Journal of Defence and Security Industries: Strategy and Technology, 2026, 1(2)

Science & Technology

Volume Over Lethality: A Matrix-Exponential Framework for Defending Drone Swarms

Mustafa Tekinay [1]

Abstract

We study large, heterogeneous drone swarms composed of both threat and decoy drones. The defender's engagement is modelled as a multi-stage stochastic process that captures uncertainty and variability in response times. A non-discriminatory defender is considered in which threat and decoy drones are treated identically. We adopt a matrix-exponential framework that yields tractable, closed-form expressions for key performance measures. The results indicate that swarm size has a substantially stronger impact on defence success than threat-decoy composition. For a typical defender, swarms exceedingly roughly ten drones may exceed single-defender swarm-size capacity under the baseline modelling assumptions, even with high decoy fractions. While a typical defender can reliably handle seven threat drones, augmenting the swarm with seven decoys reduces the probability of successful defence from 0.93 to 0.09. From a cost perspective, heterogeneous drone swarms can reduce overall cost by up to 67% when decoy drones are priced at approximately one tenth of the cost of threat drones. Finally, performance depends not only on aggregate rates but also on individual stage bottlenecks. Under idealised conditions where each engagement stage completes in approximately one second on average, single-defender capacity increases to approximately 35-40 drones.

Keywords: Drone swarms, counter-UAV systems, threat drones, decoy drones, matrix-exponential distributions, stochastic performance analysis, swarm defence capacity.

Received: 24.01.2026 Accepted: 25.05.2026 Research Article

[1] ASELSAN Inc., Ankara, Türkiye
*Corresponding author.
✉ m.tekinay@aselsan.com

Cite this article as:
https://doi.org/10.65834/jdsi.12.58
Tekinay, M. (2026). Volume Over Lethality: A Matrix-Exponential Framework for Defending Drone Swarms. *Journal of Defence and Security Industries: Science & Technology1(2),* 59-87.

1. Introduction

Unmanned Aerial Systems (UAS) have rapidly become a central element of modern warfare. Compared to manned platforms and long-range precision weapons, drones are cost-effective, low-risk, and scalable, enabling persistent operations with limited human exposure (Special Competitive Studies Project, 2023). Their relatively low unit cost allows attackers to trade individual platform survivability for mass and persistence.

The ongoing Russia-Ukraine war provides a clear and contemporary example of this shift. Both sides have repeatedly employed large-scale drone attacks, often involving hundreds of Unmanned Aerial Vehicles (UAVs) in a single engagement, aimed at overwhelming air defence systems through saturation rather than precision (Kyiv Post, 2026; Reuters, 2025a, 2025b, 2025c, 2026). These engagements demonstrate that swarm size itself has become an operational variable.

A key feature of contemporary drone warfare is the deliberate use of heterogeneous swarms, in which a fraction of the platforms is designed to act as decoys or expendable assets. These drones are significantly cheaper and intended to trigger sensors, consume interceptors, and degrade command-and-control effectiveness, thereby increasing the likelihood that a smaller number of high-value threat drones penetrate defensive layers (Center for European Policy Analysis, 2024; Espreso Global, 2024). Such tactics exploit structural limitations of air defence systems, which must sequentially perform detection, identification, tracking, engagement preparation and neutralisation for each incoming target (Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025). This sequential processing structure implies that defender performance is fundamentally governed by the cumulative time required to process incoming targets. Each engagement consists of multiple stochastic stages, and under swarm conditions, these stage durations accumulate across targets, creating a compounded timing constraint that directly limits system capacity. When confronted with sufficiently large swarms, even advanced systems experience performance degradation due to finite processing rates and resource constraints. This mechanism is analogous to Distributed Denial-of-Service (DDoS) attacks in computer networks.

Replacing threat drones with decoy drones further enables higher production rates and sustained operational tempo, contributing to prolonged pressure and operational fatigue. Decoy drones are typically reported to cost on the order of a few thousand US dollars, whereas long-range strike drones may cost on the order of two hundred thousand US dollars (Center for European Policy Analysis, 2024). Russia has reportedly introduced several decoy or imitation UAV platforms since 2023, commonly referred to as *Geran/Gerbera* and *Parodiya*, which are intended to emulate the general shape and flight behaviour of the *Shahed-136* drone (Anokhin & Faragasso, 2025; Espreso Global, 2024).

From a modelling perspective, this leads to a structured stochastic system in which individual engagements are composed of multiple sequential phases, and overall defence performance depends on the aggregation of these processes across multiple targets. Capturing such behaviour requires a representation that can model multi-stage dynamics while remaining tractable under repeated convolution and mixture. In particular, the ability to characterise cumulative engagement time in closed form is essential for assessing system effectiveness under varying swarm sizes and threat compositions.

Motivated by these observations, this work develops a heterogeneous drone swarm model to study the impact of swarm size and threat composition on defender performance. Drone defence is modelled as a multi-stage process, and the impact of the number of drones and the proportion of threat drones on system effectiveness is examined. To capture this structure, we adopt a Matrix-Exponential (ME) representation, which enables modelling of multi-stage engagement dynamics while preserving analytical tractability. While simulation-based approaches can reproduce engagement scenarios, they provide limited analytical insight into scaling behaviour and often require extensive computation across parameter spaces. In contrast, an analytical formulation enables direct characterisation of system performance and exposes the structural drivers of defence effectiveness. In this formulation, the sequential engagement stages correspond to phases of the ME representation, repeated neutralisation of multiple targets corresponds to convolution of these processes, and heterogeneous swarm composition induces a mixture structure over the number of threat drones.

Existing work on UAV swarm defence can be broadly categorised based on modelling approach, including agent-based simulation, optimisation-driven deployment strategies, and system-level dynamic models. These approaches differ in how they represent engagement dynamics and in the level of analytical tractability they provide. Overall, while the literature on UAV swarm defence is growing, analytically tractable models that capture engagement-time dynamics remain relatively limited.

A first class of work focuses on modelling interactions between attacking and defending agents at the swarm level. Brust et al. (2021) propose a swarm-based counter-UAV defence system in which coordinated defensive agents intercept intruding UAVs. Similarly, Gaertner (2013) study UAV swarm tactics using an agent-based simulation framework combined with Markov process modelling. These approaches capture detailed interaction dynamics and coordination effects among agents but rely primarily on simulation-based evaluation and do not provide closed-form characterisation of cumulative engagement time or defender capacity under large-scale swarms.

A second class of work addresses optimal deployment and strategic decision-making for countering UAV swarms. Li et al. (2023) and Wang et al. (2023) investigate optimal placement and positioning of defensive resources to improve coverage and interception effectiveness. Mutzari et al. (2025) model urban multi-drone defence as a sequential Stackelberg security game, focusing on strategic allocation and adversarial decision-making. While these studies provide valuable insights into spatial configuration and resource allocation, they do not explicitly capture the temporal accumulation of engagement processes that governs defender throughput under saturation conditions.

A third line of research emphasises high-fidelity simulation and system-level effectiveness evaluation. Hotters et al. (2025) and Yan et al. (2025) employ detailed simulation frameworks for red-blue confrontation scenarios, while NATO methodologies and related studies provide structured approaches for modelling UAV swarm engagements (Sosa, 2025). Jia et al. (2019) use a system dynamics model to evaluate operational effectiveness of swarming UAV combat systems, capturing aggregate behaviour and feedback effects over time. Although these approaches offer realistic representations of engagement scenarios, they predominantly rely on computational simulation and do not yield analytically tractable expressions for key performance measures such as engagement time distributions or success probabilities under heterogeneous swarm composition.

In contrast to the above approaches, the present work focuses explicitly on the temporal accumulation of engagement processes in drone defence. Rather than modelling detailed interactions or optimising spatial deployment, we represent each neutralisation as a multi-stage stochastic process and analyse how these processes compose under sequential servicing of multiple targets. By adopting an ME framework, the proposed model enables closed-form characterisation of cumulative neutralisation time, including its distribution, moments, and quantiles, while explicitly incorporating uncertainty in swarm composition. This analytical perspective provides complementary insight to simulation-based studies by revealing how defence performance scales with swarm size and identifying the structural bottlenecks that limit system effectiveness.

ME and phase-type distributions are well established in queueing theory and stochastic service systems. In this work, we build upon these existing tools by utilising them in a non-conventional setting. Rather than modelling a system with homogeneous service requirements, we construct a layered formulation that combines a multi-stage ME representation of sequential neutralisation with a Bernoulli-driven mixture over the unknown number of true threats embedded within a finite swarm containing decoys. This structure captures both engagement dynamics and threat uncertainty within a unified analytical framework. Furthermore, the ME representation provides structural flexibility, enabling straightforward modification of engagement stages and parameters without requiring a complete reformulation of the model.

The main contributions of this paper are as follows:

- We introduce a novel analytical framework for modelling drone swarm defence systems using ME representations, a modelling approach not previously applied in this area.
- The proposed framework yields tractable, closed-form expressions for key performance measures while explicitly distinguishing between threat and decoy drones.
- The model captures the probabilistic nature of warfare, including uncertainty in resource availability, sensing and lock-on success, and environmental effects.
- The results provide concrete operational insights, showing that swarm size has a substantially stronger impact on defence success than threat-decoy composition. Under the baseline modelling assumptions, swarms exceedingly roughly ten drones begin to

surpass single-defender swarm-size capacity. This occurs even when the swarm consists of up to 90% decoy drones.
- Despite identical average neutralisation time, higher decoy ratios increase dispersion in neutralisation durations, leading to reduced predictability in defender performance. Defender effectiveness is governed by the slowest stages, which in the adopted model correspond to tracking and neutralisation.

Nomenclature

$\boldsymbol{B}$	subgenerator matrix of the ME representation
$\boldsymbol{e}'$	summation column vector of ones
$\mathbb{E}[\cdot]$	expectation operator
$f(t)$	probability density function
$F(t)$	cumulative distribution function
$\boldsymbol{J}$	strict super diagonal shift matrix
K	index of the last threat drone in the neutralisation sequence
N	total number of drones in the initial swarm
p	probability that an incoming drone is a threat drone
$\boldsymbol{p}$	initial state row vector of the ME representation
$\mathbb{P}(\cdot)$	probability measure
$\boldsymbol{Q}$	infinitesimal generator matrix
S_i	stage i associated with the single-drone neutralisation
t_q	q-quantile of the neutralisation time
$(\cdot)'$	transpose operator
$\boldsymbol{\delta}_1$	first canonical row basis vector
λ_{i}	transition rate associated with stage i
σ	standard deviation
τ	prescribed time threshold for successful drone neutralisation

2. Methodology

This section describes the modelling approach used to represent the neutralisation of an incoming drone swarm. A single-drone neutralisation process is treated as a sequence of consecutive operational stages, namely *detection and identification*, *tracking*, *engagement preparation*, and *neutralisation*, whose combined duration is modelled using ME distributions. This representation allows complex, multi-stage processing behaviour to be captured within a unified analytical framework.

The modelling procedure is developed progressively, beginning with the neutralisation of a single drone and then extending to scenarios involving multiple drones. Within the swarm, drones may correspond either to genuine threats or to decoys that are intended to consume defensive resources without posing direct harm, and both types are incorporated within the same framework. Building on the single-drone formulation, the methodology is extended to account for arbitrary swarm sizes and mixed compositions of threat and decoy drones.

In this formulation, a non-discriminatory defence assumption is adopted, whereby threat and decoy drones are processed identically by the defender. This reflects operational settings in which decoy UAV platforms are intentionally designed to emulate the observable characteristics of threat systems, including shape, flight behaviour, and signatures, thereby limiting reliable discrimination prior to engagement (Anokhin & Faragasso, 2025; Espreso Global, 2024). Under such conditions, both threat and decoy drones consume defensive resources in the same manner and are therefore modelled within a unified framework.

2.1. ME Distributions

ME distributions are a class of probability distributions characterised by rational Laplace transforms (Bladt & Neuts, 2003; van de Liefvoort & Heindl, 2005). They generalise the family of phase-type distributions. An ME distribution that can be interpreted as the absorption time in a continuous-time Markov chain with m transient states is of phase type.

Consider a continuous-time Markov chain whose state space is partitioned into m transient states and a single absorbing state. Without loss of generality, index the transient states as $1, \dots,$ m and the absorbing state as $m + 1$.The infinitesimal generator can then be written in the standard block form

$$\boldsymbol{Q} = \begin{bmatrix} \boldsymbol{B} & -\boldsymbol{B}\boldsymbol{e}' \\ \boldsymbol{0} & 0 \end{bmatrix} \tag{1}$$

where $\boldsymbol{B}$ is an $m \times m$ matrix, $-\boldsymbol{B}\boldsymbol{e}'$ is an m-dimensional column vector, and $\boldsymbol{0}$ is an m-dimensional row vector of zeros. Matrix $\boldsymbol{B}$ describes the transitions between transient states and is called the subgenerator matrix. Vector $-\boldsymbol{B}\boldsymbol{e}'$ describes the transitions from the transient states to the absorbing state. Vector $\boldsymbol{e}'$ is a column vector of ones and is called the summation vector.

Throughout this paper, boldface lowercase letters denote row vectors, while boldface uppercase letters denote matrices.

Let $\boldsymbol{p}$ be an m-dimensional row vector that describes the initial distribution of the transient states. This vector is called the initial state vector. Starting from state i, the probability of being in state j at time t is $\exp(\boldsymbol{B}\mathrm{t})_{ij}$ (Bladt, 2005). The probability that the process has not yet jumped to the absorbing state by time t is expressed as

$$\begin{aligned} 1 - F(t) &= \sum_{i=1}^{m}\sum_{j=1}^{m} p_i \exp(\boldsymbol{B}t)_{ij} \\ &= \boldsymbol{p}\exp(\boldsymbol{B}t)\,\boldsymbol{e}' \end{aligned} \tag{2}$$

Thus, the Cumulative Distribution Function (CDF), $F(t)$, and the Probability Density Function (PDF), $f(t)$, are given as

$$\mathrm{F(t)} = 1 - \boldsymbol{p}\exp(\boldsymbol{B}t)\,\boldsymbol{e}', \qquad \mathrm{t} \geq 0 \tag{3}$$

$$\mathrm{f}(t) = -\boldsymbol{p}\,\exp(\boldsymbol{B}t)\,\boldsymbol{B}\boldsymbol{e}', \qquad t \geq 0 \tag{4}$$

The results shown in (3) and (4) hold for any $\langle \boldsymbol{p}, \boldsymbol{B}, \boldsymbol{e}' \rangle$ as long as they satisfy the properties of distribution and density functions.

ME distributions have been widely adopted in applied probability and engineering, including risk theory, reliability analysis, and survivability modelling of communication and wireless networks. They enable tractable representation of non-exponential timing behaviour while preserving Markovian structure (Fackrell, 2009; He & Zhang, 2007; Horvath et al., 2016; Kaja et al., 2021; Tekinay et al., 2020). They are also closely related to phase-type distributions and have been extensively used in queueing-based performance modelling (Asmussen, 2003; Latouche & Ramaswami, 1999; Neuts, 1981).

2.2. Single-drone Neutralisation Process

Drone neutralisation systems span a wide range of sensing and mitigation methods. Comprehensive surveys of drone neutralisation technologies are available in the literature (Joint Research Centre (European Commission), 2025; Khawaja et al., 2025; Liang et al., 2025; Petsioti et al., 2024; Tu et al., 2024; Zhu et al., 2025). Information on commercial counter-drone systems can be found in existing review studies (González-Jorge et al., 2024). Regardless of the specific technologies employed, such systems inherently operate as a sequential process, since a drone must first be detected, subsequently identified and tracked, and only then subjected to an engagement decision and neutralisation. This staged engagement logic is explicitly reflected in doctrinal counter-UAS frameworks (Joint Air Power Competence Centre, 2021).

The durations of these stages are inherently stochastic, as they depend on a range of uncertain operational factors. Detection and identification latencies are influenced by environmental conditions, clutter, propagation effects, and sensor performance, while tracking duration depends on target dynamics and sensing quality. Engagement preparation and engagement times are further affected by resource availability, system load, and operational constraints. These sources of variability motivate a probabilistic, stage-based modelling approach.

We model the time required to neutralise a single drone as a four-stage hypoexponential distribution, corresponding to the stages of *detection and identification*, *tracking*, *engagement preparation*, and *neutralisation* as illustrated in Figure 1.

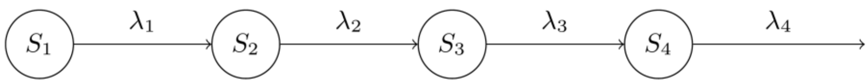

Figure 1. Markov representation of a multi-stage drone neutralisation process.

The initial state vector, the subgenerator matrix, and the summation vector corresponding to the single-drone neutralisation time are given by

$$\boldsymbol{p} = [1 \quad 0 \quad 0 \quad 0], \qquad \boldsymbol{B} = \begin{bmatrix} -\lambda_1 & \lambda_1 & 0 & 0 \\ 0 & -\lambda_2 & \lambda_2 & 0 \\ 0 & 0 & -\lambda_3 & \lambda_3 \\ 0 & 0 & 0 & -\lambda_4 \end{bmatrix}, \qquad \boldsymbol{e}' = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \tag{5}$$

From an operational perspective, the vector $\boldsymbol{p}$ indicates that each engagement starts in the detection and identification stage. The subgenerator matrix $\boldsymbol{B}$ encodes the progression of the defence system through successive stages (detection, tracking, engagement preparation, and neutralisation), where each diagonal entry represents the rate of leaving a stage and off-diagonal elements represent transitions to the next stage. The vector $\boldsymbol{e}'$ aggregates all transient states and enables computation of the probability that the engagement process has completed, i.e., that the drone has been neutralised.

The probability that a drone is neutralised by time t is given by $\mathbb{P}(\mathrm{T} \leq t)$ and is obtained by substituting the parameters in (5) into the ME distribution function in (3). Let τ denote the time threshold within which a threat-drone must be neutralised to prevent harm. The probability of successful neutralisation is then given by $F(\tau)$.

2.3. Multi-drone Neutralisation Process

An extension of the single-drone neutralisation model is developed for the case where the system is initially presented with an arbitrary number of threat drones. Neutralisation is assumed to proceed sequentially, so that the total neutralisation time for k drones is given by the sum of k independent and identically distributed single-drone neutralisation times. Leveraging the closure property of the ME class under convolution, we construct an equivalent ME representation described by a single parameter triplet $\left\langle \boldsymbol{p}_k, \boldsymbol{B}_k, \boldsymbol{e}_k' \right\rangle$ that characterises the neutralisation of k consecutive drones. This representation provides a compact state-space formulation and enables exact analysis of the resulting service-time distribution for arbitrary swarm sizes.

Let the neutralisation time of a single drone follow a ME distribution with representation $\langle \boldsymbol{p}, \boldsymbol{B}, \boldsymbol{e}' \rangle$. Then, the total time required to neutralise k consecutive drones admit a ME representation $\left\langle \boldsymbol{p}_k, \boldsymbol{B}_k, \boldsymbol{e}_k^{'} \right\rangle$ where $\boldsymbol{B}_k$ is constructed to capture the sequential progression through k identical neutralisation phases, $\boldsymbol{p}_k$ denotes the corresponding initial-state probability vector, and $\boldsymbol{e}_k'$ is a column vector of ones of appropriate dimension. Operationally, this construction represents the defender processing drones one at a time, where each phase corresponds to a full detection-to-neutralisation cycle for a single drone.

Let $\boldsymbol{J}_k$ denote a $\mathrm{k} \times \mathrm{k}$ strict superdiagonal shift matrix, defined by $(\boldsymbol{J}_k)_{i,i+1} = 1$ for $\mathrm{i} = 1, \ldots, \mathrm{k} - 1$, and zeros elsewhere. Matrix $\boldsymbol{J}_k$ has ones on the first superdiagonal and zeros on the diagonal and all other entries. We express the k-fold subgenerator matrix as

$$\boldsymbol{B}_k = \boldsymbol{I}_k \otimes \boldsymbol{B} + \boldsymbol{J}_k \otimes (-\boldsymbol{B}\boldsymbol{e}'\boldsymbol{p}) \tag{6}$$

The first term $\boldsymbol{I}_k \otimes \boldsymbol{B}$ represents k independent copies of the single-drone engagement process, while the second term $\boldsymbol{J}_k \otimes (-\boldsymbol{B}\boldsymbol{e}^{'}\boldsymbol{p})$ links these copies sequentially, ensuring that the system transitions to the next drone only after completing the current neutralisation.

The resulting $\boldsymbol{B}_k$ matrix has a block upper-bidiagonal structure, with each diagonal block equal to the single-drone subgenerator matrix $\boldsymbol{B}$ and each super diagonal block representing the transition between consecutive drone-neutralisation stages. The matrix $\boldsymbol{B}_k$ has dimension $\mathrm{km} \times \mathrm{km}$ where m denotes the dimension of the single-drone subgenerator matrix $\boldsymbol{B}$. For completeness, the corresponding block matrix form of $\boldsymbol{B}_k$ is shown below.

$$\boldsymbol{B}_k = \begin{bmatrix} \boldsymbol{B} & -\boldsymbol{B}\boldsymbol{e}'\boldsymbol{p} & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{B} & -\boldsymbol{B}\boldsymbol{e}'\boldsymbol{p} & \cdots & \boldsymbol{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{B} & -\boldsymbol{B}\boldsymbol{e}'\boldsymbol{p} \\ \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{0} & \boldsymbol{B} \end{bmatrix} \tag{7}$$

In this representation, each diagonal block corresponds to the internal stages of neutralising a single drone, while each superdiagonal block represents the transition from completing one drone engagement to initiating the next. This directly reflects the sequential nature of the defence system.

The initial state vector $\boldsymbol{p}_k$ is given in (8). The triplet $\left\langle \boldsymbol{p}_k, \boldsymbol{B}_k, \boldsymbol{e}_k^{'} \right\rangle$, where $\boldsymbol{e}_k'$ is a column vector of ones of length km, thus defines a valid ME representation for the neutralisation time of k consecutive drones.

$$\boldsymbol{p}_k = [\boldsymbol{p} \quad 0 \quad \ldots \quad 0] \tag{8}$$

This structure indicates that the process always begins with the first drone entering the detection stage, after which the system progresses through subsequent drones sequentially.

We assume that the process begins when a swarm of N drones appear at time 0. The sequential neutralisation process by a single defender is illustrated in Figure 2 where each state represents the remaining number of drones.

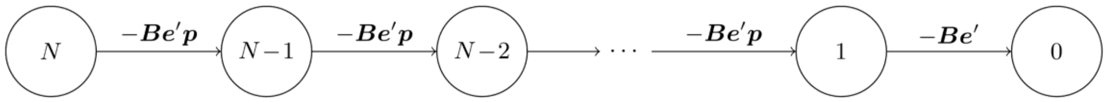

Figure 2. ME state-transition representation of the drone neutralisation process, starting from an initial swarm size of N.

In contrast to the single-drone case, the defender must neutralise all N drones in a sequential manner by time τ to be successful. The aggregate neutralisation time distribution function is obtained by substituting (7) and (8) into (3). The probability of successful neutralisation is then given by $F_k(\tau) = 1 - \boldsymbol{p}_k \exp(\boldsymbol{B}_k\tau)\ \boldsymbol{e}_k'$.

2.4. Threat and Decoy Drone Modelling

Heterogeneous drone populations are incorporated into the proposed ME framework by distinguishing between threat and decoy drones. Specifically, each incoming drone is assumed to be a threat with probability

$$p \triangleq \mathbb{P}(\text{drone is a threat}) \tag{9}$$

and a decoy with complementary probability $1 - p$.

The defender does not have the ability to distinguish between decoy and threat drones. Therefore, the defender adopts the same drone neutralisation process to both kinds of drones. At each engagement, the defender selects a drone to neutralise at random. We are interested in the probability of neutralising all the threat drones by some time τ in the presence of decoy drones. As a result, the relative ordering of threat and decoy drones within the neutralisation sequence influences the total time required to neutralise all threats. We define the random variable

$$K = \max\{i \leq N : X_i = 1\}, \quad X_i \sim \text{Bernoulli}(p), \quad \text{i.i.d.} \tag{10}$$

where K represents the index of the last threat drone in a sequence of N drones and follows the distribution of the last success index in a finite sequence of Bernoulli trials. Its Probability Mass Function (PMF) is given by

$$\mathbb{P}(K = k) = \begin{cases} (1-p)^N, & k = 0, \\ p(1-p)^{N-k}, & k = 1, 2, .., N \end{cases} \tag{11}$$

It follows that all threat drones are neutralised by time τ with probability

$$\mathbb{P}(\text{success by } \tau) = \mathbb{P}(\mathrm{K} = 0) + \sum_{\mathrm{k}=1}^{\mathrm{N}} \mathbb{P}(K = k) F_k(\tau) \tag{12}$$

As shown in (12), the overall success probability is a weighted mixture of the neutralisation time distributions corresponding to different numbers of required threat neutralisations. The success criterion in (12) can be naturally incorporated within the ME framework by allowing for a defective initial condition. In particular, with probability $(1-p)^N$, the swarm contains no threat drones, corresponding to a point mass at zero neutralisation time. With the complementary probability, at least one threat drone is present, and the time to success follows a ME distribution obtained as a mixture over the possible numbers of required threat neutralisations. This formulation preserves the analytical tractability of the ME framework while explicitly accounting for uncertainty in swarm composition.

Let $\boldsymbol{\delta}_1^{(k)}$ denote the first canonical row basis vector of dimension k. It follows that

$$\boldsymbol{p}_k = \boldsymbol{\delta}_1^{(k)} \otimes \boldsymbol{p} \tag{13}$$

Let $\boldsymbol{B}_{\text{mix}}$ denote the block-diagonal subgenerator matrix mixing $\{\boldsymbol{B}_k\}_{k=1}^{N}$. It is given by

$$\boldsymbol{B}_{\text{mix}} = \boldsymbol{B}_N \oplus \boldsymbol{B}_{N-1} \oplus \ldots \oplus \boldsymbol{B}_1 \tag{14}$$

Operationally, $\boldsymbol{B}_{\text{mix}}$ captures all possible engagement scenarios corresponding to different numbers of threat drones within the swarm, with each block representing a specific case where a given number of drones must be neutralised.

Each $\boldsymbol{B}_k$ has a dimension $\mathrm{km} \times \mathrm{km}$. Consequently, $\boldsymbol{B}_{\text{mix}}$ is a square matrix of dimension $mN(N+1)/2$. Let $w_k \triangleq \mathbb{P}(K = k)$. The initial state vector of the ME representation is then given by

$$\boldsymbol{p}_{\text{mix}} = [w_N \boldsymbol{p}_N \quad w_{N-1} \boldsymbol{p}_{N-1} \quad \cdots \quad w_1 \boldsymbol{p}_1] \tag{15}$$

The weighting coefficients w_k reflect the likelihood of each scenario, so that the overall model accounts for uncertainty in swarm composition by combining all possible threat realisations.

By stacking the k consecutive ME representations into the block-diagonal subgenerator matrix $\boldsymbol{B}_{\text{mix}}$ and the weighted initial state vector $\boldsymbol{p}_{\text{mix}}$, the mixture admits a closed-form ME representation. The success-time CDF (time to neutralise all threats) can be written as

$$F_{succ}(t) = 1 - \boldsymbol{p}_{\text{mix}} \exp(\boldsymbol{B}_{\text{mix}} t)\, \boldsymbol{e}'_{\text{mix}}, \qquad \mathrm{t} \geq 0 \tag{16}$$

The resulting success-time distribution is defective, with an atom of size w_0 at $t = 0$; the remaining continuous component is represented by the ME triplet $\left\langle \boldsymbol{p}_{\text{mix}}, \boldsymbol{B}_{\text{mix}}, \boldsymbol{e}'_{\text{mix}} \right\rangle$, where $\boldsymbol{e}'_{\text{mix}}$ is a column vector of ones of length $mN(N+1)/2$.

3. Results

This section applies the analytical expressions derived above to evaluate a drone swarm defence system. The stage-wise rate parameters are initialised as $\lambda_1 = 5$, $\lambda_2 = 0.5$, $\lambda_3 = 2$, and $\lambda_4 = 0.5$, corresponding to mean durations of 200 ms for *detection and identification*, 2 s for *tracking*, 500 ms for *engagement preparation*, and 2 s for *neutralisation*, respectively.

Modern counter-UAS systems exhibit rapid operational timelines consistent with the parameter values adopted. Implementations indicate that detection and identification are performed within sub-second timescales, with recent perception systems demonstrating single-frame processing latencies on the order of 50-100 ms and conventional approaches often exceeding 200 ms (DroneShield Ltd., 2024; Joint Air Power Competence Centre, 2021; Lockheed Martin, 2025; Zhu et al., 2025). Tracking is typically realised through successive updates of target state information, and practical systems describe track formation over short temporal windows involving multiple observations (Argustec Information Technology Co., Ltd., 2024; Blighter Surveillance Systems and Chess Dynamics and Enterprise Control Systems, 2017; Joint Research Centre (European Commission), 2025).

Engagement preparation is generally performed rapidly following track establishment, with system descriptions indicating short handoff and targeting delays prior to countermeasure deployment (Blighter Surveillance Systems and Chess Dynamics and Enterprise Control Systems, 2017; Lockheed Martin, 2025).

Countermeasure mechanisms further indicate that neutralisation timelines depend on the engagement modality: electronic attack and high-power RF techniques are often characterised by near-instantaneous or sub-second effects once applied, whereas directed energy approaches and kinetic responses may require longer but still short-duration engagements on the order of seconds (DroneShield Ltd., 2024; Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025).

3.1. Impact of N and p

This subsection examines the impact of swarm size N and per-drone threat probability p on the distribution of neutralisation time. Let $F_{succ}(t)$ in (16) denote the CDF of the swarm neutralisation time (time to neutralise all threat drones) for a swarm size of N with per-drone threat probability p. By the ME representation, the density function is readily obtained as $f_{succ}(t)$.

Figure 3 illustrates the effect of increasing swarm size under the worst-case assumption that all incoming drones are threat drones $p = 1$. As the swarm size grows, the neutralisation time distribution shifts progressively to the right, indicating that achieving high-confidence neutralisation becomes increasingly time-consuming.

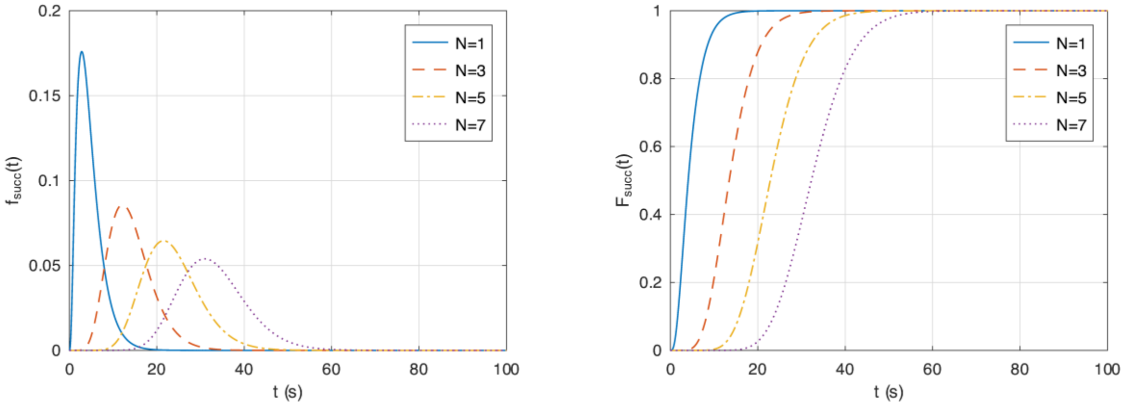

(a) **(b)**

Figure 3. Swarm neutralisation time distribution for $N \in \{1, 3, 5, 7\}$ for all-threat case $p = 1$. **(a)** Probability density function $f_{succ}(t)$ and **(b)** cumulative distribution function $F_{succ}(t)$.

Figure 4 presents a joint view of the discrete and continuous effects induced by mixed swarms with $N = 7$. Figure 4(a) shows the probability mass function of the last threat index K for $p \in \{0.15, 0.35, 0.55, 0.75\}$, while Figure 4(b) depicts the corresponding cumulative distribution functions of the neutralisation time.

A first observation is the presence of a jump in the CDF at time zero, whose magnitude increases as p decreases. This jump corresponds to the probability mass at $K = 0$, i.e., the event that no threat drones are present in the swarm. In practical terms, this represents scenarios where all drones are decoys, so the defender achieves immediate success without performing any neutralisation. The size of this atom is directly reflected in Figure 4(a) at $k = 0$ by the probability mass $(1 - p)^N$. This correspondence provides a consistent interpretation of the defective behaviour observed in the time-domain CDFs.

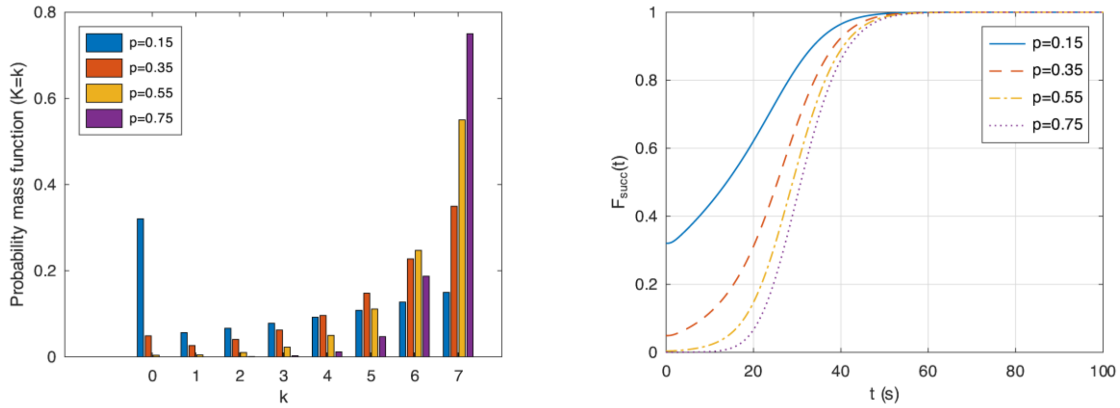

(a) **(b)**

Figure 4. Swarm size $N = 7$ with threat-drone probability $p \in \{0.15, 0.35, 0.55, 0.75\}$. **(a)** Probability mass function of the last threat index K and **(b)** cumulative distribution function $F_{succ}(t)$.

In addition, the tail behaviour of the neutralisation time distributions exhibits a notable convergence pattern. Although the CDFs begin at different initial levels, their upper tails approach that of the $p = 0.75$ case, indicating similar high-confidence neutralisation times. This behaviour contrasts with the earlier scenario of fixed p and increasing N where the upper tail progressively shifted rightward. The convergence of the tails can again be understood through Figure 4(a): the PMF of K is right-skewed (excluding $k = 0$ mass), with non-negligible probability mass at large indices $k = 6$ and $k = 7$. As a result, even swarms with relatively low per-drone threat probability can, with appreciable likelihood, induce neutralisation dynamics comparable to those of more heavily armed swarms.

3.2. Operational Significance

We are interested in answering the following question: will a defender be successful for a given swarm size N and per-drone threat probability p? To address this, we define a success measure as

$$\mathbb{P}(\text{success}) = F_{succ}(\tau) \tag{17}$$

where success corresponds to neutralising the swarm within a prescribed engagement time τ. Assuming detection at standoff ranges on the order of several hundred metres and drone speeds of tens of metres per second, this yields an engagement window on the order of tens of seconds. Accordingly, $\tau = 45$ s is adopted as a representative operational threshold.

Reported counter-UAS capabilities indicate that detection and engagement ranges in practice span from a few hundred metres to several kilometres depending on the sensing modality and system configuration. For example, compact radar systems report drone detection ranges on the order of 250-400 m, whereas integrated counter-UAS platforms provide detection ranges up to several kilometres and engagement or recognition ranges around 1 km (Aaronia AG, 2024; A-Bots, 2026; Argustec Information Technology Co., Ltd., 2024; Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025; SpotterRF / Spotter Global, 2024). At the same time, representative UAS platforms, including loitering munitions such as the Shahed-136, operate at speeds on the order of tens of metres per second (Army Recognition, 2024; Army Technology, 2024). These values correspond to engagement timelines ranging from on the order of 10-20 s for short-range sensing to over a minute for longer-range detection, placing a 45 s threshold within a realistic and representative operational regime.

Figure 5 presents the behaviour of the success probability $F_{succ}(45\text{ s})$ with respect to N and p. Note that the plots do not depict the cumulative distribution function with respect to τ; rather, they show the cumulative distribution function evaluated at $\tau = 45$ s.

When all drones are threat drones ($p = 1$), the defender can reliably neutralise swarms up to seven drones, achieving a success probability of approximately 0.93. Reducing p to 0.25 yields a success probability of 0.98. Even when p is as low as 0.1, the defender can handle a swarm size of approximately ten drones with high confidence. For swarm sizes exceedingly roughly ten drones, the deployment of additional defenders becomes necessary to maintain a high probability of successful neutralisation.

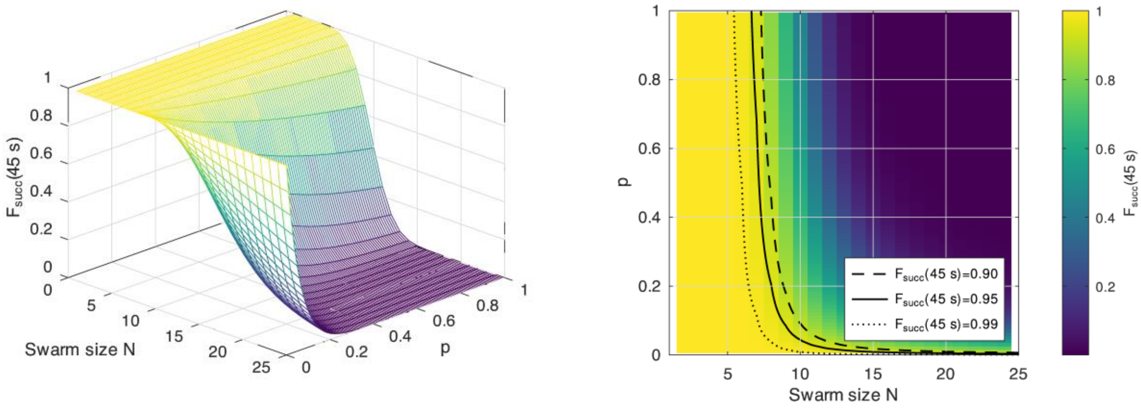

Figure 5. Probability of successful neutralisation $F_{succ}(45\text{ s})$ as a function of N and p. **(a)** Success-probability surface over the (N, p) parameter space and **(b)** corresponding heatmap with 90%, 95%, and 99% success boundaries.

Instead of a swarm of seven all-threat drones (which represents the defender's capacity for purely threatening swarms), introducing seven decoy drones to form a fourteen-drone swarm increases the probability of defence evasion sharply from 7% to 91%.

As evident from the heatmap in Figure 5(b), swarm size N has a substantially stronger impact on defence success than the threat probability p (provided p is not vanishingly small). The success boundaries are primarily driven by variations in N rather than p, indicating a higher sensitivity to swarm size than to threat-drone probability.

Table 1 presents several swarm configurations that achieve the same defence evasion probability. Two interesting observations can be made. First, the presence of a single threat drone in a swarm of size fifty can be as effective as a swarm composed entirely of ten threat drones.

Table 1. Different swarms achieving the same evasion probability.

Swarm	N	p	$\mathbb{E}[\boldsymbol{threats}]$	$\mathbb{E}[\boldsymbol{decoys}]$	$1 - F_{succ}(45\text{ s})$
1	10	1	10	0	0.56
2	11	0.48	5	6	0.56
3	15	0.14	2	13	0.56
4	50	0.02	1	49	0.56

Second, the effectiveness of decoys is strongly dependent on p. Taking Swarm 1 as the baseline configuration, Swarm 2 replaces nearly half of the threat drones with decoys. In this case, the number of introduced decoys is approximately equal to the number of removed threat drones.

Swarm 3 removes eight threat drones and introduces thirteen decoys. In the extreme case of Swarm 4, forty-nine decoys are introduced to substitute nine threat drones.

The trend indicates that adding decoys in place of threats becomes progressively less effective as p decreases. Therefore, maintaining a fixed level of evasion probability requires a disproportionately large increase in swarm size N at very small values of p. This behaviour is reflected in the right-tail structure of the boundaries shown in Figure 5(b), where increases in N are progressively less likely to cross the success boundary at low values of p.

Consequently, decoy-based strategies are most effective in an intermediate operating regime where p remains sufficiently large to avoid this tail behaviour. Beyond this regime, further reductions in p yield diminishing returns, as increasingly large swarms are required to sustain the same evasion performance.

From a cost standpoint, assuming decoys cost 10% of a threat drone, Swarm 3 is the most economical one: the total attacker cost is reduced by approximately 67% (corresponding to the equivalent of 3.3 threat drones). Even at a higher decoy cost of 20%, the cost reduction remains substantial at roughly 54%. This comparison reflects only unit platform costs and does not account for interceptor usage, sensing resources, operational workload, or command-and-control processes. A more comprehensive study would be required to capture the full cost implications.

3.3. Dispersion Effects in Swarm Neutralisation Time

This section examines the temporal behaviour of the neutralisation-time distribution, rather than evaluating performance at a fixed prescribed time. In particular, the dispersion characteristics of the neutralisation-time distribution are analysed. In this context, high quantiles and Coefficient of Variation (CV) serve as informative indicators of dispersion.

Let t_q denote the q-quantile of the neutralisation-time distribution, defined as

$$t_q \triangleq \inf\{\, t \geq 0 : F_{\text{succ}}(t) \geq q\} \tag{18}$$

The quantity t_{95} therefore specifies the minimum time required to successfully neutralise a swarm in 95% of engagements.

Figure 6 presents four mean-equivalent swarm configurations, all exhibiting nearly identical expected neutralisation times of 18 ± 0.05 s, yet displaying markedly different levels of distributional dispersion. The expected neutralisation time is obtained directly from the ME representation as

$$\mathbb{E}[T] = -\boldsymbol{p}_{\text{mix}}\, \boldsymbol{B}_{\text{mix}}^{-\mathbf{1}}\, \boldsymbol{e}_{\text{mix}} \tag{18}$$

where T denotes the neutralisation time.

Table 2 provides detailed information for the swarms shown in Figure 6. All four swarm configurations correspond to scenarios in which the defender is expected to succeed with high confidence. On average, the defender requires 18 s to finalise neutralisation in all cases. Nevertheless, the configurations exhibit substantially different t_{95} values. In particular, neutralisation is completed within ≈ 50.5 s for Swarm 1 and within ≈ 28.7 s for Swarm 4 in 95% of engagements.

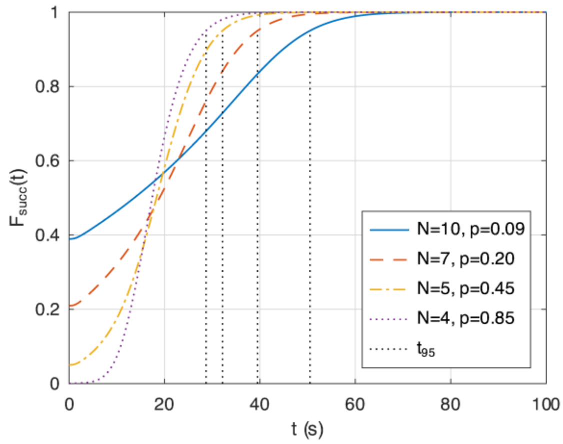

Figure 6. CDFs of four mean-equivalent configurations with varying t_{95} quantiles.

This indicates that Swarm 4 exhibits a more concentrated neutralisation-time distribution than Swarm 1. The reduced variability observed for Swarm 4 is attributable to its large p, under which nearly all drones are threat drones and must be neutralised to achieve success. In contrast, for Swarm 1, the sequence of neutralisation events plays a more significant role: threat drones may be neutralised either before or after a large number of decoy drones. This uncertainty introduces a mixture of distinct neutralisation-time scales, resulting in increased dispersion.

This effect is also reflected in Table 2 in terms of CV, defined as the ratio of the standard deviation σ to the mean $\mathbb{E}[T]$. The second moment of the neutralisation time is obtained using $\mathbb{E}[T^2] = \mathbf{2}\, \boldsymbol{p}_{\text{mix}}\, \boldsymbol{B}_{\text{mix}}^{-2}\, \boldsymbol{e}_{\text{mix}}$.

Table 2. Equivalent-mean swarm configurations with varying dispersion.

Swarm	N	p	$\mathbb{E}[T]$	t_{95}	CV
1	10	0.09	17.98 s	50.5 s	1.03
2	7	0.2	18.04 s	39.5 s	0.74
3	5	0.45	18.04 s	32.1 s	0.48
4	4	0.85	17.97 s	28.7 s	0.33

3.4. Defence Stage Sensitivity Analysis

In this section, the sensitivity of neutralisation performance to the internal defence stages is examined. Specifically, we study how variations in the stage rates $\{\lambda_i\}$ affect both expected performance and dispersion, with the objective of identifying dominant bottlenecks and stages with limited operational leverage.

Figure 1 shows the defence stages. The baseline rate values are $\lambda_1 = 5$, $\lambda_2 = 0.5, \lambda_3 = 2, \lambda_4 = 0.5$. We are interested in evaluating the relative importance of stages. To quantify their contribution to the overall neutralisation time, we sweep individual rates whilst keeping the remaining rates at their baseline values. Specifically, each baseline rate is multiplied by factors of 0.2, 0.5, 1.25 and 2. Figure 7 presents the resulting sensitivity for a swarm of size $N = 7$ with a threat probability of 0.2.

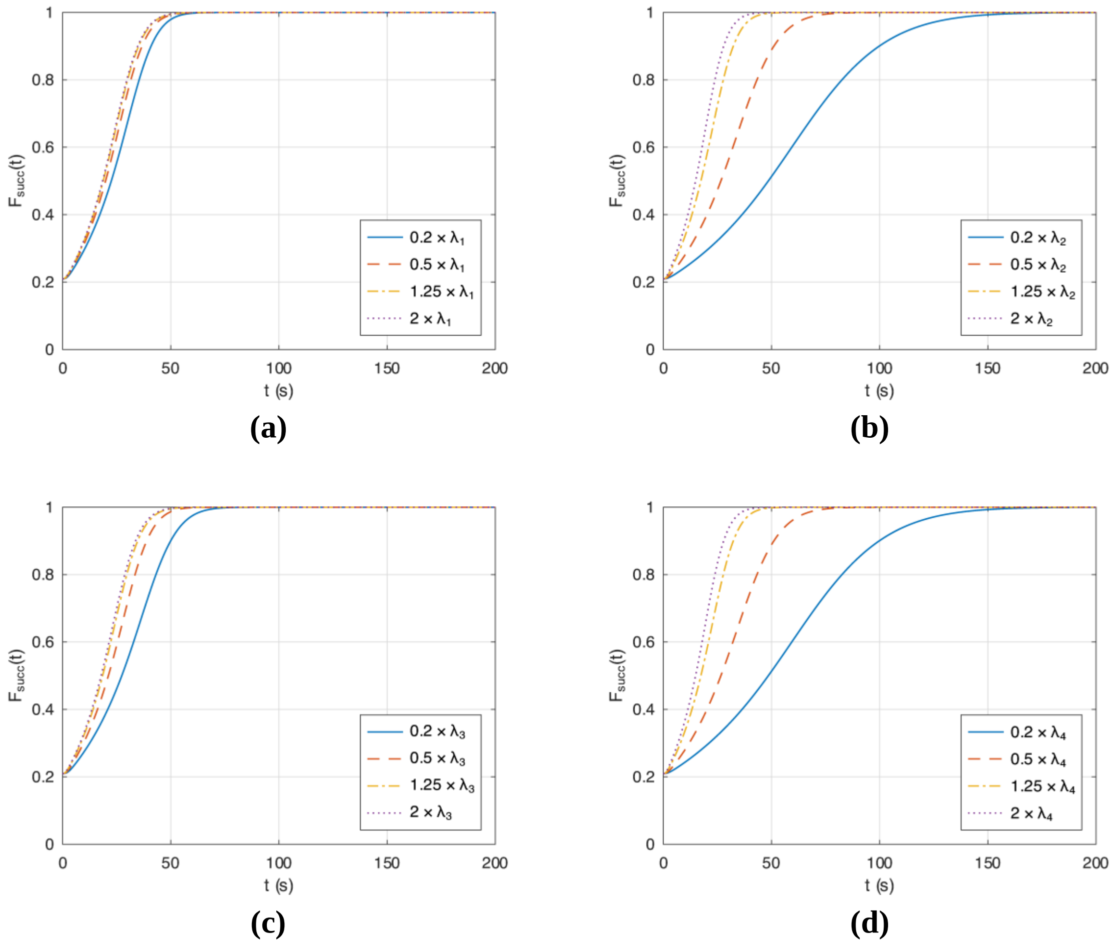

Figure 7. Defence stage sensitivity analysis for a swarm of size $N = 7$ with threat drone probability $p = 0.2$. Each subfigure shows the impact of a single defence stage. The baseline rate value is multiplied by 0.2, 0.5, 1.25, and 2, while the other parameters are fixed. **(a)** Sensitivity to λ_1, **(b)** sensitivity to λ_2, **(c)** sensitivity to λ_3 and **(d)** sensitivity to λ_4.

Sensitivity is highest for stages with the lowest baseline rates, namely stages 2 and 4. These stages behave identically because their baseline rates are the same. These two stages represent the *tracking* and *neutralisation* processes, respectively. This identical behaviour demonstrates that the ordering of stages does not affect overall neutralisation performance. Improvements to these stages yield substantially larger gains than improvements to faster stages, as they act as system bottlenecks. Once a bottleneck stage is sufficiently sped up, performance becomes dominated by the next slowest stage, and the marginal benefit diminishes.

Figure 8 presents the sensitivity analysis in terms of t_{95} and $F_{succ}(45\text{ s})$ for the same swarm configuration. The maximum observed changes in t_{95} and $F_{succ}(45\text{ s})$ over the rate sweeps are shown, further highlighting the dominant influence of bottleneck stages on overall defence performance.

With the baseline rates, are $\lambda_1 = 5,\ \lambda_2 = 0.5, \lambda_3 = 2, \lambda_4 = 0.5$, the defence success probability is 0.66 for a swarm of size $\mathrm{N} = 10$ with threat drone probability $\mathrm{p} = 0.4$. Increasing the tracking-stage rate λ_2 by 100% to a rate value of 1 raises the success probability to 0.93 while jointly increasing both λ_2 and λ_4 by 50% to rate values of 0.75 improves it to 0.97. These results demonstrate that performance is not determined solely by aggregate rate sums; rather, individual stage rates, particularly bottlenecks, play a decisive role.

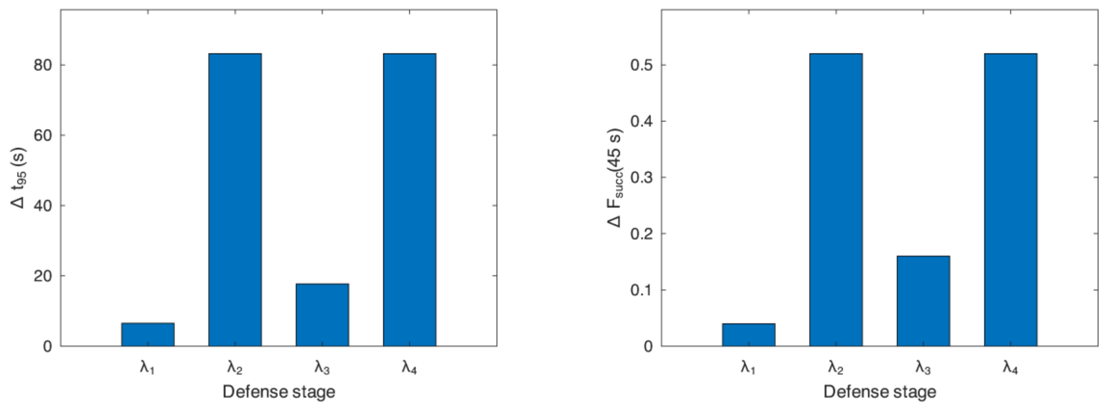

(a) (b)

Figure 8. Defence stage sensitivity analysis for a swarm of size $\mathrm{N} = 7$ with threat drone probability $p = 0.2$. **(a)** Change in t_{95} when stage rate values are multiplied by 0.2 and 2 and **(b)** change in defence success probability $F_{succ}(45\text{ s})$ when stage rate values are multiplied by 0.2 and 2.

3.5. System Performance Across Parameter Regimes

In this section, we examine the system behaviour at a holistic level. The objective is to study how the system performs as a whole and to draw practical conclusions by jointly considering swarm size N, threat probability p, and stage rates λ_i. In addition, this section examines the robustness of the observed results across different defender and attacker characteristics, as well as across varying operational time scales and scenarios.

Capacity is defined as the maximum swarm size N for which the defence success probability satisfies $\mathrm{F}_{\mathrm{succ}}(45) \geq 0.93$. This threshold represents a high-confidence operational criterion and provides a consistent basis for comparing performance across different configurations.

Table 3 presents the defender capacity under varying stage rates and threat probabilities. The results indicate that capacity is primarily governed by the stage rates, with substantial gains achieved as processing speeds increase. Faster stage configurations lead to a significant expansion in the maximum manageable swarm size. In contrast, variations in the threat probability p result in only moderate changes in capacity. Even for low values of p, the increase in capacity remains limited compared to the improvements obtained through faster stage rates, indicating that processing bottlenecks dominate system performance.

Table 3. Defender swarm-size capacity N for $\mathrm{F}_{\mathrm{succ}}(45) \geq 0.93$ under varying stage rates and threat probabilities p.

Stage Rates				Capacity N			
λ_1	λ_2	λ_3	λ_4	$p = 1$	$p = 0.7$	$p = 0.45$	$p = 0.1$
5	0.5	2	0.5	7	7	7	9
3	1	2	1	12	13	13	15
8	4	8	1	24	24	25	28
10	5	4	2	37	37	38	40

For example, when all stages—detection and identification, tracking, engagement preparation, and neutralisation—are performed on the order of one second on average, the defender's swarm-size capacity increases substantially, reaching up to 37 drones. This highlights the direct operational benefit of reducing stage durations, as even moderate improvements across stages can translate into significant gains in achievable swarm capacity.

Table 4 provides a complementary perspective by evaluating the success probability for fixed swarm sizes.

Table 4. Success probability $\mathrm{F}_{\mathrm{succ}}(45)$ under varying swarm size N, threat probability p, and stage rates.

Stage Rates				p	$\mathrm{F}_{\mathrm{succ}}(45)$		
λ_1	λ_2	λ_3	λ_4		$N = 5$	$N = 20$	$N = 35$
5	0.5	2	0.5	0.7	0.9975	<0.001	<0.001
				0.1	0.9996	0.3293	0.0678
3	1	2	1	0.7	>0.999	0.0653	<0.001
				0.1	>0.999	0.6390	0.1320
8	4	8	1	0.7	>0.999	0.9977	0.1331
				0.1	>0.999	0.9996	0.6043
10	5	4	2	0.7	>0.999	>0.999	0.9879
				0.1	>0.999	>0.999	0.9975

For small swarms (e.g., $N = 5$), the system consistently achieves near-certain success across all parameter settings, indicating operation well below capacity. However, as swarm size increases, performance degrades rapidly. For intermediate and large swarm sizes, the success probability drops sharply, often reaching values effectively indistinguishable from zero under slower stage configurations. While reducing the threat probability improves performance, this effect remains secondary and does not offset the impact of increasing N.

Overall, these results show that swarm size appears to be the dominant factor governing system performance within the considered modelling assumptions. While threat composition influences outcomes, its effect is comparatively limited. The consistency of this behaviour across all tested configurations indicates that the dominance of swarm size is a robust property of the system.

3.6. Simulation-Based Validation of the Analytical Framework

To assess the accuracy of the analytical derivations and underlying stochastic model, Monte Carlo simulations are performed and compared against the analytical results. In each trial, the number of threat drones is generated according to the Bernoulli model, and the total neutralisation time is obtained by summing independent realisations of the four-stage sequential process.

Figure 9 presents the analytical CDF together with empirical CDFs obtained from Monte Carlo simulations for increasing numbers of trials. For lower sample sizes (e.g., 100 and 200 trials), the empirical curves exhibit visible fluctuations due to sampling variability. As the number of trials increases, the Monte Carlo estimates converge toward the analytical distribution, with near-perfect agreement observed for 10,000 trials.

This agreement is observed across the full support of the distribution, indicating that the analytical model accurately captures both the central behaviour and the tail of the neutralisation time. The maximum deviation between analytical and empirical CDFs decreases with increasing sample size, demonstrating convergence of the simulation to the theoretical result.

A similar trend is observed in Figure 10, where the expected neutralisation time is evaluated as a function of the threat probability p. Even with a relatively small number of trials, the Monte Carlo estimates follow the analytical curve closely, while higher sample sizes yield nearly indistinguishable agreement across the entire range of p.

4. Discussion

This work studies heterogeneous drone swarm engagements against a single defender and characterises defender swarm-size capacity under a range of operational scenarios. Using an ME modelling framework, closed-form expressions are derived for the probability of defence success and the expected neutralisation time. The threat-drone probability, swarm size, and per-stage defence rates are treated as tunable parameters, enabling systematic exploration of their impact on system performance.

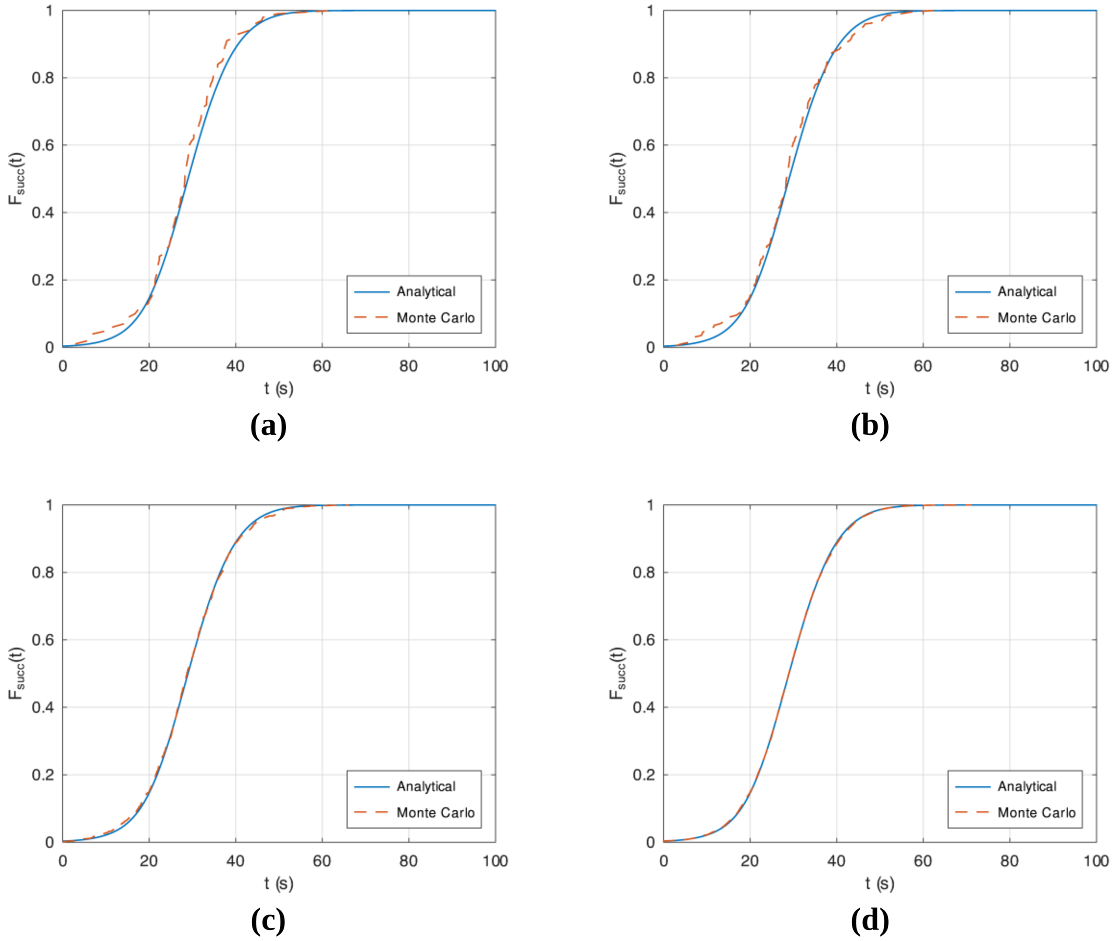

Figure 9. Analytical and Monte Carlo CDFs of the neutralisation time for a swarm of size $N = 7$ with threat probability $p = 0.55$. The Monte Carlo estimate converges to the analytical distribution as the number of trials increases. **(a)** 100 trials, **(b)** 200 trials, **(c)** 1,000 trials, and **(d)** 10,000 trials.

The results presented in the previous section should be interpreted in the context of the modelling assumptions adopted, including the sequential engagement structure and single-defender abstraction.

Our results show that swarm size is a more critical factor than the threat-drone ratio. Increasing the decoy ratio primarily increases uncertainty in defender performance rather than proportionally extending capacity. When cost considerations are taken into account, decoy compositions in the approximate range of 50%-90% appear to constitute a practically relevant regime. From a practical perspective, this range represents a reasonable operating region for the UAV types and pricing assumptions considered in this simplified model. For extremely large swarm sizes and vanishing threat probabilities, economic and logistical constraints are expected to dominate, and a more comprehensive cost analysis would be required to obtain a clearer system-level picture.

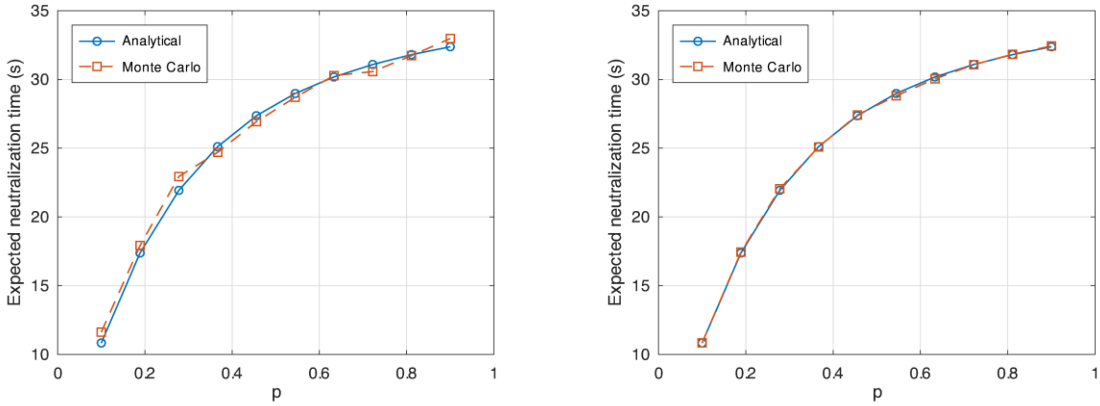

(a) **(b)**

Figure 10. Expected neutralisation time as a function of the threat probability p for a swarm of size $N = 7$, with p ∈ [0.1,0.9]. Analytical results are compared with Monte Carlo estimates. **(a)** 200 trials and **(b)** 10,000 trials, showing improved agreement with increasing sample size.

The defender's swarm-size capacity is largely governed by the bottleneck stage within the sequential defence process. Improving the slowest stage yields the most significant performance gains, whereas accelerating non-limiting stages provides diminishing returns. Given the growing number of reported drone attacks, this result suggests that each defence stage must operate, on average, on the order of one second or faster. If this requirement cannot be met, deploying additional defenders with overlapping coverage becomes necessary to maintain a high probability of successful defence.

Several modelling assumptions should be noted. The defender is assumed to be unable to distinguish between decoy and threat drones. This assumption is intentional, as the objective of this work is to quantify the impact of high-volume swarms, and it has been reported that decoy drones may exhibit signatures and footprints similar to those of threat drones in operational environments. The problem of reliable decoy-threat discrimination is therefore considered orthogonal to the present analysis. In addition, cooperative drone behaviours and advanced swarm tactics are not modelled, and the framework focuses exclusively on a single defender.

In practice, however, defence systems may exhibit partial discrimination capability. This can be incorporated within the proposed framework by augmenting the single-target service model with an initial classification stage, represented as an additional phase in the ME structure, after which non-threatening targets exit the process early while predicted threats proceed to the remaining engagement stages. In this interpretation, the modification corresponds to an adjusted initial distribution and transition structure with an early absorbing branch, while preserving the overall probabilistic mixture formulation. Such a formulation reduces the effective load on the defender and shifts the system from a volume-driven regime toward a threat-driven regime as discrimination improves. Accordingly, the present formulation can be interpreted as a conservative upper-bound stress case.

It should also be noted that, even for a single target, the neutralisation process may in practice be more complex than the four-stage sequential representation adopted here. Real-world defence chains may include feedback loops, repeated identification or tracking updates, and partial overlap between functions enabled by sensor fusion and pipelined processing. Such effects could be represented through a more general state-transition structure, for example by introducing backward transitions between phases or additional coupled states within the underlying ME model. While this would provide a richer description of single-target engagement dynamics, it would also increase model complexity substantially and move the analysis away from the parsimonious staged structure adopted in this work. For this reason, we retain the present four-stage sequential abstraction as a first-order representation of the single-target defence timeline and leave explicitly feedback-driven or partially parallel stage dynamics for future work.

Furthermore, the defence process is modelled as strictly sequential, corresponding to a single effective engagement capability. This assumption provides a conservative baseline and is consistent with a single-server interpretation in queueing terms. In practice, modern systems may support partial parallelism, including concurrent tracking or multi-target engagement capabilities, which would effectively increase service capacity and shift the system bottleneck. From a queueing perspective, transitioning to a multi-server regime reduces the effective service time per target and increases the maximum swarm size that can be handled within a given time window. Asymptotically, this leads to an approximately linear scaling in expected defender capacity with the number of parallel engagement resources. In the limiting case of high parallelism, system behaviour becomes increasingly governed by the slowest stage rather than the sum of all stages. Importantly, the qualitative insights of this work, including the identification of bottleneck stages and the dominant role of swarm size, remain unchanged, and the reported results can be interpreted as conservative lower-bound estimates of defender performance.

In addition, multi-layered or networked defence architectures can be interpreted within this framework as interconnected service stages or parallel service nodes operating across different layers. Such configurations effectively distribute the engagement process spatially and functionally, further increasing system capacity and resilience. A detailed analytical treatment of layered or networked defence structures is left for future work.

Despite these limitations, the proposed framework remains flexible and readily admits extensions. Multiple defenders, successive swarm arrivals, target differentiation, feedback mechanisms between stages, and probabilistic engagement failures can all be incorporated within the same ME modelling structure without fundamentally altering the analytical form.

In most cases, these extensions can be represented through appropriate modifications of the initial distribution vector and subgenerator matrix, which define the phase structure and transition dynamics of the process. Since all parameters are tuneable, the framework is also applicable to a wide range of operational settings, including urban warfare scenarios where reduced engagement times and tighter performance requirements are critical.

From this perspective, the present model can be viewed as one instantiation within a broader modelling class. The key advantage of this approach is that increasingly complex system behaviours can be captured through parameter and structural adjustments, rather than requiring a complete reformulation of the analysis. This provides a structured way to incorporate additional system features while preserving analytical tractability and allows the framework to be adapted to a range of defence configurations under uncertainty.

5. Conclusion

This paper introduced a ME framework for modelling a single defender engaging drone swarms composed of both decoy and threat drones, as observed in recent conflict environments. The results indicate that high-volume swarms with sizes exceeding approximately ten drones may overwhelm a typical defender under the baseline modelling assumptions, even when the decoy composition is as high as 90%. Achieving a swarm-handling capacity on the order of 40 drones requires, under the same assumptions, that the defender complete detection, identification, tracking, engagement preparation, and neutralisation stages on average within one second each.

The proposed framework provides a baseline for extending the analysis to more complex scenarios, including coordinated drone swarms, multiple defenders, and successive swarm arrivals over time. A natural extension is the incorporation of parallel or multi-target engagement capabilities, which would require extending the ME framework to multi-channel service structures and a modified mixture construction. Such an extension would enable explicit modelling of systems capable of simultaneous tracking and engagement, providing a more complete characterisation of defender capacity under modern operational conditions.

Conflict of Interest

The author declares no conflict of interest related to this work.

Funding

This research received no external funding.

Acknowledgments

The author acknowledges ASELSAN Inc. for providing research, computational, and software resources supporting this work. The author also appreciates general technical discussions with colleagues during the course of the study.

References

Aaronia AG. (2024). *Aartos anti-drone jammer*. Retrieved March 25, 2026, from https://drone-detection-system.com/sensor-types-overview/anti-drone-jammer/

A-Bots. (2026). *Drone detection in 2026*. Retrieved March 25, 2026, from https://abots.com/blog/drone-detection-2026/

Anokhin, I., & Faragasso, S. (2025). *Updated analysis of Russian Shahed-136 deployment against Ukraine*. Institute for Science and International Security. https://isis-online.org/isis-reports/may-2025-updated-analysis-of-russian-shahed-136-deployment-against-ukraine

Argustec Information Technology Co., Ltd. (2024). *Argustec wukong anti-drone system*. Argustec Information Technology Co., Ltd. Retrieved March 25, 2026, from https://www.argustecn.com/antidronesystem.html

Army Recognition. (2024). *Shahed-136 loitering munition technical data*. https://armyrecognition.com/military-products/army/unmanned-systems/unmanned-aerial-vehicles/shahed-136-loitering-munition-kamikaze-suicide-drone-technical-data

Army Technology. (2024). *Shahed-136 kamikaze UAV (Iran)*. https://www.army-technology.com/projects/shahed-136-kamikaze-uav-iran/

Asmussen, S. (2003). *Applied probability and queues*. Springer.

Bladt, M. (2005). A review on phase-type distributions and their use in risk theory. *ASTIN Bulletin*, 35(1), 145–161. https://doi.org/10.2143/AST.35.1.583170

Bladt, M., & Neuts, M. F. (2003). Matrix-exponential distributions: Calculus and interpretations via flows. *Stochastic Models*, 19(1), 113–124. https://doi.org/10.1081/STM-120018141

Blighter Surveillance Systems, Chess Dynamics, & Enterprise Control Systems. (2017). *AUDS anti-UAV defence system brochure*. https://www.atri.lt/wp-content/uploads/2022/05/auds-brochure.pdf

Brust, M. R., Danoy, G., Stolfi, D. H., & Bouvry, P. (2021). Swarm-based counter UAV defense system. *Discover Internet of Things*, 1(1), 2. https://doi.org/10.1007/s43926-021-00002-x

Center for European Policy Analysis. (2024). *The phony war: Ukraine and Russia's decoy drones*. Retrieved January 7, 2026, from https://cepa.org/article/the-phony-war-ukraine-and-russias-decoy-drones/

DroneShield Ltd. (2024). *RfAI: Radio frequency artificial intelligence*. Retrieved March 25, 2026, from https://www.droneshield.com/capabilities#rfai

Espreso Global. (2024). *Weapons of mass deception: What are decoy drones and how Russia and its allies are trying to bypass air defense*. Retrieved January 7, 2026, from https://global.espreso.tv/russia-ukraine-war-weapons-of-mass-deception-what-are-decoy-drones-and-how-russia-and-its-allies-are-trying-to-bypass-air-defense

Fackrell, M. (2009). An alternative characterization for matrix exponential distributions. *Advances in Applied Probability*, 41(4), 1005–1022. https://doi.org/10.1239/aap/1261669583

Gaertner, U. (2013). *UAV swarm tactics: An agent-based simulation and markov process analysis* (Master's thesis). Naval Postgraduate School. Monterey, CA, USA.

González-Jorge, H., Aldao, E., Veiga-López, F., Fontenla-Carrera, G., Balvís, E., & Ríos-Otero, E. (2024). Counter drone technology: A review. *Preprints, 2024*. https://doi.org/10. 20944/preprints202402.0551.v1

He, Q. M., & Zhang, H. (2007). On matrix exponential distributions. *Advances in Applied Probability*, 39(1), 271–292. https://doi.org/10.1239/aap/1175266478

Horvath, A., Scarpa, M., & Telek, M. (2016). Phase type and matrix exponential distributions in stochastic modeling. In L. Fiondella & A. Puliafito (Eds.), *Principles of performance and reliability modeling and evaluation: Essays in honor of kishor trivedi on his 70th birthday* (pp. 3–25). Springer International Publishing. https://doi.org/10.1007/978-3-319-30599-8_1

Hotters, D., Kalliatanakis, N., Kuhnert, F., Prakasha, P. S., van Heur, R., & Kaya, B. (2025). A novel approach to quantify counter-drone system effectiveness against UAS swarms. *Proceedings of CEAS Aerospace Europe Conference (CEAS 2025).*

Jia, N., Yang, Z., & Yang, K. (2019). Operational effectiveness evaluation of the swarming UAVs combat system based on a system dynamics model. *IEEE Access*, *7*, 25209–25224. https://doi.org/10.1109/ACCESS.2019.2898728

Joint Air Power Competence Centre. (2021). *A comprehensive approach to countering unmanned aircraft systems*. Retrieved January 9, 2026, from https://www.japcc.org/wp-content/uploads/A-Comprehensive-Approach-to-Countering-Unmanned-Aircraft-Systems.pdf

Joint Research Centre. (2025). *Technical developments in counter-drone technology: C-UAS detection, tracking and identification technology* (EUR Scientific and Technical Research Report). Publications Office of the European Union. https://publications.jrc.ec.europa.eu/repository/bitstream/JRC140692/JRC140692_01.pdf

Kaja, H., Paropkari, R. A., Beard, C., & Van De Liefvoort, A. (2021). Survivability and disaster recovery modeling of cellular networks using matrix exponential distributions. *IEEE Transactions on Network and Service Management*, 18(3), 2812–2824. https://doi.org/10.1109/TNSM.2021.3090355

Khawaja, W., Ezuma, M., Semkin, V., Erden, F., Ozdemir, O., & Guvenc, I. (2025). A survey on detection, classification, and tracking of UAVs using radar and communications systems. *arXiv*. https://arxiv.org/abs/2402.05909

Kyiv Post. (2026). *Russia fires iskander missile, swarms Ukraine with 95 drones overnight*. Reuters. Retrieved January 7, 2026, from https://www.kyivpost.com/post/67565

Latouche, G., & Ramaswami, V. (1999). *Introduction to matrix analytic methods in stochastic modeling*. SIAM.

Li, N., Su, Z., Ling, H., Karatas, M., & Zheng, Y. (2023). Optimization of air defense system deployment against reconnaissance drone swarms. *Complex System Modeling and Simulation*, *3*(2), 102–117. https://doi.org/10.23919/CSMS.2023.0003

Liang, X., Qi, Y., Chen, G., & Meng, G. (2025). Research on a distributed strategy for UAV detection and tracking of specific targets. *Artificial Intelligence Review*, 58(12), 388. https://doi.org/10.1007/s10462-025-11364-x

Lockheed Martin. (2025). *Counter unmanned aerial systems (C-UAS)*. Lockheed Martin Corporation. Retrieved March 25, 2026, from https://www.lockheedmartin.com/en-us/capabilities/counter-unmanned-aerial-systems.html

Mutzari, D., Deb, T., Molinaro, C., Pugliese, A., Subrahmanian, V. S., & Kraus, S. (2025). Defending a city from multi-drone attacks: A sequential stackelberg security games approach. *Artificial Intelligence*, *349*, 104425. https://doi.org/10.1016/j.artint.2025.104425

Neuts, M. F. (1981). *Matrix-geometric solutions in stochastic models: An algorithmic approach*. Dover Publications.

Petsioti, P., Życzkowski, M., Brewczyski, K., Cichulski, K., Kaminski, K., Razvan, R., Mohamoud, A., Church, C., Koniaris, A., De Cubber, G., & Doroftei, D. (2024). Methodological approach for the development of standard C-UAS scenarios. *Open Research Europe*, *4*(240). https://doi.org/10.12688/openreseurope.18339.1

Reuters. (2025a). *Russia launches war's largest drone attack on Ukraine, Kyiv says*. Retrieved January 7, 2026, from https://www.reuters.com/world/europe/russia-launches-warslargest-drone-attack-ukraine-kyiv-says-2025-02-23/

Reuters. (2025b). *Russia pounds Kyiv, other regions in mass drone and missile attack*. Retrieved January 7, 2026, from https://www.reuters.com/business/aerospace-defense/russianlaunches-major-drone-missile-attack-ukraine-still-ongoing-2025-09-28/

Reuters. (2025c). *Three killed in biggest Ukrainian drone attack on Moscow region*. Retrieved January 7, 2026, from https://www.reuters.com/world/europe/ukraine-launches-droneattacks-targeting-moscow-russia-says-2025-03-11/

Reuters. (2026). *Ukraine targets Moscow daily with drones this year, Russia says, in apparent escalation*. Retrieved January 7, 2026, from https://www.reuters.com/business/aerospace-defense/ukraine-targets-moscow-daily-with-drones-this-year-russia-saysapparent-2026-01-04/

Sosa, C. (2025). *Swarm of UAS modeling and simulation methodology* (tech. rep.) [STO-MPSAS- 192]. NATO Science & Technology Organization, Systems Analysis & Studies Panel.

https://publications.sto.nato.int/publications/STO%20Meeting%20Proceedings/STO-MP-SAS-192/MP-SAS-192-31.pdf

Special Competitive Studies Project. (2023). *The drone advantage in modern warfare*. Retrieved January 7, 2026, from https://thesvi.org/the-drone-advantage-in-modern-warfare/

Spotter Global. (2024). *360° counter-UAS radar datasheet*. https://www.milestonesys.com/globalassets/marketplace/uploaded-assets/00120000013cnzqaai/360-counter-uas-radar-datasheet-rev0042.pdf

Tekinay, M., Beard, C., & van de Liefvoort, A. (2020). Partial packet duplication: Control of fade and non-fade duration outages using matrix exponential distributions. *IEEE Transactions on Vehicular Technology*, *69*(5), 5652–5656. https://doi.org/10.1109/TVT. 2020.2978848

Tu, X., Zhang, C., Zhuang, H., Liu, S., & Li, R. (2024). Fast drone detection with optimized feature capture and modeling algorithms. *IEEE Access*, *12*, 108374–108388. https://doi.org/10.1109/ACCESS.2024.3438991

van de Liefvoort, A., & Heindl, A. (2005). Approximating matrix-exponential distributions by global randomization. *Stochastic Models*, 21(2-3), 669–693. https://doi.org/10.1081/STM-200056231

Wang, C., Wu, A., Hou, Y., Liang, X., Xu, L., & Wang, X. (2023). Optimal deployment of swarm positions in cooperative interception of multiple UAV swarms. *Digital Communications and Networks*, *9*(2), 567–579. https://doi.org/10.1016/j.dcan.2022.04.002

Yan, J., LI, J., Li, J., Liu, C., & Guo, Y. (2025). Simulation and deduction method for UAVs swarm in air-to-ground attack scenarios. *Proceedings of the 2025 17th International Conference on Computer Modeling and Simulation*, 135–144. https://doi.org/10.1145/3761668.3761690

Zhu, Y., Yin, L., & Zhao, K. (2025). Millisecond-response tracking and gazing system for UAVs: A domestic solution based on ”phytium + cambricon”. *arXiv* https://arxiv.org/abs/2509.04043