DOI: https://doi.org/10.15276/aait.09.2026.35
UDC 004.942:519.17

Supply chain resilience under stochastic disruptions

Stanislav V. Pohorielov[1)]
ORCID: https://orcid.org/0000-0002-0189-8655; Stanislav.Pohorielov@khpi.edu.ua. Scopus Author ID: 9270293800
Oleksii V. Yanovskyi[1)]
ORCID: https://orcid.org/0009-0002-4310-2843; Oleksii.Yanovskyi@cs.khpi.edu.ua
[1)] National Technical University "Kharkiv Polytechnic Institute", 2 Kyrpychova St., Kharkiv, 61002, Ukraine

ABSTRACT

The **relevance** of this research is determined by the growing vulnerability of supply chain networks to stochastic disruptions caused by local failures, capacity losses, and clustered damage. The **objective of the study** is to develop and validate a compact graph-flow model and a reproducible computational framework for assessing structural and functional resilience under uncertainty. The proposed approach represents a supply chain as a directed capacitated graph and defines network performance through normalized maximum flow from supply sources to demand nodes. The methodology combines stochastic damage generation, including node failures, edge degradation, cluster disruptions, and irreversible losses, with a baseline discrete recovery process and Monte Carlo simulation. Resilience is evaluated by a time-dependent performance trajectory, an integral resilience index, robustness, final recovery level, median resilience, interquartile range, and a lower-tail conditional risk metric. Computational experiments compare hierarchical, hub, redundant, and dense network archetypes across different disruption types, redundancy levels, network sizes, metric parameters, and damage irreversibility settings. The results show that dense topology provides the highest typical and lower-tail resilience, whereas hub topology may remain effective in median scenarios but demonstrates weaker behavior in adverse cases. Redundancy is found to be especially important for reducing tail risks rather than only improving typical performance. The analysis also shows that cluster disruptions are not automatically more destructive than random node failures when the actual damage dose is normalized, while irreversibility sharply reduces recovery completeness and lower-tail resilience. The scientific novelty of the study lies in integrating flow-based functionality, stochastic disruption modeling, recovery dynamics, and risk-oriented metrics into a single benchmark procedure. **The practical value** of the framework is its applicability as a basis for comparing supply chain network structures and for developing future prioritized, decentralized, and multi-agent recovery algorithms.

Keywords: Network resilience; supply chains; graph models; stochastic disruptions; maximum flow; network recovery

For citation: **Pohorielov S. V., Yanovskyi O. V. "Supply chain resilience under stochastic disruptions". *Applied Aspects of Information Technology*. 2026; Vol.9 No.4: 547–560. DOI: https://doi.org/10.15276/aait.09.2026.35**

1. INTRODUCTION

Modern supply chain networks are complex distributed systems in which material and information flows pass through many interconnected suppliers, warehouses, transport nodes, and consumers. The growth of global interdependence, the concentration of logistics hubs, and the drive to minimize inventories increase the efficiency of such systems under normal conditions, but at the same time increase their sensitivity to disruptions. A local failure of a node or channel may cause a disproportionately large decline in the functionality of the entire network if the damaged element is a flow bottleneck.

In the literature, supply chain resilience is usually associated with the ability of a system to prepare for unexpected events, respond to them, and recover after disruptions [1], [2]. For computational modeling tasks, this conceptual definition is insufficient: a formal procedure is required that makes it possible to compare topologies, damage scenarios, and recovery parameters on a single quantitative scale [3]. In this study, such a procedure is a graph flow model in which network functionality is defined by the maximum flow, while resilience is represented by the trajectory of normalized performance after a stochastic disruption.

The problem is that simple topological connectivity does not guarantee the functional operability of a supply chain network. A network may remain connected but lose a significant share of its throughput because of damage to critical edges or hub nodes. In addition, the stochastic nature of disruptions makes the analysis of a single scenario insufficient: different realizations of damage may produce substantially different performance trajectories. Therefore, comparison should account not only for typical behavior but also for the lower tail of the distribution, that is, the worst-case scenarios for which CVaR-like metrics are appropriate [9], [10], [11].

The aim of the article is to develop and experimentally validate a graph flow model and a computational framework for assessing the structural and functional resilience of supply chain networks under stochastic disruptions. The main contribution lies in combining a directed capacitated network,

© Pohorielov S., Yanovskyi O., 2026

This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/deed.uk)

random damage models, baseline recovery dynamics, a Monte Carlo protocol, and a system of risk-oriented metrics. Unlike studies that analyze only structural metrics or only average resilience, this article proposes a flow-stochastic benchmark with lower-tail assessment and verifies how topology, redundancy, disruption type, scale, and irreversibility of damage affect the typical and tail resilience of a network.

1.1. Relevance and problem context

The relevance of the study is determined by the fact that modern supply chains increasingly operate as distributed networked systems rather than as linear sequences of suppliers and consumers. Production globalization, the concentration of transport corridors, dependence on logistics hubs, and the reduction of safety stocks increase efficiency in normal conditions but simultaneously create structural sensitivity to disruptions. In such systems, a local disturbance may propagate through flow dependencies and become a system-level loss of service.

For this reason, supply chain resilience should be treated not only as a managerial or economic concept, but also as a problem of mathematical and computational modeling of a complex network under uncertainty. A formal computational formulation is needed because qualitative statements about preparedness, response, and recovery cannot directly answer how much functionality is preserved after a stochastic failure, which topology is more stable, or how strongly the worst scenarios differ from typical ones.

The research problem addressed in this article is therefore the lack of a compact, reproducible, and functionally interpretable procedure for comparing supply chain network structures under random disruptions. Pure connectivity indicators are insufficient: a network may remain connected while losing a considerable part of its throughput because a critical edge or a hub-like node has been degraded. Conversely, a strong initial drop may be less dangerous if the network can recover quickly within the modeled horizon.

1.2. Contribution of the study

The contribution of the article is fourfold. First, it formalizes the supply chain as a directed capacitated graph in which the main operational function is expressed by maximum flow from sources to demand nodes. Second, it introduces a stochastic disruption protocol that includes node failures, edge degradation, cluster damage, and irreversible losses. Third, it combines a baseline recovery mechanism with Monte Carlo simulation so that resilience is assessed as a distribution rather than as a single deterministic scenario. Fourth, it applies median, variability, and lower-tail metrics, which makes it possible to distinguish typical resilience from tail vulnerability.

This positioning deliberately keeps the model at an intermediate level of abstraction. It is more informative than a purely topological graph description, because it accounts for capacities and bottlenecks, but it is less parameter-heavy than a full economic simulation of inventories, multi-product flows, contractual behavior, and time-dependent demand. Such a level of abstraction is useful for a benchmark framework: it isolates structural and functional effects and prepares a basis for later controlled recovery algorithms.

2. LITERATURE REVIEW AND PROBLEM STATEMENT

Research on supply chain resilience develops at the intersection of supply chain management, risk theory, operations research, network science, and computational modeling. Ponomarov and Holcomb interpret resilience as the ability to prepare for unexpected events, respond to them, and recover after them [1], while Pettit, Fiksel, and Croxton consider it through the balance between system vulnerabilities and capabilities [2]. These studies form the conceptual foundation but do not provide a direct algorithm for comparing specific topologies in terms of flow functionality.

Quantitative reviews emphasize that formal resilience analysis can be based on optimization, simulation, system dynamics, agent-based modeling, and network methods [3]. Recent reviews from 2024-2025 further emphasize the design of resilient supply chain networks, network modeling, simulation-based digital strategies, the alignment of resilience-capacity metrics, and the combination of risk mitigation, robust optimization, and technological solutions [4], [5], [6], [7], [8]. Three directions are important for this article. The first is the graph representation of supply chain networks, where nodes correspond to infrastructure entities and edges represent logistics or information links [5], [12], [13]. The second is flow models, since the existence of a route alone does not imply sufficient capacity; maximum flow is a classical problem of combinatorial optimization [14], [15]. The third is simulation modeling of disruptions, including the ripple effect, in which a local disruption may cause systemic losses [6], [13], [14].

In studies of critical infrastructure resilience, the idea of a time-dependent performance function is widely used: after a disruption, functionality decreases and then partially or fully recovers; the integral area under this curve provides a quantitative estimate of resilience [7], [15], [16], [17]. However, for stochastic supply chain networks it is important to analyze not only average or median behavior but also the distribution of outcomes. CVaR, proposed in risk management as a characteristic of the tail of a distribution [9], [10], is adapted in this article to the lower tail of resilience: since larger metric values are preferable, the smallest values of integral resilience are of interest.

The research gap consists in the insufficient integration of these components into a single reproducible procedure: graph topology, capacities, stochastic damage, recovery dynamics, and tail risks are often analyzed separately. The proposed framework does not claim to be a complete economic model of a supply chain and does not solve the problem of optimal recovery; its role is to create a baseline computational framework on which structural network properties can be compared and decentralized and multi-agent recovery algorithms can subsequently be built.

2.1. Resilience as a dynamic network property

In the context of this study, resilience is interpreted as a dynamic structural and functional property of a network. Reliability, robustness, and recoverability are related but not identical notions. Reliability is associated with the probability of maintaining a required function over time; robustness characterizes the depth of the initial performance drop; recoverability reflects the ability to return to an acceptable state after disruption. Resilience integrates these aspects by considering the whole trajectory of performance loss and recovery.

This distinction is important for supply chain networks because the same final state can correspond to different operational histories. Two networks may reach the same recovery level at the end of the simulation, but one of them may have spent most of the horizon at a low performance level. Similarly, a network may have a high median resilience while remaining exposed to rare scenarios with severe degradation. A time-dependent and stochastic formulation is therefore necessary.

2.2. Graph and flow-based representation

A graph representation is natural for supply chains because suppliers, warehouses, distribution centers, transport nodes, and consumers can be modeled as vertices, whereas logistics or information links can be modeled as directed edges. This representation makes it possible to use graph algorithms, analyze paths, evaluate structural redundancy, and identify potentially critical elements. However, structural graph metrics alone do not determine whether the network can satisfy demand after a disruption.

The flow interpretation used in this article addresses this limitation. The existence of a path between a source and a consumer is not enough if the path has insufficient capacity. Maximum flow provides an aggregate functional measure that accounts for both route availability and throughput constraints. In this sense, the model treats the supply chain not merely as a connected graph, but as a capacitated system that must transmit resources through limited channels.

2.3. Stochastic disruptions and ripple effects

Disruptions in supply chains are stochastic by nature. The location of a failure, the intensity of damage, the affected elements, and the possibility of subsequent recovery are typically uncertain. Therefore, the analysis of one fixed damage scenario may lead to unstable conclusions. A Monte Carlo protocol is more appropriate because it generates a distribution of possible outcomes and allows the researcher to compare not only average behavior but also adverse realizations.

The ripple effect provides an additional motivation for a network approach. A local failure may change the performance of distant parts of the system through flow dependencies, capacity bottlenecks, and the loss of alternative routes. In a graph-flow model, this effect appears when the removal or degradation of one element reduces the maximum flow of the entire network. Thus, even without explicitly modeling inventories or firm-level behavior, the framework captures an important class of systemic consequences.

2.4. Research gap

The literature contains conceptual definitions of supply chain resilience, quantitative reviews, network models, simulation approaches, and risk measures. Nevertheless, these components are often applied separately. What is still needed is a single

reproducible procedure that combines graph topology, capacities, stochastic damage, recovery dynamics, and lower-tail risk into one experimental pipeline. This article addresses that gap by proposing a benchmark that can be extended toward optimization, decentralized decision making, and multi-agent recovery.

3. MATERIALS AND METHODS

The initial state of the supply chain network is specified by a directed graph with capacities:

$$G_0 = (V, E, c), \tag{1}$$

where V is the set of nodes, E is the set of directed edges, and c: $E \to R_+$ defines the capacity of each edge. The model distinguishes the sets of sources $S \subset V$ and consumers $D \subset V$. To transform the problem with multiple sources and multiple consumers into a standard maximum-flow problem, a supersource s^* and a supersink t^* are introduced.

The functionality of the current network state $G(t)$ is defined as the maximum flow $F(G(t))$ in the extended graph [14], [15]:

$$\begin{aligned} F(G) &= \max_f \sum_{(s^*,v)\in E^+} f_{s^*v}, \\ 0 &\le f_{ij} \le c_{ij} \\ \sum_i f_{iv} &= \sum_j f_{vj} \end{aligned} \quad . \tag{2}$$

A disruption is modeled as a random scenario ξ that transforms the undisturbed graph G_0 into a damaged state G_ξ. The scenario may include node failures, edge degradation, cluster damage, and a set of irreversibly damaged elements.

For an edge e that has undergone degradation, its capacity is reduced according to the rule:

$$\begin{aligned} c_\xi(e) &= (1-\gamma_e)c(e), \\ 0 &\le \gamma_e \le 1. \end{aligned} \tag{3}$$

Cluster damage is generated around a random epicenter, but in the experimental protocol its intensity is normalized through the target share of damaged nodes, which enables a correct comparison with random node failures. After the disruption, a baseline discrete recovery process is launched: at each time step, no more than one available damaged element can be restored. This strategy is not optimal; it is used as a neutral baseline procedure for assessing the influence of topology and disruption parameters.

The basic time-dependent characteristic is the normalized performance function:

$$Q(t) = \frac{F(G(t))}{F(G_0)},\ t = 0,1,\ldots,T, \tag{4}$$

where $G(0) = G_\xi$ is the state immediately after the disruption. The integral resilience of a single scenario is defined as the discrete area under the performance curve:

$$\rho(\xi) = \frac{1}{T+1}\sum_{t=0}^{T} Q(t). \tag{5}$$

To describe the shape of the trajectory in more detail, component metrics are used: robustness $R_{rb} = \min_t Q(t)$, final recovery level $R_{inf} = Q(T)$, and, where necessary, recovery rate after the moment of minimum performance. Since disruptions are stochastic, a sample of resilience values $\{\rho_1, \ldots, \rho_M\}$ is formed over M runs. Typical resilience is described by the median ρ_{med}, variability by the interquartile range ρ_{IQR}, and lower-tail risk by the metric:

$$\rho_{CVaR_\alpha} = \frac{1}{k}\sum_{i=1}^{k} \rho_{(i)},\ k = \lceil \alpha M \rceil, \tag{6}$$

where $\rho_{(i)}$ are the values of integral resilience are ordered in ascending order. Thus, the metric ρ_{CVaR_α} is the average resilience in the worst α-fraction of scenarios. In the sensitivity subseries, a composite performance function is also used:

$$Q(t;\omega) = \omega\frac{F(t)}{F(0)} + (1-\omega)\frac{LCC(t)}{N}, \tag{7}$$

where $LCC(t)/N$ is the share of nodes in the largest connected component, and lambda specifies the weight of the flow component. For statistical comparisons, the Kruskal-Wallis test, pairwise Mann-Whitney tests with Bonferroni correction, bootstrap confidence intervals, and Cliff's effect size are applied [18], [19], [20], [21].

3.1. Modeling assumptions

The proposed model treats each supply chain configuration as a directed capacitated network with distinguished source and demand nodes. The direction of an edge represents the permitted direction of material flow. If a real logistics connection is bidirectional, it can be represented by two directed edges with possibly different capacities. This convention preserves the generality of the

model while keeping compatibility with standard maximum-flow algorithms.

The model is single-commodity: different products, transport modes, and contractual relations are aggregated into one generalized flow. This simplification is intentional. It makes the role of topology, bottlenecks, and stochastic damage transparent, and it prevents the comparison of network structures from being obscured by many economic parameters. The approach therefore should be interpreted as a structural-functional benchmark rather than as a complete operational model of a specific company.

3.2. Disruption types

Node failure represents the loss of an infrastructure object such as a supplier, warehouse, logistics hub, or distribution node. When a node fails, it is removed from the functional part of the graph, and all incident edges become unavailable for flow transmission until recovery is possible. This type of disruption is especially important because node removal can simultaneously affect many routes and may disconnect supply from demand even when several edges remain intact.

Edge degradation represents a partial loss of transportation or communication capacity. Unlike full edge removal, degradation preserves the topological structure of the graph but reduces the amount of flow that can pass through the affected channel. This makes it suitable for modeling congestion, reduced transport availability, temporary restrictions, or partial service loss. The experiments show that the effect of edge degradation depends on whether the network has enough spare capacity and alternative routes.

Cluster disruption represents a localized event that affects a group of neighboring nodes. In practical terms, this may correspond to a regional shock, a disruption of a logistics area, or another event that damages several topologically close elements. The experimental protocol normalizes cluster intensity by the target share of damaged nodes, which is important because the same graph radius may cover substantially different numbers of nodes in sparse and dense topologies.

3.3. Recovery dynamics

The baseline recovery process is intentionally simple: at each discrete time step, no more than one available damaged element can be restored. This rule creates a neutral reference policy that does not optimize the repair order. Its purpose is to reveal the influence of network structure and damage characteristics before introducing more sophisticated recovery algorithms. Elements marked as irreversible are excluded from the set of feasible recovery actions during the modeled horizon.

This baseline should not be interpreted as a prescription for real-world repair planning. Rather, it is a controlled environment that produces comparable performance trajectories for different network configurations. Future work can replace the random baseline with priority-based, centrality-based, flow-criticality-based, or multi-agent strategies and then measure the improvement over the benchmark presented in this article.

3.4. Interpretation of metrics

The normalized performance function makes it possible to compare networks with different absolute capacities. The integral resilience index summarizes the entire trajectory, while robustness captures the lowest performance level and final recovery captures the state at the end of the horizon. These component metrics are necessary because the same integral value may result from different recovery shapes: a sharp drop followed by quick recovery, or a moderate but long-lasting degradation.

The lower-tail metric has a special role. Because larger resilience values are better, the relevant risk is located in the lower tail of the resilience distribution. The CVaR-like measure used here reports the average resilience among the worst scenarios. This interpretation is particularly useful for critical supply chain networks, where a design that performs well on average may still be unacceptable if its worst cases involve severe loss of functionality.

4. COMPUTATIONAL FRAMEWORK AND EXPERIMENTAL DESIGN

The framework implements the same pipeline for all series: topology generation, calculation of the baseline maximum flow, realization of a stochastic disruption, step-by-step recovery, calculation of $G(t)$ metric computation, and statistical aggregation. The modular structure separates network generators, damage models, flow computation, the recovery protocol, statistical processing, and visualization. This approach corresponds to the principles of reproducible simulation research and the contemporary emphasis on simulation-based analysis of disruptions in supply chains [6], [22], [23], [24], [25].

The study uses four topological archetypes: a hierarchical network, a redundant network with additional alternative edges, a hub network with flow concentration around central nodes, and a dense network with a larger number of alternative

routes. In the baseline series, $N=50$ and $T=10$ are used; in the scaling analysis, *N* varies over the set $\{30,50,100\}$. The baseline number of runs is 100 per configuration, while 1600 runs with a paired seed design are used for the irreversibility analysis S5.D.

Table 1.**System of experimental stages in the article**

Stage	Varied factor	Purpose
S1	Topology type: hierarchical, hub, reserved, dense	Baseline comparison of network structures
S2	Redundancy level r	Assessment of the effect of redundant links
S3	Disruption type and intensity	Comparison of node failure, edge degradation, and cluster disruption
S4	Network size *N*	Verification of scalability of the topology ranking
S5	Metric parameters and irreversible damage share $\alpha, T_{eff}, \omega, w$	Analysis of metric sensitivity and damage irreversibility

***Source*: compiled by the authors**

The concise pseudocode of one run is as follows: generate G_0, compute $F(G_0)$, generate ξ, form $G(0)=G_{\xi}$, for each $t=0$, ..., T compute $F(G(t))$ and $Q(t)$, perform one available recovery step, and after completion calculate integral resilience ρ, robustness R_{rb}, and the final recovery level R_{inf}. If $F(G_0)=0$ or if all active sources or consumers are destroyed after the disruption, the run is marked as invalid and is not included in the aggregated metrics, but it is retained in the raw results.

4.1. Topology generation and experimental control

The computational experiments use controlled topological archetypes rather than empirical networks of a specific enterprise. This choice is methodologically useful because it isolates structural effects. A hierarchical topology represents multilevel supply organization with potential bottlenecks between levels. A reserved topology adds alternative links to such a structure. A hub topology concentrates flows around a limited number of central nodes. A dense topology increases the number of alternative paths and reduces dependence on individual elements.

All topology generators are parameterized so that the same experimental logic can be repeated for different network sizes and disruption intensities. This is essential for reproducibility: conclusions should not depend on a single randomly generated graph. The framework therefore treats topology type, network size, redundancy level, disruption model, recovery horizon, and random seeds as explicit experimental parameters.

4.2. Monte Carlo protocol and data output

For a fixed configuration, the simulation is repeated many times. Each run generates or selects a network, calculates the baseline maximum flow, samples a disruption scenario, records the damaged state, performs step-by-step recovery, and stores the full normalized performance trajectory. Runs for which the baseline maximum flow is zero or for which all active sources or all active consumers are destroyed are flagged as invalid for normalized aggregation but remain available in raw results for transparency.

The raw output of the framework includes run-level metrics and time series. Aggregated files summarize median resilience, lower-tail resilience, interquartile range, robustness, final recovery, and the share of non-recovered runs. Such separation between raw and aggregated data supports auditability: each reported table and figure can be traced back to the underlying Monte Carlo realizations.

4.3. Statistical validation

The use of nonparametric tests is motivated by the shape of simulation outcomes. Resilience values are bounded, can be asymmetric, and may contain extreme scenarios. For this reason, the Kruskal-Wallis test is used for global comparisons, the Mann-Whitney test is applied for pairwise comparisons, and Bonferroni correction controls the family-wise error rate. Cliff's delta complements p-values by indicating the practical magnitude and direction of stochastic dominance.

4.4. Reproducibility and quality control

The reproducibility of the computational study is supported by separating the experimental configuration from the program logic. Each series is defined by a set of parameters: topology class,

network size, disruption type, disruption intensity, recovery horizon, number of Monte Carlo runs, and random seed structure. This separation makes it possible to rerun an entire series or to modify only one factor while keeping all other assumptions unchanged.

Quality control is implemented at two levels. At the model level, each generated network is checked for a positive baseline maximum flow so that the normalized performance function is mathematically meaningful. At the data level, every run stores validity flags and exclusion reasons. This prevents silent removal of problematic cases and makes it possible to distinguish invalid normalizations from genuine low-resilience scenarios.

The paired-seed design used in the irreversibility subseries is especially important. By applying the same random seed grid for all values of the irreversible damage share, the comparison isolates the effect of irreversibility from the effect of the initial stochastic shock. As a result, changes in median resilience, lower-tail resilience, final recovery, and the non-recovered share can be interpreted as consequences of limited recoverability rather than as random differences between disruption realizations.

5. RESEARCH RESULTS

5.1. Effect of topology

Stage 1 showed that topology is a statistically significant factor of resilience: the Kruskal-Wallis test for the distributions of integral resilience ρ produced a statistically significant result $H = 164.3$ ($\rho < 0.001$). The main indicators are presented in Table 2. The dense topology has the highest values of median resilience $\rho_{med} = 0.9900$ and lower-tail resilience $\rho_{CVaR} = 0.9657$, as well as the lowest variability by interquartile range $IQR = 0.0092$ (Fig. 1). The hub structure ranks second in median resilience but has a noticeably worse lower tail, which indicates hidden vulnerability under adverse scenarios. The hierarchical topology demonstrates the lowest robustness.

Table 2. **Descriptive statistics for Stage 1**

Topology	valid	ρ_{med}	$\rho_{CVaR_{0.10}}$	ρ_{IQR}	R_{rb}
hierar-chical	95	0.9661	0.9142	0.0254	0.3939
hub	100	0.9851	0.9294	0.0269	0.7000
reserved	98	0.9688	0.9188	0.0304	0.4545
dense	100	0.9900	0.9657	0.0092	0.8000

***Source*: compiled by the authors**

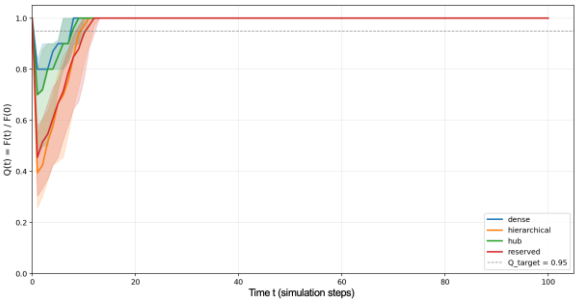


Fig. 1. **Mean curves of normalized performance *Q*(*t*) and *IQR* for the topologies in Stage 1**

***Source:* compiled by the authors**

Thus, the advantage of the dense topology is manifested not only in the typical scenario but also in the lower tail of the distribution (Fig. 2). This is consistent with the role of alternative routes: a larger number of independent paths reduce the dependence of global flow on individual nodes and edges.

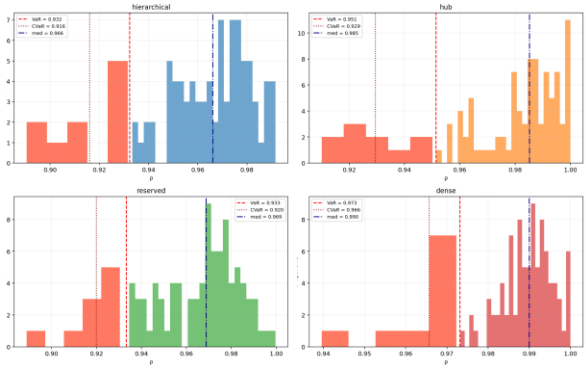


Fig. 2. **Comparison of lower-tail resilience CVaR for the topologies in Stage 1**

***Source:* compiled by the authors**

5.2. Effect of redundancy

Stage 2 examines the redundant topology under different values of the parameter *r*. The results show a positive but not strictly monotonic effect: the highest median resilience ρ_{med} and the highest lower-tail resilience ρ_{CVaR} are achieved at different levels of *r* – 3 and 1 respectively. The relationship is shown graphically in Fig. 3. This means that the number of redundant edges is not a self-sufficient indicator of usefulness: their location relative to flow bottlenecks is important.

The most important conclusion of Stage 2 is that redundancy has a stronger effect on tail scenarios than on median behavior. Additional links may only slightly change the typical level of resilience, but they can noticeably reduce the risk of catastrophic performance drops. Therefore, redundancy should be evaluated not only by the

change in median resilience ρ_{med} but also by the change in ρ_{CVaR}.

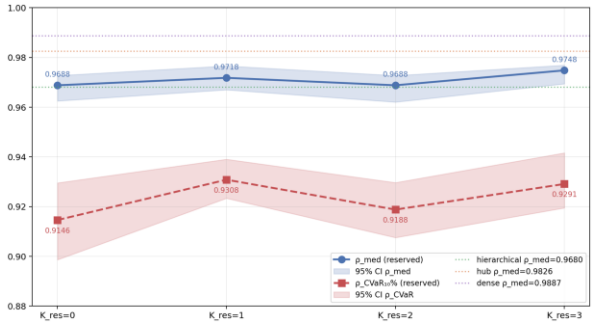


Fig. 3. **Dependence of resilience indicators of the reserved topology on the redundancy level *r***

Source: **compiled by the authors**

5.3. Type and intensity of disruptions

Stage 3 compares node failures, edge degradation, and cluster damage for dense and hierarchical topologies. The results confirm that edge degradation in the selected parameterization is largely compensated by capacity reserves, whereas node damage has a stronger impact on flow (Fig. 4). For the dense topology, all levels of edge degradation yield almost perfect resilience $\rho_{med} = 1$, whereas node failure and cluster disruption lead to a decline as intensity increases. For the hierarchical topology, a high level of node failure reduces median resilience to 0.9307 and robustness to 0.2424.

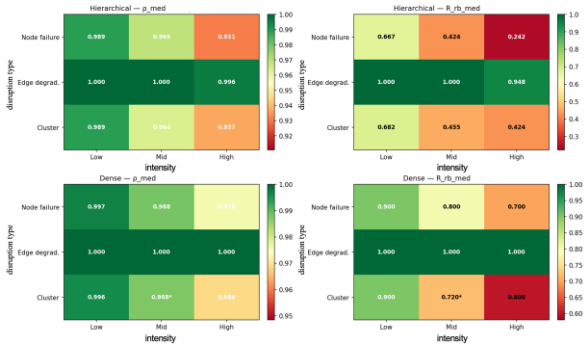


Fig. 4. **Heat map of resilience values depending on disruption type and intensity**

Source: **compiled by the authors**

A methodologically important point is that, after normalizing the "dose" of cluster disruption by the target share of damaged nodes, the hypothesis that cluster disruptions are automatically worse was not confirmed: in the "cluster vs node_failure" pairs, no statistically significant decrease in integral resilience was obtained. Thus, a cluster disruption may appear more dangerous not because of the damage type itself, but because of the larger actual number of damaged nodes. The component robustness metric captures the immediate depth of the performance drop more effectively.

5.4. Network scaling

Stage 4 verifies the stability of the conclusions for $N \in \{30, 50, 100\}$. For all topologies, an increase in N is accompanied by a decrease in median resilience because, at a fixed failure probability, the expected number of damaged elements increases. At the same time, the ranking by median resilience is preserved for all sizes: *dense* > *hub* > *reserved* > > *hierarchical* (Fig. 5). The Spearman coefficient between the rankings for all pairs of scales equals 1.

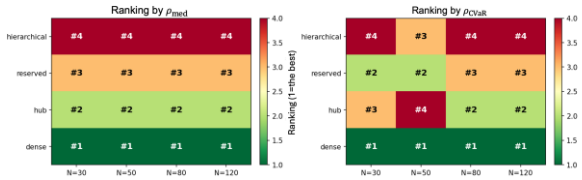


Fig. 5. **Stability of topology ranking under changes in network size *N***

Source: **compiled by the authors**

Tail resilience is more sensitive to scale: the dense topology retains first place, but the relative positions of hub and reserved may change. This means that the stability of the median ranking does not eliminate the need to analyze the lower tail of the distribution separately, especially for large networks.

5.5. Metric sensitivity and damage irreversibility

Stage 5 checks whether the main conclusion is an artifact of the selected metric parameters. For S5.A-S5.C, the canonical results of Stage 1 are used. Under variations of $\alpha \in \{0.05, 0.10, 0.20\}$, the integration horizon $T_{eff} \in \{30, 50, 100\}$, and the weight $\omega \in \{0, 0.5, 1\}$ in the composite performance function, the topology ranking remains unchanged: *dense* > *hub* > *reserved* > *hierarchical* (Fig. 6). Therefore, the advantage of the dense topology does not depend on the specific lower-tail level, integration horizon, or method of combining the flow and structural components.

Subseries S5.D assesses the influence of the irreversible share of damage w and requires new simulations. For hierarchical and reserved topologies, a paired seed design is used: within each topology, the same seed grid is applied for all *w* values, so the effect of the initial shock is not mixed with the effect of irreversibility. The key results are presented in Table 3 and in Fig. 7 and Fig. 8.

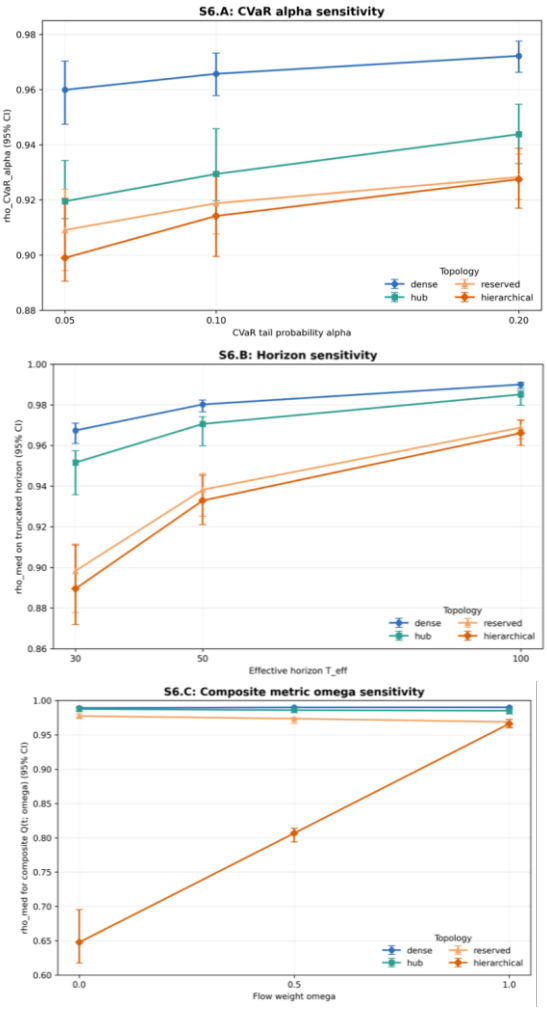


Fig. 6. **Sensitivity of the results to α, T_{eff}, and ω**

Source: **compiled by the authors**

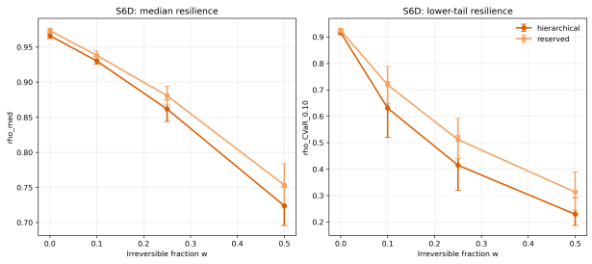


Fig. 7. **Effect of w on median resilience ρ_{med} and ρ_{CVaR} for hierarchical and reserved topologies**

Source: **compiled by the authors**

Table 3. **Effect of the irreversible damage share w on the resilience of hierarchical and reserved topologies**

Topology	w	valid	non_rec%	ρ_{med}	ρ_{CVaR}	R_{inf}
hierarchical	0.00	194	0.0	0.9662	0.9164	0.9697
hierarchical	0.10	194	30.9	0.9301	0.6308	0.9697
hierarchical	0.25	194	99.5	0.8615	0.4148	0.8788
hierarchical	0.50	194	100.0	0.7238	0.2293	0.7273
reserved	0.00	191	0.0	0.9733	0.9244	0.9697
reserved	0.10	191	30.4	0.9382	0.7189	0.9697
reserved	0.25	191	95.3	0.8806	0.5117	0.9091
reserved	0.50	191	99.5	0.7534	0.3131	0.7576

Source: **compiled by the authors**

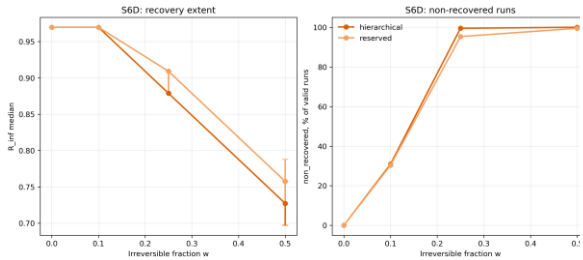


Fig. 8. **Effect of w on final recovery R_{inf} and the share of non_recovered runs**

Source: **compiled by the authors**

Irreversibility has a strong monotonic effect. For the hierarchical topology, median resilience decreases from 0.9662 to 0.7238, while lower-tail resilience decreases from 0.9164 to 0.2293 when w changes from 0 to 0.50. For the reserved topology, the corresponding values change from 0.9733 to 0.7534 and from 0.9244 to 0.3131. Lower-tail resilience reacts to irreversibility much more sharply than the median, meaning that the worst scenarios deteriorate first. At $w \geq 0.25$, almost all valid runs become non_recovered, which demonstrates the difference between "median resilience is not yet catastrophic" and "recovery completeness has already failed."

5.6. Synthesis of experimental findings

Taken together, the five experimental stages show that the same network can be evaluated differently depending on whether the focus is placed on typical behavior, worst-case scenarios, initial robustness, or final recovery. Dense topology is consistently the strongest archetype, but the comparison between hub and reserved structures depends more strongly on whether median or lower-tail resilience is emphasized. This confirms that a single aggregate indicator is not sufficient for a complete assessment of supply chain resilience.

The experiments also demonstrate that resilience is shaped by the interaction between

topology and disruption mechanism. Edge degradation has limited impact when capacity reserves and alternative routes are sufficient, whereas node failure directly removes flow-transmitting structures and therefore affects maximum flow more strongly. Cluster disruption requires careful interpretation because its apparent severity depends on the actual number of damaged nodes. Normalizing damage by target fraction makes the comparison more methodologically consistent.

Finally, the irreversibility subseries shows that the completeness of recovery cannot be inferred from integral resilience alone. A network may keep an acceptable average level of performance during the horizon while still failing to return to the initial functional state. Therefore, future benchmark studies should report both trajectory-based resilience and explicit recoverability indicators, especially when some damaged elements may remain unavailable for long periods.

6. DISCUSSION OF RESULTS

The obtained results should be interpreted as an analysis of structural and functional resilience, not as an economic optimization problem. The dense topology demonstrates the best values of typical and tail metrics, but this does not imply its universal practical optimality: a larger number of links may require higher costs, more complex coordination, and greater information transparency. Within this article, the dense network is rather an upper structural-functional boundary among the archetypes considered.

The hub topology has high median resilience but worse tail behavior. This is important for critical supply chain networks: evaluation based only on a typical scenario may underestimate the risk of rare but substantial performance declines. Redundancy, by contrast, often shows its greatest effect precisely in the lower tail. This conclusion is consistent with recent reviews in which redundancy, additional capacity, flexibility, and technological solutions are considered key directions of mitigation strategies for resilient supply chain networks [4], [8]. Therefore, when designing resilient networks, it is advisable to evaluate not only the increase in average or median performance but also the reduction of CVaR risk.

The results regarding disruption types emphasize the importance of properly normalizing damage intensity. If a cluster is defined only by radius, different topologies may receive different actual “doses” of destruction. Normalization by the target share of damaged nodes showed that cluster disruptions are not automatically worse than random node failures in terms of integral resilience, although they may produce a deeper immediate decline in robustness.

The scaling analysis confirms that the main topology ranking is not an artifact of the selected N=50. However, absolute resilience values decrease as network size increases, and the tail behavior of the middle-ranking positions becomes less stable. Therefore, large networks require a separate analysis not only of typical performance but also of the lower tail of the distribution.

The most important additional conclusion is provided by S5.D. The irreversibility of damage does not merely reduce integral resilience; it changes the qualitative recovery regime: at $w \geq 0.25$, almost all runs remain incompletely recovered. This means that resilience as the area under the $G(t)$ curve and recoverability as the ability to reach an acceptable final state are not identical characteristics. For future decentralized and multi-agent algorithms, this creates a natural prioritization problem: when it is impossible to restore everything, it is necessary to determine which nodes or edges make the largest contribution to global flow and lower-tail resilience.

The limitations of the study are related to the abstract nature of the model. First, a single-commodity maximum flow is used without explicit modeling of multiple product types, inventories, time delays, or the behavior of independent participants. Second, recovery is specified by a baseline random strategy rather than an optimal algorithm. Third, the topologies are controlled archetypes rather than empirical networks of specific companies. These limitations are intentional: they make it possible to isolate structural effects and form a benchmark for the next algorithmic stages of the research.

6.1. Implications for network design

The experiments suggest that network design should be evaluated not only by typical throughput but also by the distribution of post-disruption outcomes. Dense structures provide the highest structural-functional upper bound among the considered archetypes because they offer more opportunities for rerouting flow. However, the practical selection of such a structure must also account for cost, coordination complexity, information transparency, and operational feasibility.

The behavior of hub networks is more nuanced. They can be efficient in typical scenarios because centralized routes support high flow under normal or moderately damaged conditions. At the same time,

their lower-tail behavior reveals vulnerability when disruptions affect critical hubs or their incident links. This result is important for critical supply chains, where decisions based only on median performance can underestimate the risk of rare but severe performance drops.

6.2. Redundancy as tail-risk mitigation

The results on redundancy show that additional links should be assessed by their functional placement rather than merely counted. A redundant edge contributes to resilience only if it creates an alternative path that can carry meaningful flow when usual routes are damaged. This explains why the effect of redundancy is not strictly monotonic: some added links are more valuable than others because they bypass bottlenecks or connect otherwise vulnerable parts of the network.

The strongest role of redundancy appears in the lower tail. Additional routes may produce only a small increase in median resilience, but they can prevent catastrophic drops in adverse realizations. For decision makers, this means that the economic value of redundancy may be underestimated if the assessment focuses only on expected or typical performance. Tail-sensitive metrics should therefore be included in the evaluation of resilience investments.

6.3. Irreversibility and recoverability

The irreversibility analysis shows that resilience and recoverability are related but distinct characteristics. Integral resilience may remain moderate when the average trajectory is not fully collapsed, whereas the share of non-recovered runs can already be very high. This indicates a qualitative regime change: the system may still operate at a reduced level, but it no longer returns to its initial functional state within the modeled horizon.

This distinction has direct implications for future recovery algorithms. When all damaged elements can be restored, the main problem is the order and speed of repair. When some elements are irreversible or unavailable for a long time, the problem becomes one of prioritization under scarcity: which nodes or edges should receive attention first, and which losses can be tolerated with minimal impact on global flow and lower-tail resilience.

6.4. Limitations and future extensions

Several limitations should be emphasized. The model uses a single-commodity flow and does not explicitly represent inventories, lead times, multi-product interactions, contracts, prices, or strategic behavior of independent firms. The recovery process is a baseline random procedure, not an optimized plan. The topologies are controlled archetypes rather than empirical networks. These limitations are deliberate because they make it possible to isolate structural mechanisms before adding further operational complexity.

Future extensions can develop this benchmark in several directions. The first is to introduce targeted attacks and compare them with stochastic failures. The second is to replace the baseline recovery policy with centralized and decentralized algorithms that use flow-criticality indicators. The third is to add cost constraints and evaluate trade-offs between resilience gains and additional network complexity. The fourth is to calibrate topology generators on empirical supply chain data when such data are available.

6.5. Benchmark value for future recovery algorithms

The benchmark value of the proposed framework lies in its ability to provide a neutral reference point before controlled recovery policies are introduced. Since the current recovery process is intentionally simple, any future improvement obtained by a priority-based or multi-agent algorithm can be measured against the same topology, disruption, and metric structure. This is important for distinguishing the effect of an algorithm from the effect of a favorable network architecture.

For centralized recovery, the framework can be extended by ranking damaged elements according to their marginal contribution to maximum flow. For decentralized recovery, local agents may use limited information about neighboring nodes, local capacity, and approximate bottleneck indicators. For multi-agent recovery, coordination mechanisms can be evaluated by their ability to improve lower-tail resilience, not only median performance. In all cases, the proposed metrics provide a common evaluation language.

A particularly promising direction is to optimize recovery with respect to CVaR rather than only expected resilience. Such an objective would explicitly prioritize the reduction of severe outcomes in the lower tail of the distribution. This is consistent with the practical logic of critical supply chains, where preventing catastrophic loss may be more important than obtaining a small improvement in typical scenarios.

7. CONCLUSIONS

The article proposes a graph flow model and a reproducible computational framework for assessing the resilience of supply chain networks under stochastic disruptions. The network is represented as a directed graph with capacities, functionality is measured by normalized maximum flow, and resilience is evaluated through integral, component, and risk-oriented metrics. This formulation makes it possible to compare topologies not only by structural connectivity but also by their ability to maintain and restore functional flows.

The computational experiments showed that: (1) the dense topology consistently has the highest median and tail resilience; (2) the hub topology is effective in the typical regime but has a worse lower tail; (3) redundancy primarily reduces tail risks and does not always monotonically increase median resilience; (4) cluster disruptions are not automatically worse than random node failures after normalizing the damage "dose"; (5) the topology ranking by median resilience is preserved under scaling; (6) the main conclusion is robust to the choice of α, T_{eff} and ω; (7) the irreversibility of damage sharply worsens lower-tail resilience and recovery completeness.

The proposed framework forms a basis for further studies of controlled recovery. The next natural step is to replace the random baseline strategy with algorithms of prioritized, decentralized, and multi-agent recovery that use local information about the network state but are aimed at improving global flow resilience and tail resilience.

Conflicts of Interest: The authors declare that there is no conflict of interest

Funding: This research received no external funding

Accessibility data: The simulation code, experiment configurations, and raw and aggregated results supporting this study are publicly available in Zenodo at https://doi.org/10.5281/zenodo.23843543. The code repository is available at https://github.com/svpog/supply-chain-resilience-benchmark

Using artificial intelligence tools: The authors used ChatGPT (OpenAI) to assist with language editing, manuscript preparation, and checks of the research code and archived computational results. The authors reviewed and verified the AI-assisted content and take full responsibility for the accuracy and integrity of the article

REFERENCES

1. Ponomarov, S. Y. & Holcomb, M. C. "Understanding the concept of supply chain resilience". *The International Journal of Logistics Management*. 2009; 20 (1): 124–143. DOI: https://doi.org/10.1108/09574090910954873.

2. Pettit, T. J., Fiksel, J. & Croxton, K. L. "Ensuring supply chain resilience: development of a conceptual framework". *Journal of Business Logistics*. 2010; 31 (1): 1–21. DOI: https://doi.org/10.1002/j.2158-1592.2010.tb00125.x.

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

4. Roshani, A., Walker-Davies, P. & Parry, G. "Designing resilient supply chain networks: a systematic literature review of mitigation strategies". *Annals of Operations Research*. 2024; 341: 1267–1332. DOI: https://doi.org/10.1007/s10479-024-06228-6.

5. Ma, C., Zhang, L., You, L. & Tian, W. "A review of supply chain resilience: a network modeling perspective". *Applied Sciences*. 2025; 15 (1): 265. DOI: https://doi.org/10.3390/app15010265.

6. Korder, B., Maheut, J. & Konle, M. "Simulation methods and digital strategies for supply chains facing disruptions: insights from a systematic literature review". *Sustainability*. 2024; 16, (14): 5957. DOI: https://doi.org/10.3390/su16145957.

7. Bruckler, M., Wietschel, L., Messmann, L., Thorenz, A. & Tuma, A. "Review of metrics to assess resilience capacities and actions for supply chain resilience". *Computers & Industrial Engineering*. 2024; 192: 110176. DOI: https://doi.org/10.1016/j.cie.2024.110176.

8. Hosseini Shekarabi, S. A., Kiani Mavi, R. & Romero Macau, F. "Supply chain resilience: a critical review of risk mitigation, robust optimisation, and technological solutions and future research directions". *Global Journal of Flexible Systems Management*. 2025; 26: 681–735. DOI: https://doi.org/10.1007/s40171-025-00458-8.

9. Rockafellar, R. T. & Uryasev, S. "Optimization of conditional value-at-risk". *Journal of Risk*. 2000; 2 (3): 21–41. DOI: https://doi.org/10.21314/JOR.2000.038.

10. Rockafellar, R. T., & Uryasev, S. “Conditional value-at-risk for general loss distributions”. *Journal of Banking & Finance*. 2002; 26 (7): 1443–1471. DOI: https://doi.org/10.1016/S0378-4266(02)00271-6.

11. Acerbi, C. & Tasche, D. “On the coherence of expected shortfall”. *Journal of Banking & Finance*. 2002; 26 (7): 1487–1503. DOI: https://doi.org/10.1016/S0378-4266(02)00283-2.

12. Barabási, A.-L. “Network Science”. *Cambridge University Press*. 2016.

13. Dolgui, A., Ivanov, D. & Sokolov, B. “Ripple effect in the supply chain: an analysis and recent literature”. *International Journal of Production Research*. 2018; 56 (1–2): 414–430. DOI: https://doi.org/10.1080/00207543.2017.1387680

14. Ivanov, D. & Dolgui, A. “New disruption risk management perspectives in supply chains: digital twins, ripple effect, and resileanness”. *IFAC-PapersOnLine*. 2019; 52 (13): 337–342. DOI: https://doi.org/10.1016/j.ifacol.2019.11.138

15. Bruneau, M., et al. “A framework to quantitatively assess and enhance the seismic resilience of communities”. *Earthquake Spectra*. 2003; 19 (4): 733–752. DOI: https://doi.org/10.1193/1.1623497.

16. Cimellaro, G. P., Reinhorn, A. M. & Bruneau, M. “Framework for analytical quantification of disaster resilience”. *Engineering Structures*. 2010; 32 (11): 3639–3649. DOI: https://doi.org/10.1016/j.engstruct.2010.08.008.

17. Henry, D. & Ramirez-Marquez, J. E. “Generic metrics and quantitative approaches for system resilience as a function of time”. *Reliability Engineering & System Safety*. 2012; 99: 114–122. DOI: https://doi.org/10.1016/j.ress.2011.09.002.

18. Efron, B. & Tibshirani, R. J. “An Introduction to the Bootstrap”. B*oca Raton, FL, USA: Chapman & Hall/CRC*. 1993. DOI: https://doi.org/10.1201/9780429246593.

19. Kruskal, W. H. & Wallis, W. A. “Use of ranks in one-criterion variance analysis”. *Journal of the American Statistical Association*. 1952; 47 (260): 583–621. DOI: https://doi.org/10.1080/01621459.1952.10483441.

20. Mann, H. B. & Whitney, D. R. “On a test of whether one of two random variables is stochastically larger than the other”. *The Annals of Mathematical Statistics*. 1947; 18 (1): 50–60. DOI: https://doi.org/10.1214/aoms/1177730491.

21. Cliff, N. “Dominance statistics: ordinal analyses to answer ordinal questions”. *Psychological Bulletin*. 1993; 114 (3): 494–509. DOI: https://doi.org/10.1037/0033-2909.114.3.494.

22. Rubinstein, R. Y., & Kroese, D. P. “Simulation and the Monte Carlo Method”. *3rd ed. Hoboken, NJ, USA: Wiley*. 2016. DOI: https://doi.org/10.1002/9781118631980.

23. Law, A. M. “Simulation modeling and analysis”. *5th ed. New York, NY, USA: McGraw-Hill Education*. 2015.

24. Wilson, G., et al. “Best practices for scientific computing”. *PLoS Biology*. 2014; 12 (1): e1001745. DOI: https://doi.org/10.1371/journal.pbio.1001745.

25. Sandve, G. K., et al. “Ten simple rules for reproducible computational research”. *PLoS Computational Biology*. 2013; 9 (10): e1003285. DOI: https://doi.org/10.1371/journal.pcbi.1003285.

Received 21.08.2026
Received after revision 15.10.2026
Accepted 23.10.2026

DOI: https://doi.org/10.15276/aait.09.2026.35
УДК 004.942:519.17

Стійкість мереж поставок за стохастичних збоїв

Погорєлов Станіслав Вікторович[1)]
ORCID: https://orcid.org/0000-0002-0189-8655; Stanislav.Pohorielov@khpi.edu.ua. Scopus Author ID: 9270293800
Яновський Олексій Васильович[1)]
ORCID: https://orcid.org/0000-0002-4310-2843; Oleksii.Yanovskyi@cs.khpi.edu.ua
[1)] Національний технічний університет «Харківський політехнічний інститут», вул. Кирпичова, 2. Харків, 61002, Україна

АНОТАЦІЯ

Актуальність дослідження зумовлена зростанням вразливості мереж поставок до стохастичних збоїв, локальних відмов, втрати пропускної здатності та кластерних пошкоджень. **Метою** роботи є розроблення й перевірка графової

потокової моделі та відтворюваного обчислювального фреймворку для оцінювання структурно-функціональної стійкості таких мереж. Запропонований підхід подає мережу поставок як орієнтований граф із пропускними здатностями, а продуктивність визначає через нормований максимальний потік від джерел до споживачів. Методика поєднує моделювання випадкових пошкоджень, відмов вузлів, деградації ребер, кластерних збоїв і незворотних втрат із базовим дискретним процесом відновлення та протоколом Monte Carlo. Стійкість оцінюється за часовою траєкторією продуктивності, інтегральною метрикою, робастністю, фінальним рівнем відновлення, медіанною стійкістю, міжквартильним розмахом і нижньохвостовою ризик-орієнтованою метрикою. Результати обчислювальних експериментів показали, що щільна топологія забезпечує найвищу типову та нижньохвостову стійкість, резервування насамперед зменшує ризик найгірших сценаріїв, а незворотність пошкоджень різко погіршує повноту відновлення. **Практична цінність** фреймворку полягає у можливості його використання як benchmark-основи для подальших алгоритмів пріоритетного, децентралізованого та мультиагентного відновлення мереж поставок.

Ключові слова: стійкість мереж; мережі поставок; графові моделі; стохастичні збої; максимальний потік; відновлення мереж

ABOUT THE AUTHORS



Stanislav V. Pohorielov - Dr. Sci. (Phys.-Math.), Professor, Department of Computer Mathematics and Data Analysis, National Technical University “Kharkiv Polytechnic Institute”, 2, Kyrpychova St. Kharkiv, 61002, Ukraine
ORCID: https://orcid.org/0000-0002-0189-8655; Stanislav.Pohorielov@khpi.edu.ua. Scopus Author ID: 9270293800
Research field: Data analysis, computational mathematics, quantum physics

Погорєлов Станіслав Вікторович - доктор.фізико-матем. наук, професор кафедри Комп’ютерної математики та аналізу даних. Національний технічний університет «Харківський політехнічний інститут», вул. Кирпичова, 2. Харків, 61002, Україна
Галузь досліджень: аналіз даних, обчислювальна математика, квантова фізика



Oleksii V. Yanovskyi - postgraduate student, Department of Computer Mathematics and Data Analysis, National Technical University “Kharkiv Polytechnic Institute”, 2, Kyrpychova St. Kharkiv, 61002, Ukraine
ORCID: https://orcid.org/0000-0002-4310-2843; Oleksii.Yanovskyi@cs.khpi.edu.ua
Research field: Supply chain resilience, stochastic disruptions, graph-flow models, network recovery.

Яновський Олексій Васильович - аспірант кафедри Комп’ютерної математики та аналізу даних. Національний технічний університет «Харківський політехнічний інститут», вул. Кирпичова, 2. Харків, 61002, Україна
Галузь досліджень: стійкість ланцюгів поставок, стохастичні збої, графово-потокові моделі, відновлення мережі