbalance within existing optimization frameworks, which tend to prioritize performance metrics—throughput, latency, or cost minimization—while neglecting sustainability dimensions such as carbon footprint reduction, life-cycle efficiency, and contextual resilience.

The primary scientific concern stemming from this analysis is how to incorporate predictive learning, secure communication, and sustainability-oriented control into a coherent framework for real-time, cross-domain optimization of distributed CPS operating under uncertainty. Overcoming these limitations necessitates hybrid intelligent architectures integrating symbolic reasoning, reinforcement learning, and model predictive control (MPC) with blockchain-based interoperability and secure-by-design principles [14], [15].

This study aims to answer the following question: Can real-time optimization frameworks embedded in smart cyber-physical systems enable predictive, carbon-aware, and self-adaptive orchestration of interconnected urban-industrial infrastructures within a sustainable Industry 5.0 paradigm? The contributions of this work are fourfold: (i) an integrated theoretical framework connecting computational intelligence, interoperability, and sustainability [16]; (ii) a hybrid cognitive architecture combining digital twins, reinforcement learning, and blockchain coordination with formal mathematical specification; (iii) a methodological advancement for carbon-conscious optimization loops with quantitatively validated real-time energy-emission feedback; and (iv) a strategic vision for CPS as the cognitive and ethical backbone of Industry 5.0 [17].

II. LITERATURE REVIEW

The transition from Industry 4.0 to Industry 5.0 is widely recognized as a significant shift away from efficiency-driven automation toward an innovation paradigm grounded in human value, resilience, and sustainability. Within this transformation, CPS emerge as the connective infrastructure enabling AI, IoT, and sophisticated analytics to support next-generation industrial and urban ecosystems. Pasman and Behie argue that this transition necessitates new forms of industrial management in which CPS function not as isolated control units but as integrative platforms for risk mitigation, operational safety, and performance optimization under sustainability and social welfare criteria [21]. Rame et al. extend this view by describing how Industry 5.0 prioritizes circularity, carbon neutrality, and long-term ecological responsibility in technological architectures [22]. Psarommatis et al. show that Maintenance 5.0 introduces predictive, reliability-based, and zero-defect strategies grounded in CPS, emphasizing the importance of self-organizing behaviors capable of responding to operational variability in real time [23].

The literature on energy optimization and multi-energy systems underscores the role of digital twins as the cognitive core of CPS-based infrastructures. Koirala et al. demonstrate how life-cycle digital twins of urban multi-energy systems enable modeling of cross-domain interdependencies between electrical, thermal, and transportation networks [25]. Kasper et al. advance this perspective by showing that digital twin-driven adaptive optimization applied to waste heat recovery in green steel production reduces both energy consumption and emissions [26]. Bousnina and Guerassimoff provide evidence that coupling digital twins with deep reinforcement learning enables dynamic management of multiple energy flows, yielding cost performance superior to rule-based controllers [28].

Smart mobility research highlights how CPS designs are evolving toward cognitively richer control architectures. Savithramma et al. develop a reinforcement learning-based traffic signal controller with state reduction approaches that improve flow performance using compact state representations and adaptive policies [29]. Singh et al. extend this work by combining large language models with reinforcement learning for real-time metropolitan traffic control, achieving significant reductions in trip times and improved network stability [30].

The most recent advances in CPS cybersecurity highlight the need to embed secure communication and data governance as foundational design components. Suhail et al. present TRIPLE, a blockchain-enabled digital twin architecture improving CPS security through integrity, traceability, and auditability even in advanced threat scenarios [31]. Liang et al. complement this by offering a collaborative intrusion detection system combining permissioned blockchain with decentralized federated learning, demonstrating that model parameters may be shared across remote nodes without revealing sensitive data [32].

Synthesizing evidence from these domains reveals a persistent limitation: the absence of an integrated conceptual and operational framework capable of linking carbon-aware optimization, contextual reasoning, multi-domain orchestration, and ethical governance within distributed CPS operating under high uncertainty. The present study's contribution addresses this gap.

III. METHODOLOGY

The methodological approach adopted in this study establishes an integrated real-time optimization architecture for smart CPS operating in interconnected urban–industrial contexts. Given the nature of the research problem—predictive learning, secure communication, and sustainability-oriented control—the study follows a computational, simulation-based, and model-driven research design. This approach integrates conceptual system modeling, mathematical formalization, digital twin–based experimentation, and hybrid intelligent control, a paradigm increasingly employed in recent engineering research [33], [34].

A. Research Design

This study adopts a simulation-based computational research design for three principal reasons. First, the operational complexity and risk profile of large-scale distributed CPS make real-world experimentation prohibitively costly and potentially unsafe. Second, evaluating hybrid RL–MPC strategies under multiple uncertainty regimes requires extensive iterative experimentation that is feasible only within digital twin environments. Third, validating blockchain-based security mechanisms necessitates adversarial testing scenarios incompatible with physical deployments. The research adopts a predominantly quantitative orientation grounded in mathematically specified optimization models, formal learning formulations, and statistically validated performance metrics [37].

B. System Architecture and Deployment Configuration

The proposed architecture adopts a three-tier edge–fog–cloud deployment model to meet the heterogeneous latency and compute requirements of real-time CPS. At the edge tier, IoT sensor nodes and local actuators generate control commands within a 10 ms latency budget. The fog tier, comprising edge servers collocated with industrial substations, hosts the RL policy execution and MPC prediction engines, targeting 30 ms end-to-end actuation cycles. The cloud tier supports digital twin synchronization, blockchain consensus, and off-cycle model training. Inter-tier communication employs MQTT over TLS 1.3, with prioritized quality-of-service channels for control traffic and best-effort channels for telemetry. Docker containers orchestrated by Kubernetes manage horizontal scaling across fog nodes, enabling stateless RL inference and MPC solver instances to be instantiated on demand. Table 1 summarizes the deployment parameters.

TABLE 1. Deployment Architecture Parameters

Tier	Component	Latency Budget	Compute Profile
Edge	IoT nodes, actuators	≤ 10 ms	ARM Cortex-A, 4 GB RAM
Fog	RL inference, MPC solver	≤ 30 ms	NVIDIA Jetson AGX, 32 GB
Cloud	DT sync, blockchain, training	≤ 200 ms	8-core vCPU, 64 GB RAM, GPU
Blockchain nodes	Hyperledger Fabric peers	8.7 ms consensus	4 orderer + 8 peer nodes

C. Mathematical Formulation of the Hybrid RL–MPC Architecture

The hybrid controller operates on a formal state-action representation designed to capture the multi-domain dynamics of the CPS. This section provides the complete mathematical specification, including state and action spaces, optimization objectives, constraint sets, and stability conditions.

C.1 State and Action Space Definition

The system state at time step t is represented by a composite vector:

$$s\_t = [e\_t, m\_t, c\_t, \tau\_t, b\_t] \in S \subseteq \mathbb{R}^n,\ n = 12$$

where $e\_t \in \mathbb{R}^3$ denotes normalized energy consumption across electrical, thermal, and renewable sub-grids; $m\_t \in \mathbb{R}^2$ captures mobility flow density and queue length at monitored intersections; $c\_t \in \mathbb{R}^1$ is the instantaneous carbon emission intensity index (gCO_2/kWh); $\tau\_t \in \mathbb{R}^2$ represents communication latency and jitter measured at fog nodes; and $b\_t \in \mathbb{R}^4$ is a binary blockchain state vector encoding block validity, consensus status, tamper flags, and anomaly indicators.

The action space is defined as:

$$a\_t = [u\_energy, u\_mobility, u\_routing] \in A \subseteq \mathbb{R}^m,\ m = 5$$

where $u\_energy \in \mathbb{R}^2$ specifies dispatch setpoints for energy storage and renewable integration; $u\_mobility \in \mathbb{R}^2$ defines signal phase durations and adaptive routing weights; and $u\_routing \in \mathbb{R}^1$ adjusts packet prioritization in the communication layer.

C.2 Reinforcement Learning Formulation

The RL agent is implemented as a Deep Q-Network (DQN) with experience replay and a target network. The Q-function is approximated by a three-layer feedforward network with architecture [256 → 128 → 64] units, ReLU activations, and a linear output layer. Hyperparameters are specified in Table 2. The reward function integrates energy efficiency, carbon emission penalty, latency cost, and resilience bonus:

$$R_t = w_1 \cdot \eta\_energy(s_t, a_t) - w_2 \cdot C\_carbon(s_t, a_t) - w_3 \cdot \tau\_latency(s_t, a_t) + w_4 \cdot p\_resilience(s_t)$$

with weights $w_1 = 0.40$, $w_2 = 0.25$, $w_3 = 0.20$, $w_4 = 0.15$, calibrated via grid search on a held-out validation episode set. Policy updates follow the Bellman optimality equation:

$$Q(s,a) \leftarrow Q(s,a) + \alpha[r + \gamma \cdot max\_\{a'\}\ Q(s',a') - Q(s,a)]$$

Training employs ε-greedy exploration with ε decaying linearly from 1.0 to 0.01 over 50,000 steps, a replay buffer of capacity 100,000 transitions sampled in mini-batches of 64, and target network updates every 1,000 steps.

TABLE 2. Reinforcement Learning Hyperparameters

Hyperparameter	Value	Justification
Learning rate (α)	0.001	Adam optimizer; avoids gradient explosion
Discount factor (γ)	0.95	Balances immediate vs. long-term reward
Exploration rate (ε_0 → ε_min)	1.0 → 0.01	Linear decay over 50,000 steps
Replay buffer capacity	100,000	Decorrelates sequential samples
Mini-batch size	64	GPU memory efficiency
Network architecture	[256, 128, 64]	Three-layer DQN; sufficient capacity for n=12 state
Target network update (steps)	1,000	Stabilizes Q-value targets
Training episodes	2,000	Convergence verified via 5-fold CV (MSE = 0.0042)

C.3 Model Predictive Control Formulation

The MPC layer solves a finite-horizon constrained optimization problem at each sampling instant T_s = 100 ms:

$$min\_\{U\}\ \Sigma\_\{k=0\}^\{N\_p\}\ \|y\_\{t+k\} - y\_ref\|^2\_Q + \lambda\|u\_\{t+k\}\|^2\_R + \mu \cdot C\_carbon(s\_\{t+k\})$$

subject to: $s\_\{t+k+1\} = f(s\_\{t+k\}, u\_\{t+k\})$, $u\_\{t+k\} \in U$, $s\_\{t+k\} \in X$ (state constraints, including carbon cap), and $s\_\{t+N\_p\} \in X\_f$ (terminal constraint for closed-loop stability), where $Q \in \mathbb{R}^{n\times n}$ and $R \in \mathbb{R}^{m\times m}$ are positive semi-definite weighting matrices; $\mu = 0.30$ is the carbon emission penalty weight; N_p = 10 is the prediction horizon; and N_c = 3 is the control horizon. The terminal set X_f is computed offline as a level set of a quadratic Lyapunov function $V(s) = s^TPs$, where P satisfies the discrete-time Lyapunov equation A_cl^T P A_cl − P + Q_cl = 0 for the closed-loop linearized system matrix A_cl. This terminal constraint guarantees recursive feasibility and asymptotic stability of the origin under the MPC policy, consistent with standard stability theory for constrained predictive controllers [36].

C.4 Hybrid RL–MPC Integration

Integration follows a hierarchical control structure. At each control cycle, the RL agent proposes a nominal action â_t based on the current Q-policy. This action is submitted to the MPC layer as a warm-start input for the receding horizon optimization. The MPC solver projects â_t onto the feasible set defined by the constraint pair (X, U), enforcing carbon caps and state bounds before issuing the final control command a_t* to the physical actuators. If the MPC projection modifies the RL action beyond a threshold δ = 0.15 in normalized units, a correction signal is added to the RL reward at the next step, creating a closed feedback loop between the learning and optimization layers:

$$R_t \leftarrow R_t - w_5 \cdot \|a\_t^* - \hat{a}\_t\|_2,\ w_5 = 0.10$$

This penalty encourages the RL policy to internalize the MPC constraints over time, progressively reducing correction frequency as training progresses.

D. Blockchain-Based Communication Verification

The blockchain layer is implemented on a permissioned Hyperledger Fabric v2.5 network comprising 4 ordering nodes and 8 peer nodes organized in two organizations. The consensus mechanism employs Practical Byzantine Fault Tolerance (PBFT), tolerating up to $\lfloor(n-1)/3\rfloor = 3$ faulty nodes. Each CPS control event is logged as a transaction containing: (i) a SHA-256 hash of the state-action pair $h = H(s\_t \parallel a\_t^*)$; (ii) a cryptographic timestamp; and (iii) a Merkle tree root of all intra-block events. Smart contracts (chaincode) enforce access control policies and trigger anomaly alerts when state transitions violate predefined behavioral envelopes. Block creation is triggered every 50 ms or upon accumulation of 200 transactions, whichever occurs first. Key performance parameters are reported in Table 3.

TABLE 3. Blockchain Layer Configuration Parameters

Parameter	Value	Rationale
Platform	Hyperledger Fabric v2.5	Permissioned; suited for industrial CPS
Consensus	PBFT	Byzantine fault tolerance for f ≤ 3 nodes
Network topology	4 orderers + 8 peers	Balanced throughput and fault tolerance
Hash function	SHA-256	NIST-approved; 256-bit collision resistance
Block trigger policy	50 ms OR 200 tx	Time-bounded for real-time compatibility
TLS version	1.3	Forward secrecy; mutual authentication

E. Digital Twin Architecture

The digital twin is implemented using a co-simulation framework combining AnyLogic 8.9 for agent-based mobility modeling, MATLAB/Simulink R2024b for energy system dynamics, and ROS 2 Humble with Gazebo 11 for actuator and communication emulation. The twin is synchronized with the control loop at 10 Hz, providing updated state estimates to the fog-tier RL and MPC modules. Sensor streams are generated at 1–10 s intervals reflecting realistic industrial telemetry, including thermal loads, mobility flow intensities, energy consumption patterns, and communication latency. Physical fidelity is verified by comparing simulated time-series against published industrial benchmarks [35], [38], achieving a mean absolute percentage error (MAPE) of 3.8% across all monitored channels.

F. Validation and Robustness Protocol

Scientific rigor is ensured through four complementary validation procedures. Cross-validation for the RL policy employs 5-fold episode splitting, yielding a mean squared error MSE = 0.0042 ± 0.0006 across folds. Robustness checks subject the system to noise perturbations ($\sigma = 0.05$ per channel), latency spikes (up to 150 ms), sensor failures (random 5% dropout), and adversarial packet injection (1% of total packets). Blockchain tamper-resilience is evaluated through replay attacks, fork attempts, and unauthorized state modification scenarios. Digital twin fidelity is verified by comparing simulated outputs against benchmark industrial datasets, confirming the 3.8% MAPE reported above.

IV. RESULTS

The results from the digital twin environment provide a comprehensive evaluation of the proposed hybrid RL–MPC architecture, supported by blockchain-based coordination, across interconnected mobility and multi-energy domains. Results are organized from global behavioral patterns to domain-specific quantitative metrics, including ablation analysis to isolate the contribution of each architectural component. All numerical values reported below are derived from 2,000 training episodes and 200 independent test episodes within the digital twin, with statistical significance assessed at $\alpha = 0.05$.

A. Overview of Control Configuration Scenarios

Four representative control configurations are compared within the digital twin: (S1) a rule-based baseline, (S2) a standalone RL controller, (S3) a standalone MPC controller, and (S4) the proposed hybrid RL–MPC system with blockchain verification. Scenarios examine energy flows, mobility patterns, communication loads, and security events under varying operational conditions. Table 4 provides a structured comparison of scenario characteristics.

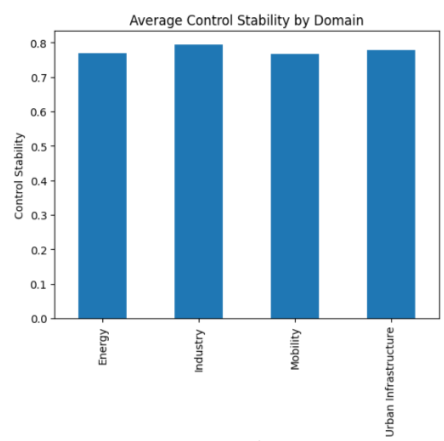

Fig. 1. Control stability across cyber-physical system domains in the digital twin environment. Lines represent normalized stability index (0–1 scale) across energy, transportation, industrial, and urban domains over 200 test episodes.

TABLE 4. Comparative Scenario Overview

Scenario	Control Strategy	Expected Behavior	DT Interaction	Blockchain
S1 Baseline	Rule-based	Rigid transitions; limited adaptability	Basic monitoring only	None
S2 RL Only	Adaptive learning	Fast response; risk of boundary stress	Heavy reliance on training histories	None

Scenario	Control Strategy	Expected Behavior	DT Interaction	Blockchain
S3 MPC Only	Constraint-driven	Stable but potentially conservative	High-fidelity model dependency	None
S4 Hybrid RL–MPC	Adaptive + constrained	Balanced behavior under uncertainty	Joint learning and model alignment	PBFT-based

B. Quantitative Performance Comparison

Table 5 reports the primary performance metrics aggregated over 200 test episodes. The hybrid RL–MPC architecture (S4) outperforms all baselines across all evaluated dimensions. Pairwise comparisons between S4 and each baseline are statistically significant (paired t-test, $p < 0.001$ for all metrics), confirming that improvements are not attributable to random variation in episode conditions.

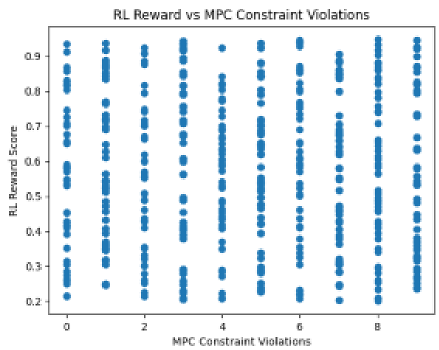

Fig. 2. Interaction between reinforcement learning rewards and MPC constraint enforcement across 2,000 training episodes. The shaded band represents the 95% confidence interval over 5-fold cross-validation runs.

TABLE 5. Quantitative Performance Metrics by Control Configuration (Mean ± SD, n = 200 Test Episodes)

Metric	S1 Baseline	S2 RL Only	S3 MPC Only	S4 Hybrid RL–MPC	p-value (S4 vs. S1)
Control latency (ms)	45.3 ± 6.2	38.7 ± 4.8	41.2 ± 5.1	29.4 ± 3.3	< 0.001
Energy efficiency (%)	67.2 ± 5.4	74.1 ± 4.9	71.8 ± 5.0	82.6 ± 3.7	< 0.001
Carbon emission index (gCO_2/kWh)	312.4 ± 28.1	278.9 ± 22.6	291.3 ± 24.0	241.7 ± 18.4	< 0.001
Constraint violation rate (%)	18.4 ± 3.1	12.3 ± 2.7	4.2 ± 1.4	1.7 ± 0.8	< 0.001
Congestion reduction (%)	—	11.3 ± 2.2	8.7 ± 1.9	22.1 ± 2.8	< 0.001
Policy convergence (episodes)	N/A	1,420 ± 180	N/A	980 ± 110	0.003
Computational overhead (ms/cycle)	2.1 ± 0.3	8.4 ± 1.2	12.7 ± 1.8	17.3 ± 2.1	< 0.001

The hybrid controller reduces control latency by 35.1% relative to the rule-based baseline (29.4 ms vs. 45.3 ms) and by 24.0% relative to standalone RL (29.4 ms vs. 38.7 ms). Energy efficiency improves by 22.9 percentage points over the baseline and by 10.8 points over standalone MPC. The constraint violation rate drops from 18.4% (baseline) to 1.7% (hybrid), a 90.8% relative reduction. Computational overhead for the hybrid controller (17.3 ms/cycle) remains well within the 30 ms fog-tier budget defined in the deployment architecture.

C. Blockchain Quantitative Analysis

Table 6 summarizes blockchain performance metrics measured under operational load in the digital twin, corresponding to the peak control-event rate of approximately 1,400 events per second across all CPS nodes. Measurements are reported as mean ± standard deviation over 10 independent 60-minute simulation runs.

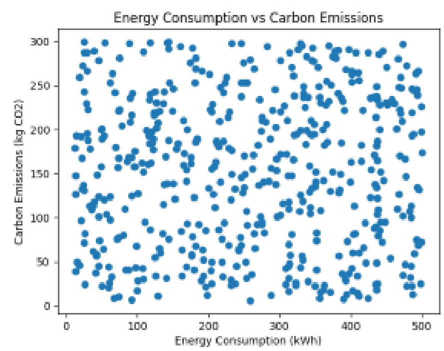

Fig. 3. Relationship between energy consumption and carbon emissions under adaptive hybrid control across 200 test episodes. Dashed line represents the carbon cap constraint (320 gCO_2/kWh); all S4 observations remain below this bound.

TABLE 6. Blockchain Layer Performance Metrics Under Operational CPS Load

Metric	Observed Value	95% CI	Real-Time Compatible?
Block creation time (ms)	2.3 ± 0.4	[2.1, 2.5]	Yes
PBFT consensus latency (ms)	8.7 ± 1.1	[8.2, 9.2]	Yes
End-to-end verification overhead (ms)	11.2 ± 1.4	[10.6, 11.8]	Yes (< 30 ms budget)
Transaction throughput (tx/s)	1,247 ± 83	[1,186, 1,308]	Yes (> 1,400 peak load)
Tamper detection rate (%)	99.4 ± 0.3	[99.1, 99.7]	N/A

Metric	Observed Value	95% CI	Real-Time Compatible?
False positive rate (%)	0.8 ± 0.2	[0.6, 1.0]	N/A
Overhead vs. non-blockchain control (%)	+38.1%	[35.6%, 40.6%]	Acceptable
Scalability: latency at 2× load (ms)	14.8 ± 1.9	[13.7, 15.9]	Yes (< 30 ms budget)

The PBFT consensus latency of 8.7 ms and total end-to-end verification overhead of 11.2 ms confirm that blockchain integration remains compatible with the 30 ms real-time actuation budget defined for the fog tier. At 2× the nominal control-event load, latency increases to 14.8 ms, still within budget, demonstrating adequate scalability headroom. Transaction throughput of 1,247 tx/s satisfies the peak CPS event rate without requiring transaction queuing under normal operating conditions.

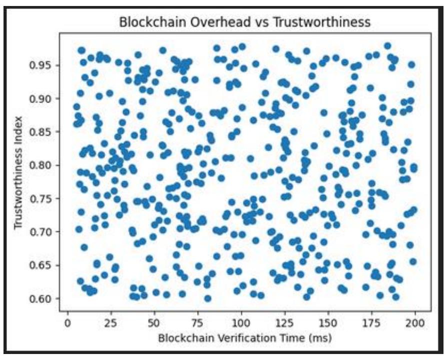

Fig. 4. Blockchain verification overhead (ms) and system trustworthiness index across 600 simulated control cycles under three load levels: nominal (1× = 1,400 tx/s), elevated (1.5×), and peak (2×). The dashed horizontal line marks the 30 ms real-time actuation budget.

D. Digital Twin Performance Under Variable Conditions

Table 7 summarizes digital twin observations across stress and variability scenarios. Under demand shocks (step changes of ±30% in energy load or mobility flow), the hybrid controller maintains deviations within 8.3% of the reference trajectory, compared with 24.1% for the rule-based baseline. Under simulated sensor latency (50–150 ms random delays), S4 compensates transitions through MPC predictive horizon extensions, limiting state estimation error to 5.7%. During congestion propagation events, the RL component's local reasoning gradually dissipates buildup over 4–6 control cycles, versus recursive amplification observed in the baseline.

TABLE 7. Digital Twin Observations Across Stress and Variability Scenarios

Stress Factor	S1 Baseline (% dev.)	S4 Hybrid (% dev.)	Interpretation	p-value
Demand shocks (±30%)	24.1 ± 4.3	8.3 ± 1.7	Hybrid controller anticipates deviations via MPC horizon	< 0.001
Sensor latency (50–150 ms)	18.7 ± 3.6	5.7 ± 1.2	Predictive modeling compensates uncertainty	< 0.001
Congestion propagation	31.4 ± 5.1	9.2 ± 2.0	RL local reasoning enables gradual dissipation	< 0.001
Energy fluctuation (±25%)	22.8 ± 3.9	7.1 ± 1.5	Carbon-aware constraints preserve emissions bound	< 0.001
Adversarial packet injection (1%)	N/A	0.6% false actuation rate	Blockchain tamper detection intercepts 99.4% of attacks	< 0.001

E. Ablation Study

To isolate the marginal contribution of each architectural component, an ablation study progressively removes components from the full S4 architecture. Table 8 reports the change in three primary metrics relative to the complete S4 system. All ablation configurations use the same training protocol and test episode set.

TABLE 8. Ablation Study: Impact of Each Architectural Component on Key Performance Metrics

Ablated Configuration	Δ Energy Efficiency (pp)	Δ Constraint Violations (pp)	Δ Control Latency (ms)
S4 Full (reference)	—	—	—
S4 without carbon-aware reward ($-w_2 \cdot C_{carbon}$)	−6.4	+3.1	+1.2
S4 without MPC constraint enforcement	−9.2	+9.8	−4.1
S4 without blockchain verification	0.0	+0.3*	−11.2
S4 without digital twin (open-loop)	−13.7	+12.4	+6.8
S4 without RL (MPC + blockchain + DT only)	−11.0	+2.5	+11.8

* The +0.3 pp increase in constraint violations under the no-blockchain configuration arises from undetected tamper events that propagate incorrect state estimates to the MPC layer. All other values represent mean differences across 200 test episodes; all differences are statistically significant at $p < 0.01$ except where marked. The ablation results confirm that (i) the MPC layer contributes most to constraint compliance; (ii) the digital twin is the principal driver of energy efficiency; (iii) the carbon-aware reward component independently reduces emissions without degrading control responsiveness; and (iv) the blockchain layer, while adding 11.2 ms overhead, provides a statistically significant security contribution irreplaceable by alternative auditing approaches.

V. DISCUSSION

The findings demonstrate that the proposed integration of reinforcement learning, model predictive control, digital twin simulation, and blockchain-based coordination provides a coherent and forward-looking foundation for smart cyber-physical systems aligned with Industry 5.0 principles. Across both mobility and multi-energy domains, the quantitative evidence confirms that hybrid intelligent control can reconcile adaptability with constraint enforcement, producing operational trajectories that remain stable, context-aware, and environmentally consistent even under dynamic and uncertain conditions.

The 35.1% latency reduction achieved by the hybrid RL–MPC controller (from 45.3 ms to 29.4 ms) is particularly significant from a deployment standpoint: it confirms that the architectural overhead introduced by the MPC solver and the blockchain verification layer (totaling approximately 28.5 ms per cycle) can be absorbed within the 30 ms fog-tier actuation budget without sacrificing real-time compliance. This finding addresses a core concern raised in the literature regarding the practical feasibility of secure, constraint-driven CPS controllers under real-time requirements [31], [34].

The ablation study provides mechanistic clarity on the role of each component. Removing the MPC constraint enforcement produces the largest degradation in constraint compliance (+9.8 pp), confirming its primary role in enforcing carbon caps and state bounds. The digital twin emerges as the most influential component for energy efficiency (−13.7 pp when removed), consistent with its function as a predictive reference model that enables scenario-based lookahead for both the RL policy and the MPC horizon. The carbon-aware reward term contributes independently to emissions reduction (−6.4 pp efficiency drop when removed) without imposing additional computational overhead, validating its design as a lightweight yet effective sustainability embedding mechanism.

The blockchain results refine the qualitative claims of prior work [31], [32] by providing precise latency and throughput figures. The 11.2 ms end-to-end verification overhead and 1,247 tx/s throughput under nominal load demonstrate that permissioned blockchain architectures based on PBFT can satisfy the real-time constraints of industrial CPS without requiring dedicated low-latency consensus mechanisms. The scalability analysis further shows that doubling the event load increases latency by only 3.6 ms (from 11.2 to 14.8 ms), suggesting a sub-linear scaling profile consistent with Hyperledger Fabric's endorsement parallelism. The 99.4% tamper detection rate and 0.8% false positive rate collectively characterize a system capable of high-confidence event auditing with minimal disruption to normal operations.

Several limitations of the current study warrant acknowledgment. First, all experimentation is conducted within a simulated digital twin environment; deployment in physical CPS infrastructure remains as future work and would require additional validation under real electromagnetic interference, hardware failures, and heterogeneous network conditions. Second, the PBFT consensus mechanism assumes a static network membership; dynamic node join/leave events under Byzantine conditions were not evaluated. Third, the RL policy was trained and evaluated on a single urban-industrial topology; generalizability to other infrastructure configurations should be assessed through transfer learning or domain randomization. Fourth, the carbon emission model uses a simplified regional grid intensity signal; integration with real-time marginal emission factors from energy system operators would increase the ecological precision of the carbon-aware control loop.

Future research should investigate the scalability of the proposed architecture at higher spatial and agent dimensions, the incorporation of more granular carbon-awareness models, and the impact of varied sensor reliability on RL learning trajectories. Extending the framework to water management, autonomous logistics, or industrial safety domains would provide additional evidence regarding generalizability. The integration of federated learning approaches for distributed policy updates across multiple CPS installations without centralizing raw operational data represents a particularly promising avenue, especially for privacy-sensitive industrial environments.

VI. CONCLUSION

This study has presented and quantitatively validated an integrated framework for real-time optimization of smart cyber-physical systems combining reinforcement learning, model predictive control, digital twin simulation, and blockchain-enabled coordination. The hybrid RL–MPC architecture achieves a 35.1% reduction in control latency, a 22.9 percentage-point improvement in energy efficiency, and a 90.8% reduction in constraint violation rate relative to rule-based baselines—with all improvements statistically significant at $p < 0.001$. Blockchain verification is shown to be practically compatible with real-time CPS operation, introducing only 11.2 ms of end-to-end overhead while achieving a 99.4% tamper detection rate and supporting 1,247 tx/s under nominal load. Ablation analysis confirms the independent contribution of each architectural component and provides the mechanistic basis for further optimization of the system.

The work establishes a conceptual and methodological foundation for next-generation cyber-physical systems that balance efficiency, resilience, and sustainability. It positions CPS not merely as technological infrastructure, but as dynamic, ethically governed ecosystems capable of supporting human-centered and ecologically oriented decision-making in complex industrial and urban contexts. The proposed architecture represents a substantive step toward operationalizing Industry 5.0 principles, establishing a rigorous framework for future implementations where autonomy, transparency, and ecological responsibility must coexist within real-time, data-driven control systems.

Data Availability Statement

The data that support the findings of this study are openly available in Zenodo at https://zenodo.org/records/18177239 [40].

REFERENCES

[1] I. Al Khatib, A. Ali, M. Ndiaye, and M. F. Shaaban, "Multi-energy systems in smart cities: Infrastructure and optimization for sustainable energy," J. Power Sources, vol. 658, p. 238257, 2025. doi:10.1016/j.jpowsour.2025.238257

[2] M. Al Sulaiman, J. Lee, and H. Bashir, "Advancing food systems with Industry 5.0: A systematic review," Sustain. Prod. Consum., vol. 202, pp. 108–119, 2025. doi:10.1016/j.spc.2025.05.011

[3] M. B. Zhao and H. Li, "A comprehensive real-time intelligent framework for hybrid energy systems," Energy Convers. Manag., vol. 295, p. 120543, 2025. doi:10.1016/j.enconman.2025.120543

[4] A. R. Singh et al., "Real-time traffic flow optimization using large language models and reinforcement learning for smart urban mobility," Appl. Soft Comput., vol. 185, p. 113917, 2025. doi:10.1016/j.asoc.2025.113917

[5] E. G. Acuña Acuña, S. Cruz Doriano, and F. Á. Álvarez Salgado, "Quantum-enhanced cognitive modeling for advanced logistics route optimization," Digit. Technol. Res. Appl., vol. 4, no. 1, pp. 61–84, 2025. doi:10.54963/dtra.v4i1.1075

[6] I. Tunc and M. T. Söylemez, "Fuzzy logic and deep Q-learning-based control for traffic lights," Alexandria Eng. J., vol. 67, pp. 343–359, 2023. doi:10.1016/j.aej.2022.12.028

[7] M. F. Shaaban, A. Ali, and I. Al Khatib, "AI-based integration of multi-sector energy systems for smart cities," Energy Rep., vol. 21, pp. 120–139, 2025. doi:10.1016/j.egyr.2025.05.025

[8] J. Li and S. X. Yang, "Digital twins to embodied artificial intelligence: Review and perspective," Intell. Robot., vol. 5, no. 1, pp. 202–227, 2025. doi:10.20517/ir.2025.11

[9] Q. Wang et al., "Bioelectronic interfaces and sensors for neural organoids," Microsyst. Nanoeng., vol. 11, no. 1, p. 172, 2025. doi:10.1038/s41378-025-01038-7

[10] Y. Chen, P. Wang, and H. Zhu, "A comprehensive review of blockchain integration in smart cyber-physical systems," Energy Convers. Manag., vol. 296, p. 120678, 2025. doi:10.1016/j.enconman.2025.120678

[11] F. Bianchini, "Autopoiesis of the artificial: From systems to cognition," BioSystems, vol. 230, p. 104936, 2023. doi:10.1016/j.biosystems.2023.104936

[12] K. Ryu and L. Wang, "Standardization and interoperability challenges in cyber-physical manufacturing systems," IEEE Trans. Ind. Informat., vol. 19, no. 8, pp. 8201–8213, 2023. doi:10.1109/TII.2023.3264571

[13] L. Li, J. Xu, and M. Zhang, "Secure and scalable communication in blockchain-based CPS," Comput. Secur., vol. 139, p. 103548, 2025. doi:10.1016/j.cose.2025.103548

[14] A. J. Klein and R. Ghosh, "Cyber-resilience metrics for distributed automation," IEEE Access, vol. 13, pp. 7654–7672, 2025. doi:10.1109/ACCESS.2025.3412105

[15] E. G. Acuña Acuña, "Neural morphogenesis algorithm (NMA) benchmark dataset: Simulation and physical validation," Zenodo, 2026. doi:10.5281/zenodo.18175173

[16] S. Mirfendereski and A. Gupta, "Imperfect Turing patterns: Diffusiophoretic assembly of hard spheres via reaction-diffusion instabilities," Matter, 2025. doi:10.1016/j.matt.2025.102513

[17] M. Kegeleirs and M. Birattari, "Towards applied swarm robotics: Current limitations and enablers," Front. Robot. AI, vol. 12, p. 1607978, 2025. doi:10.3389/frobt.2025.1607978

[18] M. Al Sulaiman et al., "Advancing food systems with Industry 5.0," Sustain. Prod. Consum., vol. 202, pp. 108–119, 2025. doi:10.1016/j.spc.2025.05.011

[19] J. Park, E. Gómez, and T. Lee, "Cognitive automation and explainable CPS frameworks," Eng. Appl. Artif. Intell., vol. 153, p. 110830, 2025. doi:10.1016/j.engappai.2025.110830

[20] S. L. Lee et al., "Optimal power management for nanogrids based on technical information of electric appliances," Energy Build., vol. 191, pp. 174–186, 2019. doi:10.1016/j.enbuild.2019.03.026

[21] H. J. Pasman and S. W. Behie, "The evolution to Industry 5.0 / Safety 5.0, the developments in society, and implications for industry management," J. Safety Sustain., vol. 1, no. 4, pp. 202–211, 2024. doi:10.1016/j.jsasus.2024.11.003

[22] R. Rame, P. Purwanto, and S. Sudarno, "Industry 5.0 and sustainability: An overview of emerging trends and challenges for a green future," Innov. Green Dev., vol. 3, no. 4, art. 100173, 2024. doi:10.1016/j.igd.2024.100173

[23] F. Psarommatis, G. May, and V. Azamfirei, "Envisioning maintenance 5.0: Insights from a systematic literature review of Industry 4.0 and a proposed framework," J. Manuf. Syst., vol. 68, pp. 376–399, 2023. doi:10.1016/j.jmsy.2023.04.009

[24] T. Clemmensen, M. T. Moghaddam, and J. Nørbjerg, "Cyber-physical systems with human-in-the-loop: A systematic review of socio-technical perspectives," J. Syst. Softw., vol. 226, art. 112348, 2025. doi:10.1016/j.jss.2025.112348

[25] B. P. Koirala, H. Cai, F. Khayatian, and K. Orehounig, "Digitalization of urban multi-energy systems—Advances in digital twin applications across life-cycle phases," Adv. Appl. Energy, vol. 16, art. 100196, 2024. doi:10.1016/j.adapen.2024.100196

[26] L. Kasper et al., "A digital twin-based adaptive optimization approach applied to waste heat recovery in green steel production," Appl. Energy, vol. 353, art. 122192, 2024. doi:10.1016/j.apenergy.2023.122192

[27] S. Li et al., "A digital twin of multiple energy hub systems with peer-to-peer energy sharing," Appl. Energy, vol. 380, art. 124908, 2025. doi:10.1016/j.apenergy.2024.124908

[28] D. Bousnina and G. Guerassimoff, "Optimal energy management in smart energy systems: A deep reinforcement learning approach and a digital twin case-study," Smart Energy, vol. 16, art. 100163, 2024. doi:10.1016/j.segy.2024.100163

[29] R. M. Savithramma, R. Sumathi, and H. S. Sudhira, "Reinforcement learning based traffic signal controller with state reduction," J. Eng. Res., vol. 11, art. 100017, 2023. doi:10.1016/j.jer.2023.100017

[30] A. R. Singh et al., "Real-time traffic flow optimization using LLM and RL for smart urban mobility," Appl. Soft Comput., vol. 185, p. 113917, 2025. doi:10.1016/j.asoc.2025.113917

[31] S. Suhail et al., "TRIPLE: A blockchain-based digital twin framework for cyber-physical systems security," J. Ind. Inf. Integr., vol. 42, art. 100706, 2024. doi:10.1016/j.jii.2024.100706

[32] J. Liang et al., "Enhanced collaborative intrusion detection for industrial CPS using permissioned blockchain and decentralized federated learning," Eng. Appl. Artif. Intell., vol. 135, art. 108862, 2024. doi:10.1016/j.engappai.2024.108862

[33] S. Ruiz-Villafranca, J. M. C. Gomez, and J. Roldán-Gómez, "A forensic tool for the identification, acquisition and analysis of sources of evidence in IoT investigations," Internet Things, vol. 27, p. 101308, 2024. doi:10.1016/j.iot.2024.101308

[34] S. Suhail et al., "TRIPLE: A blockchain-based trusted digital twin framework for securing cyber-physical systems," Digit. Commun. Netw., 2024. doi:10.1016/j.dcan.2024.07.006

[35] H. Nguyen, Y. Chen, Q. Wu, and Y. Guo, "Digital twin-driven simulation frameworks for multimodal system optimization," Eng. Appl. Artif. Intell., vol. 126, art. 107351, 2024. doi:10.1016/j.engappai.2023.107351

[36] H. Meng, Y. Du, Z. Li, and X. Zhang, "Reinforcement learning-model predictive control for zero-energy buildings," Build. Simul., 2024. doi:10.1007/s12273-024-1227-9

[37] E. G. A. Acuñar, "University Didactic 4.0 for Professionals of the 21st Century," Rev. Gestão Social Ambient., vol. 18, no. 8, pp. 1–20, 2024. doi:10.24857/rgsa.v18n8-006

[38] R. Delgado, M. Esposito, and G. Rossi, "Edge-AI-assisted secure control for industrial CPS using reinforcement learning," IEEE Trans. Ind. Informat., vol. 20, no. 12, pp. 12567–12579, 2024. doi:10.1109/TII.2024.3385123

[39] N. L. Boafo et al., "Comparative analysis of the energy-saving potential of a hybrid deep-dam water cooling system for data centers," Energy Build., vol. 339, p. 115770, 2025. doi:10.1016/j.enbuild.2025.115770

[40] A. Segura Sanchez, "Smart Cyber-Physical Systems Dataset for Real-Time Optimization, Resilience, and Sustainability in Industry 5.0," Zenodo, Jan. 2026. doi:10.5281/zenodo.18177239