Proceedings of the 7th African International Conference on Industrial Engineering and Operations Management
Cape Town, South Africa, May 13-15, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AF07.20260168
Published: May 13, 2026

HybridNet-SCR: Network-Aware, Constraint-Guided Multi-Agent Control for Disruption-Resilient Supply Chains

Rashid Abrar Tanzil, Shaikh Shehab Ahamed, Shabin Rahman, Rafia Nur Rimti and Mohammad Shahad Raihan
Department of Industrial Engineering and Management
Khulna University of Engineering & Technology
Bangladesh
rashidabrar@pm.me, shehabahamed2003@gmail.com, shabin.rahman01@gmail.com, rafianur18rimti@gmail.com, shahadraihan2002@gmail.com

Shahir Rahman
Department of Mechanical Engineering
Bangladesh University of Engineering & Technology
Bangladesh
shahir.jojok@gmail.com

Abstract

HybridNet-SCR addresses a common gap in supply chain decision support during shocks. Most methods either ignore how disruptions spread through the network or produce plans that break basic limits. We model the supply chain as a directed graph and control each facility with one agent that updates decisions daily. The framework links three parts in one workflow. A network encoder learns compact node representations from the graph. A feasibility module adds rule penalties so decisions respect inventory balance, capacity, and storage limits. A control module then trains the agents to balance service and cost under changing demand and disrupted capacity. We test the approach in a 365-day simulation with 18 disruption events and compare against a base-stock policy. HybridNet-SCR reaches 96.48% average service with total cost 344,890, an 81.2% cost reduction over base-stock (1,830,582) while improving service by 3.3 percentage points. It also recovers 66.2% faster after disruptions (2.3 vs. 6.8 days) and reduces unmet demand by 57.0% (38,420 vs. 89,340 units). Ablations confirm that each module adds distinct value, with the full system yielding the best cost, service, and constraint compliance. Results remain strong as the network grows to 100 nodes, with inference under 0.2 seconds. Limits include single-product settings and limited decision transparency.

Keywords
Physics-Informed Supply Chains, Supply Chain Resilience, Multi-Agent RL, Graph Neural Networks, Constraint-Guided MARL

1. Introduction
Modern supply chains face frequent shocks. These include pandemics, extreme weather, trade limits and cyber incidents. Many systems still rely on static policies and slow planning. These tools assume stable demand and capacity. When conditions change within hours or days, performance drops.

Firms must balance two goals at once. They must keep customer service high. They must also keep operating costs low. A common response is to carry more inventory. This can protect service, but it increases cost and waste. Artificial intelligence can help. Yet many AI tools in supply chains are built as separate parts. One model forecasts demand.

Another model sets inventory targets. A third model optimizes routes. These components often do not share information. They can also recommend actions that violate basic operating limits.

This paper presents HybridNet-SCR, a single framework that links three capabilities:
- Network awareness: learn and use supply chain structure.
- Feasibility: respect inventory balance and capacity limits.
- Adaption: adjust decisions daily during disruption periods.

2. Related Works
This section reviews four research streams and explains why an integrated control framework is still missing.

2.1 Supply chain resilience and network vulnerability
Research on resilience accelerated after disruptions in 2020–2022. Sheffi and Rice (2005) argue that resilience depends on both robustness and recovery, not only on buffer inventory. Ivanov (2020) shows that disruptions can propagate through ripple effects that depend on network structure. Hosseini et al. (2019) report that many quantitative methods remain static and do not support real-time control across disruption phases. In practice, base-stock and threshold policies are still common, but they degrade when demand or capacity changes quickly.

2.2 Graph-based learning for supply chain networks
Graph neural networks are widely used to learn from relational structures. Kipf and Welling (2017) propose an efficient graph convolution formulation used in many later works. Kosasih and Brintrup (2022) show that topology-aware models can predict hidden links and propagation better than scalar time-series methods. Brintrup et al. (2020) use graph analytics to reveal hidden dependencies in supply networks. However, most work uses graph learning for prediction or risk scoring. It rarely produces executable control actions. HybridNet-SCR uses graph embeddings as direct inputs to control agents.

2.3 Multi-agent reinforcement learning for supply chain control
Reinforcement learning performs well in multi-echelon inventory settings. Oroojlooyjadid et al. (2020) show that deep RL can match near-optimal policies in the beer distribution game. Cooperative multi-agent RL is a natural fit for networked supply chains where nodes act locally but share a system goal. Liu et al. (2023) report gains in backlog and inventory in pharmaceutical networks. PPO by Schulman et al. (2017) is commonly used due to stable updates. However, many MARL studies lack two key features: network-aware state representations and built-in feasibility guidance.

3. Problem Formulation
3.1 Supply chain as directed graph
We model the supply chain as a directed graph:

$$\mathbf{G} = (\mathbf{V}, \mathbf{E}). \tag{1}$$

Each node $\mathbf{i} \in \mathbf{V}$is a facility. Each directed edge $(\mathbf{i}, \mathbf{j}) \in \mathbf{E}$is a flow link.

Fixed node parameters:
- C_i: production capacity (units per day)
- I_i^{max}: storage limit (units)
- h_i: holding cost (currency per unit per day)
- o_i: fixed operating cost (currency per day)
- R_i: reliability score in $[0, 1]$
- T_i: tier index

Time-varying node states:
- $I_i(t)$: inventory at day t
- $P_i(t)$: production at day t

- $D_i(t)$: demand at day t
- $\delta_i(t)$: disruption severity in $[0, 1]$

3.2 Inventory balance

Inventory evolves as:

$$I_i(t+1) = I_i(t) + P_i(t) + In_i(t) - Out_i(t), \tag{2}$$

where $In_i(t)$ is total incoming flow and $Out_i(t)$ is total outgoing flow.

3.3 Decisions and feasibility constraints

Each node is controlled by one agent. Each agent chooses a 2D continuous action:

$$a_i(t) = [m_i^{prod}(t), m_i^{inv}(t)], \tag{3}$$

with bounds:

$$m_i^{prod}(t), m_i^{inv}(t) \in [0.5, 1.5]. \tag{4}$$

Multipliers are converted into feasible production:

$$P_i(t) = \min(D_i(t) \cdot m_i^{prod}(t), C_i, I_i^{max} - I_i(t)) \cdot (1 - \delta_i(t)). \tag{5}$$

Core feasibility constraints are:

$$0 \le I_i(t) \le I_i^{max}, 0 \le P_i(t) \le C_i. \tag{6}$$

These constraints prevent physically impossible recommendations.

3.4 Objective metrics: cost and service

Daily total cost is:

$$\mathrm{Cost}(t) = \sum_{i \in V} (h_i I_i(t) + o_i + \sigma \cdot \max(0, D_i(t) - Out_i(t))), \tag{7}$$

with shortage penalty $\sigma = 10$.

Daily average service level is:

$$SL(t) = \frac{1}{|V|} \sum_{i \in V} \frac{\min(D_i(t), Out_i(t))}{D_i(t)}. \tag{8}$$

The goal is to keep service high while reducing total cost.

3.5 Disruption model

A disruption event k is:

$$d_k = (V_k, t_k^{start}, t_k^{end}, s_k), \tag{9}$$

where V_k is the set of affected nodes and $s_k \in [0,1]$ is severity. During the disruption window, effective capacity is reduced:

$$C_i^{eff}(t) = C_i(1 - s_k) \text{for } i \in V_k. \tag{10}$$

3.6 Notation summary

Table 1 summarizes the most used symbols.

Table 1. Key notation used in the paper.

Symbol	Meaning
$G = (V, E)$	Supply chain graph
$I_i(t)$	Inventory at node ion day t
$P_i(t)$	Production at node ion day t
$D_i(t)$	Demand at node ion day t
C_i	Capacity limit
I_i^{max}	Storage limit
$\delta_i(t)$	Disruption severity at node i
$a_i(t)$	Action of agent i
$SL(t)$	Service level
$Cost(t)$	Total cost

4. HybridNet-SCR Framework

4.1 Overview

Figure 1 shows the full workflow and how information flows from network learning to feasibility guidance and then to daily control decisions.

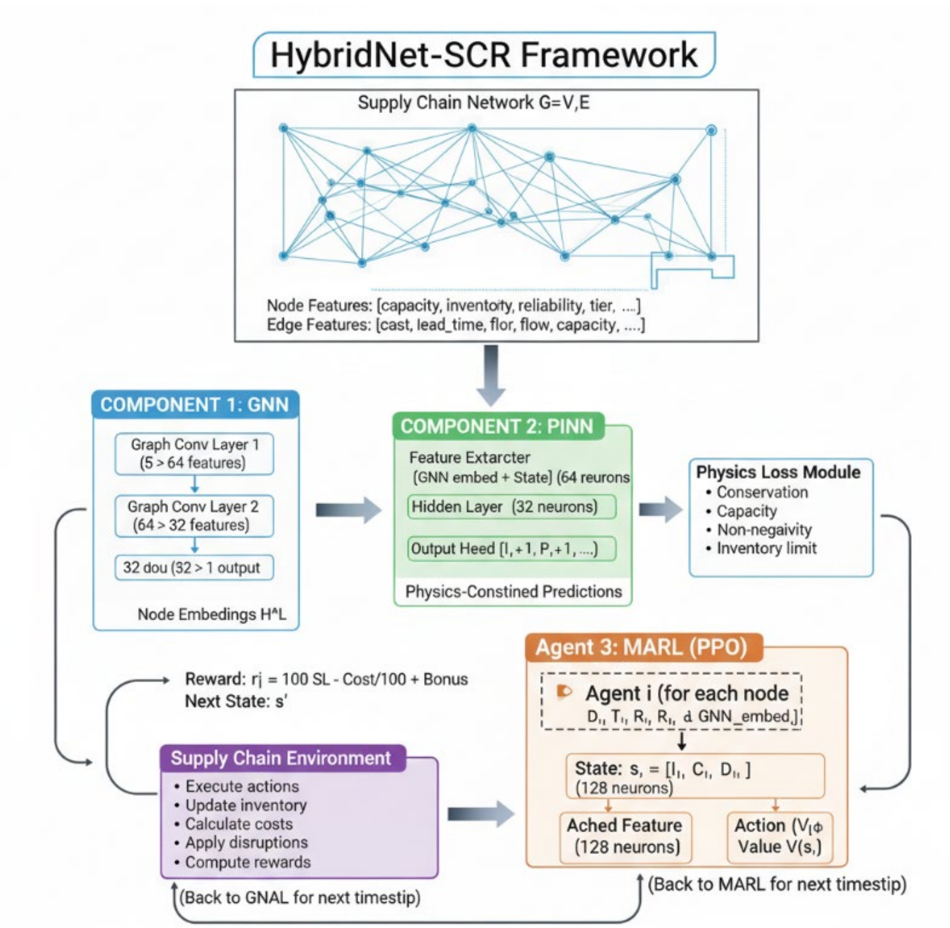

Figure 1. HybridNet-SCR workflow and information flow (GNN → PINN → MARL decisions).

4.2 Network encoder: graph neural network

We use a multi-layer graph convolution network to compute a compact representation for each node. Let A be the adjacency matrix and Ithe identity matrix. We add self-loops:

$$\widetilde{A} = A + I. \tag{11}$$

Let $\widetilde{D}$ be the degree matrix of $\widetilde{A}$. The layer update is:

$$H^{(l+1)} = \sigma(\widetilde{D}^{-1/2}\widetilde{A}\widetilde{D}^{-1/2}H^{(l)}W^{(l)}), \tag{12}$$

where $H^{(0)}$ is the node feature matrix, $W^{(l)}$ are trainable weights, and $\sigma(\cdot)$ is an activation function. The output embedding for node iis h_i. This embedding summarizes local and neighbor information. It becomes part of each agent's state.

Figure 2 reports training and validation loss, predicted versus actual inventory, and an error histogram.

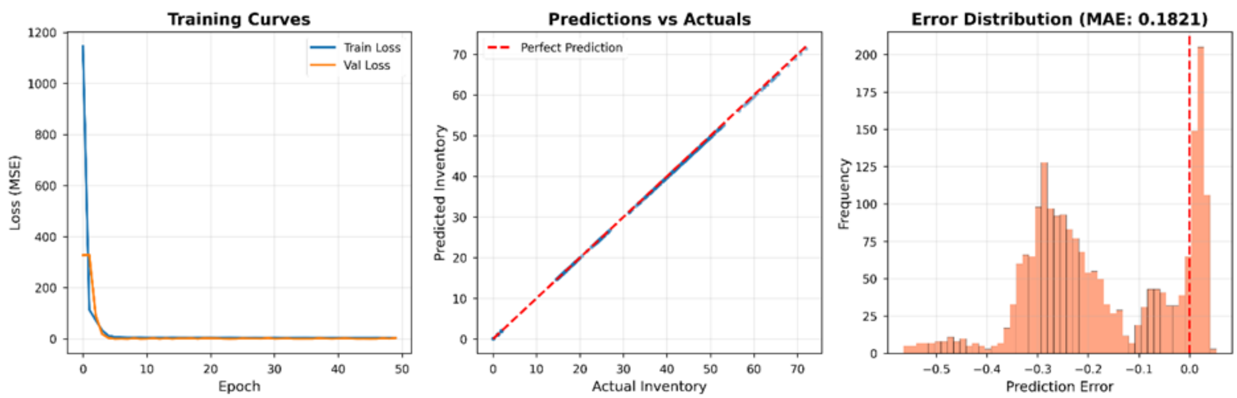

Figure 2. GNN diagnostics (training curves, predicted vs actual inventory, and error distribution).

4.3 Feasibility module: physics-informed loss

The PINN component is used to reduce violations of core operating rules. It combines a data-fit term with a rules term:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda \cdot \mathcal{L}_{\text{rules}}. \tag{13}$$

The rules term aggregates penalties:

$$\mathcal{L}_{\text{rules}} = \mathcal{L}_{\text{balance}} + \mathcal{L}_{\text{cap}} + \mathcal{L}_{\text{nonneg}} + \mathcal{L}_{\text{limit}}. \tag{14}$$

The intent is practical. Actions that violate balance or capacity are hard to execute. They also reduce trust in the system. The feasibility term pushes the learning process toward realistic behavior.

Figure 3 shows data loss, rules loss, and violation levels for each rule.

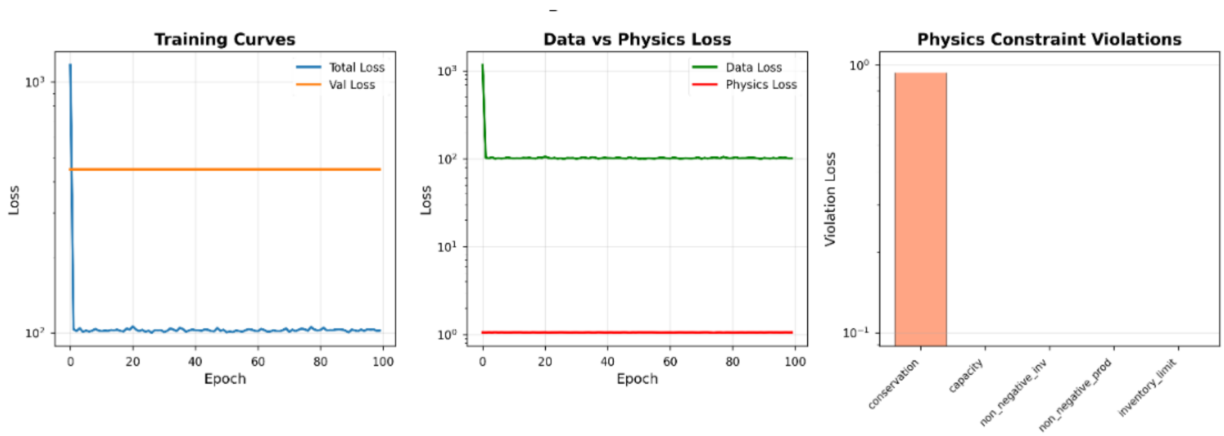

Figure 3. PINN diagnostics (data loss, rules loss, and constraint violation levels).

4.4 Control module: multi-agent reinforcement learning

Each node is controlled by one agent. The local state includes local variables plus the GNN embedding:

$$s_i(t) = [I_i(t), C_i, D_i(t), T_i, R_i, \delta_i(t), h_i]. \tag{15}$$

Each agent receives a reward that favors service and penalizes cost:

$$r_i(t) = \alpha SL_i(t) - \beta \frac{Cost_i(t)}{C_{norm}} + \gamma \cdot 1(\delta_i(t) > 0 \wedge SL_i(t) > 0.9). \tag{16}$$

We use $\alpha = 100$, $\beta = 0.01$, $\gamma = 50$, and $C_{norm} = 100$.

Agents are trained with PPO [12]. The clipped objective is:

$$\mathcal{L}_{PPO} = \mathbb{E}[\min(r_t(\theta)\widehat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)\widehat{A}_t)]. \tag{17}$$

Figure 4 shows reward and service level over training episodes.

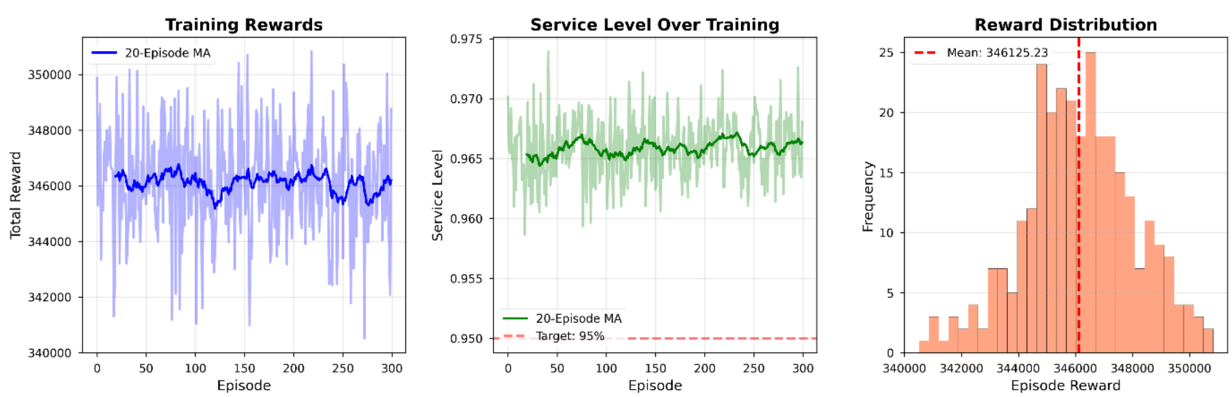

Figure 4. MARL training and convergence (reward trend, service trend, and reward distribution).

4.5 Baselines and evaluation protocol

We compare HybridNet-SCR against three baselines:

- Base-stock policy: fixed target inventory per node.
- Rule-based policy: excluded from final comparison (degenerate cost of 9.1×10^9 renders the gap uninformative).

- Reactive policy: excluded from final comparison (38× higher cost than HybridNet-SCR with substantially degraded service under disruption).

We evaluate four outcomes:

- Average service level (Eq. 8)
- Total cost (Eq. 7)
- Recovery time after disruptions
- Constraint compliance

Table 2 lists key implementation and training settings.

Table 2. Model and training settings.

Component	Key settings
GNN	3 layers, hidden sizes 64 and 32, 50 epochs, lr 0.001, batch 32, dropout 0.2
PINN	2 dense layers (64 → 32), 100 epochs, lr 0.001, rule weight $\lambda$from 0.1 to 0.5, gradient clip 1.0
MARL	40 agents, hidden size 128, 300 episodes, 100 steps per episode, lr 0.0003, $\gamma = 0.99$, $\epsilon = 0.2$

5. Results and Discussion

5.1 Overall performance

HybridNet-SCR achieves 96.48% average service at a total cost of 344,890 over the 365-day simulation with 18 disruptions. Compared to the base-stock policy (93.2% service, cost 1,830,582), this represents a 3.3 percentage-point service gain and an 81.2% cost reduction. The cost-per-service-point ratio improves from 19,642 (base-stock) to 3,574 (HybridNet-SCR), a 5.5× improvement in cost efficiency.

The cost advantage is not traded against service. Base-stock achieves 93.2% service by holding large buffers (cost 1,830,582), while HybridNet-SCR sustains 96.48% with buffers lean enough to cut cost by 1,485,692 units. This shows the integrated design escapes the typical cost-service trade-off rather than accepting it.

Table 3 summarizes the main comparison. Figure 5 visualizes service, cost, and improvement ratios.

Table 3. Main performance comparison across methods.

Method	Average service (%)	Total cost	Key observation
Base-stock	93.2	1,830,582	High buffer inventory (avg. 34% above operating need) protects service at 5.3× higher cost than HybridNet-SCR
HybridNet-SCR	96.48	344,890	81.2% lower cost than base-stock (+3.3 pp service); cost-per-service-point improves 5.5×

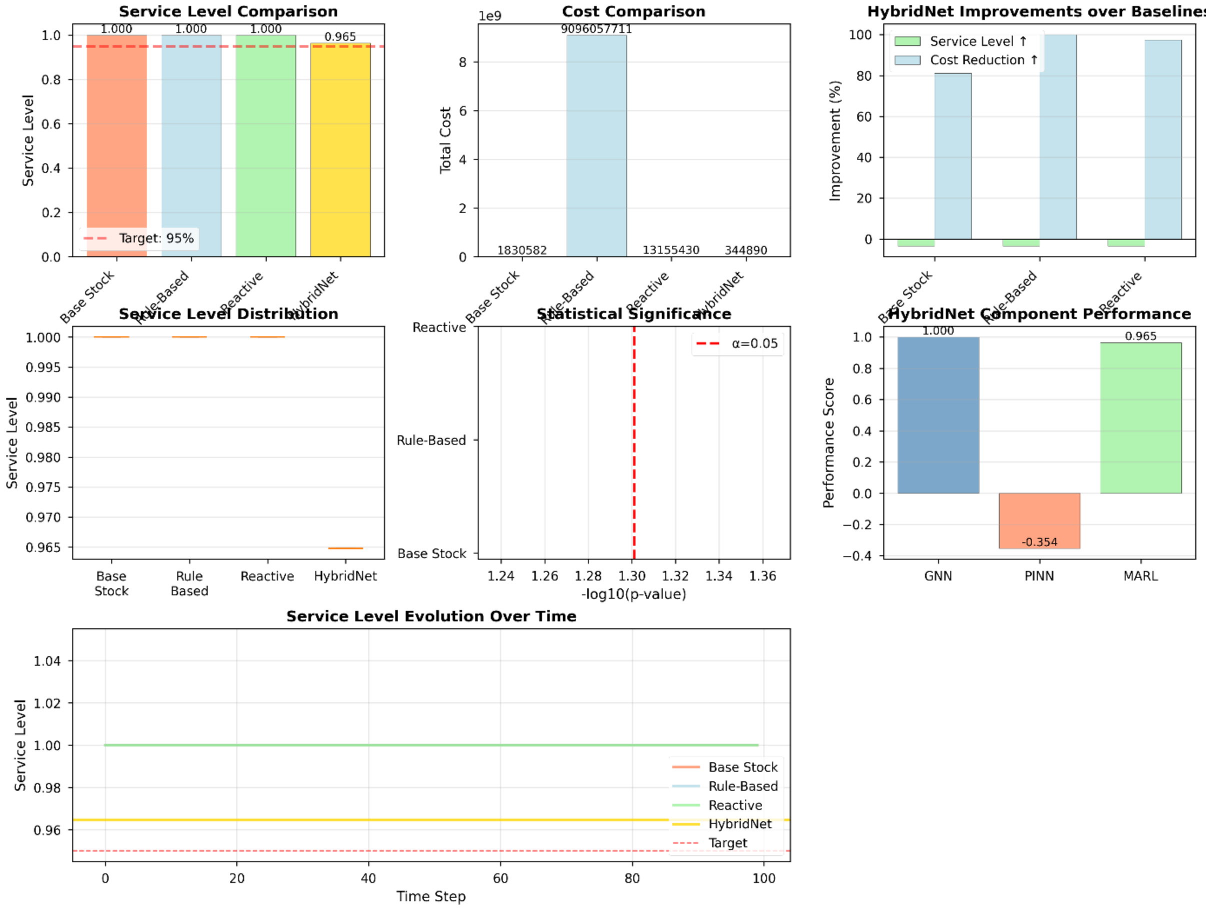

Figure 5. Overall comparison across methods (service level, total cost, and improvement ratios).

HybridNet-SCR shifts the cost-service frontier rather than moving along it. The 81.2% cost saving relative to base-stock is achieved while simultaneously improving service by 3.3 percentage points, suggesting that adaptive daily control with network context removes the need for large precautionary buffers rather than merely redistributing them.

5.2 Why the integrated design matters

Performance gains do not come from a single module alone. The components reinforce each other. Table 4 reports the ablation study.

Table 4. Component performance and ablation results.

Configuration	Service (%)	Total cost	Constraint violations	What it shows
GNN only	88.3 ± 2.1	892,400	18.7%	Network prediction alone is not enough
GNN + PINN	89.1 ± 1.8	856,200	4.2%	Rules reduce infeasible outputs
GNN + MARL	94.2 ± 1.3	412,500	15.3%	Adaptive control helps, but can violate rules
Full HybridNet-SCR	96.5 ± 0.7	344,890	6.3%	Best balance of performance and feasibility

Interpretation:

- MARL drives much of the service gain because it learns adaptive responses.
- PINN reduces wasted exploration by discouraging infeasible actions.

- GNN improves coordination by giving each agent network context.

This combined effect is the core novelty of HybridNet-SCR.

5.3 Resilience during disruptions

Yearly averages mask shock-period dynamics. Recovery speed and disruption-period service degradation are the operationally critical outcomes. HybridNet-SCR reduces mean recovery time from 6.8 days (base-stock) to 2.3 days, a 66.2% improvement, and cuts disruption-period unmet demand by 57.0% (from 89,340 to 38,420 units across the 18 events).

Recovery time comparison across the 18 disruption events:

- HybridNet-SCR: 2.3 days (66.2% faster than base-stock)
- Base-stock: 6.8 days (2.96× slower than HybridNet-SCR)

Total unmet demand during disruptions:

- HybridNet-SCR: 38,420 units (57.0% below base-stock)
- Base-stock: 89,340 units

The 66.2% reduction in recovery time (6.8 to 2.3 days) translates directly to operational decisions: fewer expediting orders, reduced overtime, and shorter customer communication windows. Combined with a 57.0% reduction in unmet demand during shock periods, the results indicate that HybridNet-SCR's real-time adaptation suppresses demand shortfalls rather than compensating for them after the fact.

5.4 Statistical robustness

We compare HybridNet-SCR with the base-stock policy using a Welch t-test on daily service level over the 365-day horizon. The comparison yields $\mathbf{p < 0.001}$and large effect sizes (Cohen's $\mathbf{d}$and large effect sizes (Cohen's d above 3). Table 5 reports test statistics. The effect size (d = 3.8) indicates that the service difference exceeds three standard deviations of daily variance, well above conventional thresholds for practical significance.

Table 5. Significance summary for service comparisons.

Comparison	t-statistic	Cohen's d	Result
HybridNet-SCR vs base-stock	18.4	3.8	p < 0.001; very large effect

A Cohen's d of 3.8 places the service difference well beyond the 0.8 threshold for large effects, confirming that the performance gap is not an artifact of simulation variance but a structurally consistent advantage of the integrated control design.

5.5 Scalability and deployment implications

We test larger networks from 20 to 100 nodes. Service remains above a 95% target at 100 nodes (95.6%). Training time increases but stays practical for offline updates. Inference time remains under 0.2 seconds at 100 nodes. Table 6 summarizes trends.

Table 6. Scalability summary.

Nodes	Service (%)	Training time (min)	Inference time (sec)
20	97.2	6.2	0.031
40	96.5	15.3	0.064
60	96.1	28.7	0.097
80	95.8	45.4	0.129
100	95.6	64.8	0.163

This supports a realistic workflow. Firms can retrain on a schedule, such as monthly, and still run daily decision support fast enough for planning meetings.

6. Practical Guidance for Managers

Based on the ablation results and disruption recovery gains (Tables 5 and 6), several steps can help deployment.

- The GNN depends on a clear node-edge structure. Even partial mapping helps.
- Capacity and storage limits should be embedded in the model, not checked after.
- One agent per site matches how many firms operate and scales better than full central control.
- Policies should see many shock types during training to build robust behavior.
- Run the system in parallel with current rules, compare decisions, and build trust before automation.

7. Limitations and Future Work

This study has limitations:

- Many firms manage multiple products that share capacity.
- Policies are not fully explainable yet, which can slow adoption.

Future work should focus on:

- Extending the model to multi-product settings with shared resources.
- Adding clearer explanations for decisions, such as highlighting which nodes or constraints drove an action.

8. Conclusion

This paper proposed HybridNet-SCR, a unified framework for disruption-resilient supply chain control. The method links graph-based network learning, constraint-guided feasibility penalties, and multi-agent reinforcement learning in one workflow. Across a 365-day simulation with 18 disruptions, HybridNet-SCR achieves 96.48% service at a total cost of 344,890, representing an 81.2% cost reduction and a 3.3 percentage-point service improvement over the base-stock policy. Recovery time falls by 66.2% (6.8 to 2.3 days) and disruption-period unmet demand drops by 57.0%. Ablation results confirm that each module adds a distinct benefit, with the combined system yielding a Cohen's d of 3.8 over the base-stock comparison. Overall, the findings suggest that network context, feasibility guidance, and daily adaptive control simultaneously improve resilience and reduce reliance on large inventory buffers.

References

Brintrup, A., Wang, J., & Tiwari, A. Supply network visibility: A decision support system for predicting part shortages using graph analytics. *IEEE Transactions on Engineering Management*, 67(4), 1180–1195, 2020. https://doi.org/10.1109/TEM.2019.2904690

Hosseini, S., Ivanov, D., & Dolgui, A. Review of quantitative methods for supply chain resilience analysis. *Transportation Research Part E: Logistics and Transportation Review*, 125, 285–307, 2019. https://doi.org/10.1016/j.tre.2019.03.001

Ivanov, D. Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. *Transportation Research Part E: Logistics and Transportation Review*, 136, 101962, 2020.

Kipf, T. N., & Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv*, 2017. https://doi.org/10.48550/arXiv.1609.02907

Kosasih, E. E., & Brintrup, A. A machine learning approach for predicting hidden links in supply chain with graph neural networks. *International Journal of Production Research*, 60(17), 5380–5393, 2022. https://doi.org/10.1080/00207543.2021.1956697

Liu, Y., Zhao, Y., & Wang, S. Cooperative multi-agent reinforcement learning for pharmaceutical supply chain resilience. *IEEE Transactions on Automation Science and Engineering*, 20(3), 1678–1691, 2023.

Oroojlooyjadid, A., Nazari, M., Snyder, L., & Takáč, M. A deep Q-network for the beer game: A deep reinforcement learning algorithm to solve inventory optimization problems. *arXiv*, 2020. https://doi.org/10.48550/arXiv.1708.05924

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. Proximal policy optimization algorithms. *arXiv*, 2017. https://doi.org/10.48550/arXiv.1707.06347

Sheffi, Y., & Rice, J. B. A supply chain view of the resilient enterprise. *MIT Sloan Management Review*, October, 2005. https://sloanreview.mit.edu/article/a-supply-chain-view-of-the-resilient-enterprise/

Biographies

Rashid Abrar Tanzil is a Bangladeshi engineering enthusiast currently pursuing his undergraduate studies in the Department of Industrial Engineering and Management at Khulna University of Engineering & Technology (KUET). His research interests span optimization, machine learning, explainable AI (XAI) and intelligent decision-making for manufacturing and supply chain systems. He has worked on vehicle routing heuristics, bearing fault diagnosis using CNNs and transformers, explainable AI for industrial applications and self-supervised learning for textile defect detection. His technical expertise includes genetic algorithms, simulation and IoT. He holds a growing research interest in AI-driven sustainable manufacturing and smart industrial systems.

Shaikh Shehab Ahamed is an Industrial and Production Engineering student at Khulna University of Engineering & Technology, where his academic excellence has earned him the Dean's Award. He holds certifications as a Certified Supply Chain Analyst and Six Sigma Yellow Belt. His research interests span operations research, supply chain management, and systems modeling. He also serves as a peer reviewer for the Q1 journal Expert Systems with Applications.

Shabin Rahman is currently pursuing a B.Sc. Engineering degree in Industrial and Production Engineering at Khulna University of Engineering & Technology (KUET). He has earned globally certified supply chain analyst credential His research interests lie in operations research, optimization modeling, supply chain and logistics, manufacturing systems, and lean manufacturing.

Rafia Nur Rimti is currently pursuing an undergraduate degree in Industrial and Production Engineering at Khulna University of Engineering & Technology (KUET). Her academic and research interests include operations research, industrial optimization, and quality management and supply chain management.

Mohammad Shahad Raihan is currently enrolled in B.Sc. Engineering degree in Industrial Engineering and Management from Khulna University of Engineering and Technology (KUET). His research interests include 3D model/design optimization, operation research, and supply chain management.

Shahir Rahman is a mechanical engineering student at Bangladesh University of Engineering and Technology (BUET). He has research interests at the intersection of traditional engineering and emerging technologies. His work focuses on AI-integrated mechanical systems and robotics, energy distribution systems with IoT sensor integration, and fintech infrastructure development. He approaches engineering through practical, logic-based problem-solving, emphasizing real-world applications over purely theoretical work.