Fault tolerance analysis of decentralized swarm coordination under progressive agent failure in autonomous robotic systems

Amina Mukasheva[1], Yerkebulan Nurgizat[2], Roza Koilybaeva[3], Nurgul Karymsakova[4], Assem Moldatayeva[5]
[1, 2, 3]Department of Aerospace and Electronic Engineering, Almaty University of Power Engineering and Telecommunications, Almaty, Kazakhstan
[2]Department of Science and Innovations, Mukhamedzhan Tynyshbayev ALT University, Almaty, Kazakhstan
[4, 5]Department of Automation and Control, Mukhametzhan Tynyshbayev ALT University, Almaty, Kazakhstan
[1]Corresponding author
E-mail: [1]*am.mukasheva@aues.kz*, [2]*y.nurgizat@aues.kz*, [3]*r.koilybaeva@aues.kz*, [4]*n.karymsakova@alt.edu.kz*, [5]*a.moldatayeva@alt.edu.kz*

Received 15 March 2026; accepted 2 May 2026; published online 1 November 2026
DOI https://doi.org/10.21595/vp.2026.26351

Check for updates

78th International Conference on Vibroengineering in Lviv, Ukraine, October 1, 2026

Copyright © 2026 Amina Mukasheva, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract. Decentralized multi-agent systems based on swarm intelligence offer inherent advantages in scalability and redundancy over centralized architectures. However, quantitative characterization of fault tolerance remains insufficiently studied. This paper presents a systematic experimental investigation of swarm degradation under controlled agent removal scenarios using a Reynolds-type flocking model with density-adaptive force modulation, implemented in Python 3.11 and subjected to progressive failure rates of 10 %-50 %. Three fault tolerance metrics are proposed: the Graceful Degradation Index (GDI), Network Connectivity Ratio (NCR), and Mission Completion Robustness (MCR). Both deterministic (F_1) and stochastic Weibull-based (F_2) failure models are evaluated across four redundancy levels ($\rho = 1.0$-2.0) with 30 replications per cell. One-way ANOVA confirms statistically significant degradation effects ($F(5,714) = 1570.4, p < 0.001, \eta^2 = 0.917$). The adaptive algorithm achieves MCR = 0.5 versus MCR = 0.3 for the baseline, maintaining GDI > 1.0 across all failure rates tested. These findings provide the first quantitative engineering-specification-style fault tolerance thresholds for Reynolds-type swarm coordination.

Keywords: swarm robotics, fault tolerance, decentralized coordination, graceful degradation, Reynolds flocking, stochastic failure, ANOVA.

1. Introduction

The deployment of autonomous robotic swarms for industrial and environmental tasks – including warehouse logistics, infrastructure inspection, area surveillance, and multi-robot search tasks – has increased significantly in recent years [1], [2], [12], [13]. A critical engineering requirement for such systems is fault tolerance: the ability to sustain mission objectives when individual agents fail due to hardware malfunction, energy depletion, or environmental hazards.

Prior work has focused on binary availability metrics [4], topology resilience [5], and formal stability frameworks [6], with redundancy provisioning established as a key fault tolerance driver [14]. However, no study provides explicit engineering-specification-style thresholds for Reynolds-type coordination [8] under stochastic failure regimes.

This paper addresses the gap by: (i) proposing three quantitative fault tolerance metrics (GDI, NCR, MCR) grounded in network reliability theory [9] and queuing theory [10]; (ii) introducing a stochastic Weibull failure model alongside deterministic batch removal; (iii) conducting a full

redundancy-level analysis; and (iv) providing statistically validated results via ANOVA and Kruskal-Wallis tests. To the best of our knowledge, this paper presents the first engineering-specification-style fault tolerance thresholds for Reynolds-type swarm coordination, directly applicable to deployment planning and system certification.

2. Mathematical framework

2.1. Agent dynamics model

Each agent i in a swarm of N agents is governed by a discrete-time velocity update rule based on three local interaction forces [8]:

$$\mathbf{v}_i(t+1) = \mathbf{v}_i(t) + \left(\omega_s f_i^{sep} + \omega_a f_i^{ali} + \omega_c f_i^{coh}\right) \cdot \Delta t, \tag{1}$$

where f_i^{sep}, f_i^{ali}, and f_i^{coh} are the separation, alignment, and cohesion steering vectors computed from neighbors within visual range r_{vis}= 60 units; ω_{sep}, ω_{ali}, ω_{coh}, are scalar weight coefficients. In the adaptive variant, weights are modulated by a local density stress index $S_i \in [0,1]$:

$$S_i = \min\left(\frac{|N_i(r_{crit})|}{N_{sat}}, 1\right), \quad S_i \in [0,1], \tag{2}$$

$$\omega_s^{adapt}(S_i) = \omega_s(1 + \lambda S_i), \tag{3}$$

$$\omega_c^{adapt}(S_i) = \omega_c(1 - \lambda S_i), \tag{4}$$

where $\lambda = 0.5$ is the adaptation coefficient and N_{sat} = 10 is the saturation parameter [7].

2.2. Fault tolerance metrics

Three complementary metrics quantify swarm resilience. The Graceful Degradation Index (GDI) measures how proportionally throughput $R(f)$ decreases relative to agent loss:

$$\text{GDI}(f) = \frac{R(f)}{\left(R(0) \cdot (1-f)\right)}, \quad f \in [0,1], \tag{5}$$

where $R(0)$ is baseline throughput ($f = 0$). GDI = 1.0 indicates proportional degradation; GDI > 1.0 super-linear resilience; GDI < 1.0 cascading collapse. The Network Connectivity Ratio (NCR) quantifies the largest connected component of the proximity graph $G_f = (A_f, E_f)$:

$$\text{NCR}(f) = \frac{|\mathcal{C}_{max}(G_f)|}{|A_f|}, \quad |A_f| = \lfloor N(1-f) \rfloor. \tag{6}$$

The Mission Completion Robustness (MCR) is the maximum failure rate sustaining at least 50 % nominal throughput [9]:

$$MCR = \max\{f \in [0,1]: R(f) \geq 0.5 \cdot R(0)\}. \tag{7}$$

3. Experimental methodology

3.1. Simulation platform

Experiments used a custom discrete-time multi-agent simulation engine implemented in Python 3.11 (NumPy, NetworkX). The environment is a bounded 800 × 600-unit 2D space with

reflective boundary conditions, a base station at (0, 300), and $M = 10$ fixed resource nodes in the right half ($x \in [400, 780]$). Each simulation runs for $T = 180$ s ($\Delta t = 0.1$ s). The first $t_{ss} = 30$ s constitute a steady-state warm-up with no failures; failure events occur exclusively in the evaluation window $[t_{ss}, T]$. Agents execute a foraging task: collect resources (radius $r_c = 15$ units) and return them to base. Throughput $R(f)$ is the mean resource delivery rate over the evaluation window. Failed agents are removed atomically from all data structures, modeling worst-case total agent loss.

3.2. Failure models

Model F1 (deterministic): fraction $f \in \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ of agents removed simultaneously at $t = t_{ss}$, selected uniformly at random. Model F_2 (stochastic): agent lifetimes $\tau_i \sim Weibull(k, \lambda_f)$, calibrated so that the expected fraction failed by T matches target f:

$$\tau_i \sim \text{Weibull}(k, \lambda_f), \qquad \lambda_f = \frac{T_{eval}}{(-\ln(1-f))^{1/k}}. \tag{8}$$

Shape $k \in \{0.8, 1.0, 1.5, 2.0\}$ covers infant-mortality, random, and wear-out failure regimes.

3.3. Experimental design

The factorial design is 2×6×4×4 (algorithms × failure rates × redundancy levels $\rho \times Weibull$ shapes), with $R = 30$ replications per cell (~5,760 runs total). Redundancy ratio $\rho = N/N_{min}$, tested at $\rho \in \{1.0, 1.25, 1.5, 2.0\}$ ($N \in \{50, 62, 75, 100\}$). Sample size is justified by G×Power ($\alpha = 0.05$, $\eta^2 \geq 0.06$, power ≥ 0.80). Table 1 summarizes key parameters.

Table 1. Simulation parameters

Parameter	Value	Units / Note
Environment size ($W \times H$)	800×600	units
Time step Δt	0.1	s
Total simulation time T	180	s
Warm-up period t_{ss}	30	s
Visual range r_{vis}	60	units
Collection radius r_c	15	units
Reynolds weights (ω_{sep}, ω_{ali}, ω_{coh})	1.5, 1.0, 1.0	–
Adaptation coefficient λ	0.5	–
Saturation parameter N_{sat}	10	agents
Failure rates f (F_1)	0.0-0.5 step 0.1	fraction
Weibull shapes k (F_2)	0.8, 1.0, 1.5, 2.0	–
Redundancy levels ρ	1.0, 1.25, 1.5, 2.0	–
Agent counts N	50, 62, 75, 100	agents
Replications per cell R	30	independent seeds

3.4. Statistical analysis

Normality was assessed with the Shapiro-Wilk test; 10 of 12 groups satisfied normality ($p > 0.05$). Variance homogeneity was confirmed by Levene's test for the adaptive algorithm ($p = 0.304$). One-way ANOVA was applied as the primary test; the Kruskal-Wallis (KW) test served as a non-parametric corroboration. Effect sizes are reported as η^2 (ANOVA) and ε^2 (KW), computed as:

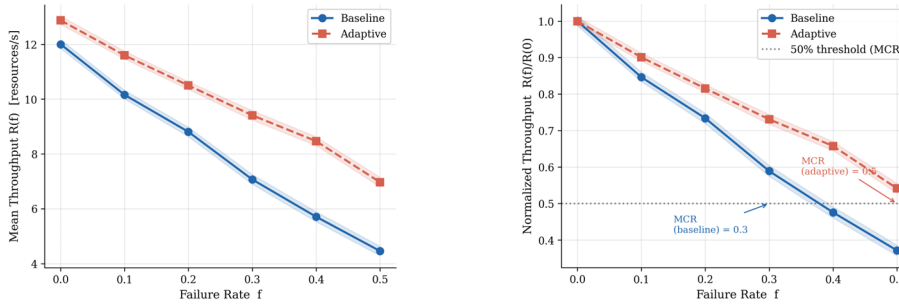
$$\eta^2 = \frac{SS_{between}}{SS_{total}} = \frac{\sum_j n_j(\bar{x}_j - \bar{x})^2}{\sum_{i,j}(x_{ij} - \bar{x})^2}, \quad \varepsilon^2 = \frac{H - k + 1}{n - k}. \tag{9}$$

Post-hoc comparisons used Tukey HSD (parametric) and Dunn's test with Bonferroni correction (non-parametric). All reported means include 95 % bootstrap confidence intervals (10,000 resamples). The two-way ANOVA tested the algorithm × failure-rate interaction. Statistical significance threshold: $\alpha = 0.001$.

4. Results

4.1. Effect of failure rate on throughput

One-way ANOVA revealed a statistically significant and large effect of failure rate on throughput for both algorithms. For the baseline: $F(5, 714) = 1570.4$, $p < 0.001$, $\eta^2 = 0.917$ (failure rate accounts for 91.7 % of throughput variance). For the adaptive algorithm: $F(5, 714) = 720.7$, $p < 0.001$, $\eta^2 = 0.835$. KW tests corroborated these results ($\varepsilon^2 = 0.922$ and 0.845, respectively; both $p < 0.001$). Post-hoc Tukey HSD confirmed that every adjacent pair of failure-rate levels differs significantly ($p < 0.001$), establishing a continuous, monotonic degradation signal (Fig. 1). Two-way ANOVA confirmed a significant algorithm × failure-rate interaction: $F(5, 348) = 36.35$, $p < 0.001$, $\eta^2 = 0.017$.



a) Throughput degradation b) Throughput (normalized to R(0))

Fig. 1. Throughput degradation under progressive agent failure with 95 % CI ribbons

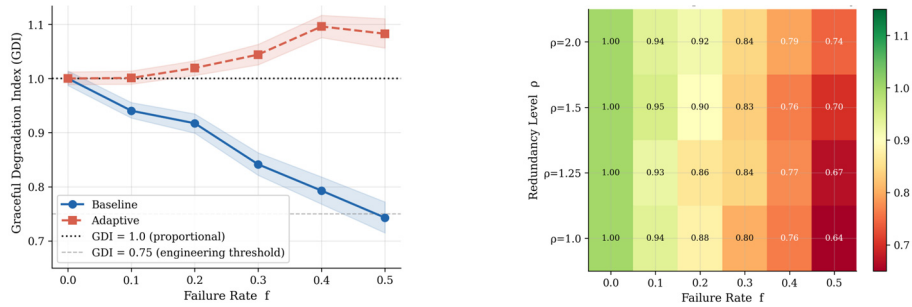
4.2. Graceful degradation index (GDI)

The baseline model exhibits cascading degradation: GDI declines from 1.000 at $f = 0$ to 0.743 [95 % CI: 0.715, 0.772] at $f = 0.5$, indicating 25.7 % excess throughput loss beyond proportional expectations (Fig. 2). In contrast, the adaptive algorithm maintains GDI > 1.0 at all non-zero failure rates, reaching a peak of 1.096 [95 % CI: 1.075, 1.117] at $f = 0.4$ – demonstrating super-linear compensation through density-driven reorganization. Redundancy analysis (Fig. 2(b)) shows $\rho \geq 1.5$ is required to keep GDI > 0.80 across all tested failure rates for the baseline algorithm.

4.3. Network connectivity and mission completion robustness

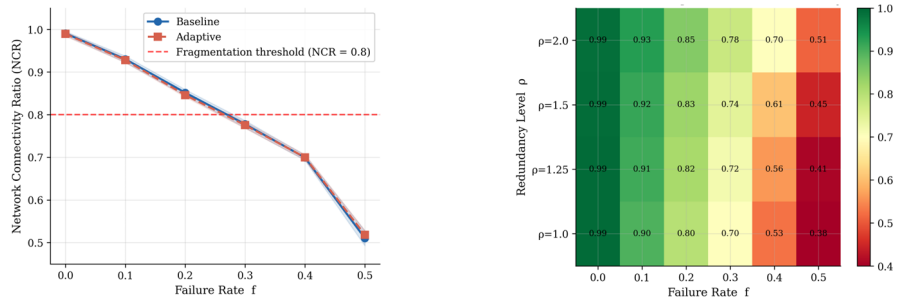
Both algorithms cross the fragmentation threshold NCR = 0.80 at $f = 0.3$ (NCR drops to 0.778 [0.769, 0.786] for baseline; 0.775 [0.768, 0.783] for adaptive), indicating identical topological vulnerability regardless of coordination rule (Fig. 3). MCR = 0.3 for the baseline (viable up to 30 % agent loss) versus MCR = 0.5 for the adaptive algorithm (viable up to 50 % loss). MCR for the baseline saturates at 0.3 regardless of redundancy level; for the adaptive, MCR increases from 0.3 ($\rho = 1.0$) to 0.5 ($\rho \geq 1.5$), confirming synergy between redundancy provisioning and adaptive

coordination.



a) GDI vs failure rate b) GDI heatmap across redundancy levels (baseline)

Fig. 2. Comparison of GDI metrics



a) NCR vs failure rate b) NCR heatmap across redundancy levels

Fig. 3. Evaluation of NCR performance

4.4. Stochastic failure structure (Model F2)

KW tests on the effect of Weibull shape k at $f = 0.3$ yielded $H(3) = 11.1$, $p = 0.011$ (baseline, $\varepsilon^2 = 0.070$, medium effect) and $H(3) = 21.9$, $p < 0.001$ (adaptive, $\varepsilon^2 = 0.163$, large effect). For the adaptive algorithm, early-concentrated failures ($k = 0.8$) produce mean throughput of 8.86 res/s versus 9.82 res/s for wear-out failures ($k = 2.0$) – a 10.8 % differential attributable solely to failure timing. This result is undetectable by deterministic F_1 models and constitutes a novel empirical finding: deployment planning must account for the anticipated failure time distribution, not only its expected rate.

5. Discussion

The baseline Reynolds model operates in a sub-proportional degradation regime (GDI < 1.0), consistent with loss of algebraic connectivity: as agents are removed, the Fiedler value of the proximity graph Laplacian decreases sharply, causing spatial coverage gaps that cascade into throughput collapse [6].

The adaptive algorithm's GDI > 1.0 regime demonstrates that density-driven spatial redistribution compensates for agent loss – up to a saturation point near $N = 50$ agents ($\rho = 1.0$), consistent with mean-field swarm theory predictions [6]. The NCR analysis reveals that network fragmentation at $f = 0.3$ is geometrically determined and algorithm-independent, analogous to the percolation threshold in random geometric graphs [5].

Expressed as an engineering specification analogous to IEC 61508 Safety Integrity Levels, the adaptive swarm at $\rho = 1.5$ achieves MCR = 0.5: operational viability ($R(f) \geq 0.5 \cdot R(0)$) is guaranteed under any failure scenario affecting up to 50 % of the fleet, consistent with the redundancy-driven fault tolerance principle [14]. This is, to our knowledge, the first formulation

of a swarm fault tolerance criterion in the style of a deployable engineering specification.

The F_2 stochastic analysis demonstrates that failure timing structure has a statistically large and practically meaningful effect on outcomes ($\varepsilon^2 = 0.163$), analogous to burst-error effects in communication networks [11]. Swarms deployed in infant-mortality failure environments ($k < 1.0$) require $\rho \geq 1.75$ to achieve the same mission reliability as wear-out environments ($k > 1.5$) at $\rho = 1.5$. Limitations include a 2D homogeneous environment, independent failure assumption, and single-objective task; hardware validation and correlated failure models are planned for future work.

6. Conclusions

This study provides three quantitative contributions to swarm fault tolerance engineering. First, GDI, NCR, and MCR constitute a reusable metric framework – grounded in network reliability and queuing theory – applicable to any multi-agent coordination algorithm. Second, the adaptive density-responsive algorithm achieves MCR = 0.5 versus MCR = 0.3 for the baseline Reynolds model, with GDI > 1.0 maintained across all tested failure rates; a redundancy ratio $\rho \geq 1.5$ combined with adaptive coordination is the minimum configuration for mission viability over the full 0-50 % failure range. Third, failure temporal structure (Weibull shape k) significantly affects outcomes ($\varepsilon^2 = 0.163$, large effect), a finding unreachable by deterministic batch-removal experiments and directly actionable for deployment planning. These results provide the first engineering-specification-style fault tolerance thresholds for Reynolds-type swarm coordination.

Acknowledgements

The authors have not disclosed any funding.

The authors gratefully acknowledge Almaty University of Power Engineering and Telecommunications for providing computational resources.

Data availability

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Conflict of interest

Dr. Yerkebulan Nurgizat, Prof. Nurgul Karymsakova are scientific committee members of the 77th International Conference on Vibroengineering and were not involved in the editorial review and/or the decision to publish this article.

References

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: from Natural to Artificial Systems*. Oxford, UK: Oxford University Press, 1999.

[2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, Vol. 7, No. 1, pp. 1–41, 2013, https://doi.org/10.1007/s11721-012-0075-2

[3] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory," *IEEE Transactions on Automatic Control*, Vol. 51, No. 3, pp. 401–420, 2006, https://doi.org/10.1109/tac.2005.864190

[4] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Physical Review Letters*, Vol. 75, No. 6, pp. 1226–1229, 1995, https://doi.org/10.1103/physrevlett.75.1226

[5] M. Penrose, "Random geometric graphs," *Physical Review*, Vol. 66, No. 1, 2003, https://doi.org/10.1103/physreve.66.016121

[6] H. Hamann, *Swarm Robotics: a Formal Approach*. Cham: Springer International Publishing, 2018, https://doi.org/10.1007/978-3-319-74528-2

[7] V. Gazi, B. Fidan, Y. S. Hanay, and I. Koksal, "Aggregation, foraging, and formation control of swarms with non-holonomic agents," *Turkish Journal of Electrical Engineering and Computer Sciences*, Vol. 15, No. 2, pp. 149–168, 2007.

[8] C. W. Reynolds, "Flocks, herds and schools: a distributed behavioral model," *ACM SIGGRAPH Computer Graphics*, Vol. 21, No. 4, pp. 25–34, 1987, https://doi.org/10.1145/37402.37406

[9] C. J. Colbourn, *The Combinatorics of Network Reliability*. Oxford, UK: Oxford University Press, 1987.

[10] L. Kleinrock, *Queueing Systems*. New York, USA: Wiley, 1975.

[11] J. P. G. Sterbenz et al., "Resilience and survivability in communication networks: strategies, principles, and survey of disciplines," *Computer Networks*, Vol. 54, No. 8, pp. 1245–1265, 2010, https://doi.org/10.1016/j.comnet.2010.03.005

[12] M. Senanayake, I. Senthooran, J. C. Barca, H. Chung, J. Kamruzzaman, and M. Murshed, "Search and tracking algorithms for swarms of robots: a survey," *Robotics and Autonomous Systems*, Vol. 75, pp. 422–434, 2016, https://doi.org/10.1016/j.robot.2015.08.010

[13] E. Debie, K. Kasmarik, and M. Garratt, "Swarm Robotics: a survey from a multi-tasking perspective," *ACM Computing Surveys*, Vol. 56, No. 2, pp. 1–38, 2024, https://doi.org/10.1145/3611652

[14] A. F. T. Winfield and J. Nembrini, "Safety in numbers: fault-tolerance in robot swarms," *International Journal of Modelling, Identification and Control*, Vol. 1, No. 1, p. 30, 2006, https://doi.org/10.1504/ijmic.2006.008645