

Digital Tides: A Fluid-Dynamic Framework for Flux-Aware Infrastructure Provisioning in UAV Logistics Networks

Wen-Yu Dong¹, Song Zhao¹, Rui-Si Han¹, Qi Bi, *Fellow, IEEE*, and Sheng Chen², *Life Fellow, IEEE*

Abstract—The emergence of high-frequency pulsating logistics uncrewed aerial vehicle (UAV) swarms gives rise to “Digital Tides”, i.e., complex traffic dynamics that challenge sustainable resource provisioning in mobile computing networks. Conventional infrastructure provisioning strategies, which typically rely on static snapshot-based analysis and localized density estimation, fail to capture the macroscopic advection of computational workloads. As a result, reactive resource activation suffers from inherent hysteresis, yielding nominal efficiency gains at the cost of mission-critical service loss at the advancing wavefront. To address this issue, we develop a fluid-based spatiotemporal framework by explicitly solving the continuity equation to characterize the macroscopic velocity field of the workload flow. Building on this framework, we propose a flux-aware asymmetric activation strategy that leverages the derived information flux vector as a kinematic precursor of demand propagation. Unlike symmetric thresholding, the proposed control logic decouples activation and deactivation dynamics. Theoretical analysis confirms the intrinsic spatial phase-lead of the flux signal and shows that the proposed strategy generates a proactive guard ring to compensate for service setup latency, including delays caused by mobile edge computing container cold-starts. We further derive closed-form expressions for instantaneous service availability and period-average energy efficiency. In addition, we formulate a quality-of-service-penalized metric to evaluate effective energy efficiency under strict outage constraints. Numerical results show that the proposed flux-driven strategy enables zero-latency tracking of the mobile wavefront and achieves a Pareto-optimal trade-off between service reliability and energy consumption, outperforming reactive baselines in dynamic logistics corridors.

Index Terms—Digital tides, UAV logistics, fluid-dynamic modeling, infrastructure provisioning, information flux.

I. INTRODUCTION

THE rapid proliferation of logistics uncrewed aerial vehicles (UAVs) for food delivery, medical transport, etc. is reshaping the low-altitude economy, establishing aerial mobile

computing as a cornerstone of future autonomous systems [1], [2], [3], [4]. Unlike terrestrial users with stochastic mobility, logistics UAVs exhibit highly structured macroscopic flow patterns driven by on-demand delivery cycles. This creates a phenomenon we term the low-altitude digital tide, a pulsating traffic pattern where swarms periodically disperse from distribution centers and return to base. This tidal traffic pattern involves two critical kinematic phases: the expansion phase where the spatial footprint of the swarm grows rapidly, and the contraction phase where the swarm recedes. Supporting this complex tidal traffic poses a significant energy challenge. Ground infrastructures, such as mobile edge computing (MEC)-enabled base stations (BSs), typically employ always-on service interfaces to guarantee availability. Given that infrastructure provisioning accounts for over 80% of network energy consumption [5], maintaining idle computing resources during the tide’s ebb periods fundamentally contradicts the sustainability goals of mobile computing systems.

Simply applying conventional energy-saving strategies to this scenario creates a critical reliability conflict. Existing methods, often rooted in static snapshot-based analysis [6], [7], [8], predominantly rely on detecting local workload density. While effective for stationary hotspots, this static perspective is inherently ill-suited for the application scenario considered. The core failure mechanism lies in the mismatch between the resource provisioning speed and the workload advection speed. Specifically, as a UAV swarm propagates outward, it forms a rapidly moving demand wavefront. However, due to the inevitable service setup latency specified by the necessary transition time from sleep to active mode, including MEC container cold-starts, a node triggered solely by local density will activate too late, often after the leading UAVs have already entered its service area. This hysteresis results in a wavefront outage, creating a moving service void that compromises the reliability of mission-critical logistics operations.

While data-driven prediction methods attempt to mitigate this lag, they often incur prohibitive computational overheads for training and lack the generalizability to handle emergent swarm behaviors. To overcome this latency-induced failure, resource orchestration must shift from observing static snapshots to predicting dynamic workload flows. This paper introduces a fluid-dynamic provisioning framework for sustainable mobile computing infrastructure. Instead of restricting the analysis to independent static realizations common in traditional stochastic

Received 31 December 2025; revised 26 March 2026; accepted 25 April 2026. Date of publication 29 April 2026; date of current version 2 September 2026. Recommended for acceptance by O. Onireti. (Corresponding author: Wen-Yu Dong.)

Wen-Yu Dong, Song Zhao, and Qi Bi are with Future Technology Research Center, China Telecom Research Institute, Beijing 102209, China (e-mail: dongwy@chinatelecom.cn; zhaosong1@chinatelecom.cn; qibi@chinatelecom.cn).

Rui-Si Han is with the Cloud Network Operating System R&D Center, China Telecom, Beijing 102209, China (e-mail: hanruisi@chinatelecom.cn).

Sheng Chen is with School of Electronics and Computer Science, University of Southampton, SO17 1BJ Southampton, U.K. (e-mail: sqc@ecs.soton.ac.uk).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TMC.2026.3688690>, provided by the authors.

Digital Object Identifier 10.1109/TMC.2026.3688690

modeling, we model the logistics swarm as a continuous compressible fluid governed by conservation laws [9], [10], [11]. This allows us to derive an information-flux vector field that quantifies the momentum of computational demand. Moreover, unlike the scalar density, which peaks only upon arrival, the flux vector naturally points in the direction of flow and reaches its threshold ahead of the density peak. By exploiting this intrinsic kinematic phase lead, our strategy effectively generates a time buffer, waking up infrastructure resources just-in-time to reliably serve the incoming wavefront.

A. Related Works

Research on energy efficiency (EE) in non-terrestrial networks (NTNs) primarily focuses on two optimization domains: aerial agents and terrestrial infrastructure. The majority of existing literature prioritizes UAV endurance, extensively exploring joint trajectory design and power control to minimize propulsion and operational energy consumption [12], [13]. Conversely, studies focusing on terrestrial networks typically exploit UAVs as aerial edge nodes to offload workloads, thereby enabling ground BSs (GBSs) to enter sleep modes [14], [15]. However, the reciprocal challenge—optimizing the energy consumption of ground infrastructure dedicated to serving high-density aerial logistics—has received limited attention. Although early works such as [16] investigated GBS sleeping for UAV command and control, they relied on static access techniques like coordinated multi-point and non-orthogonal multiple access. These approaches treat aerial workloads as stationary spatial distributions, failing to capture the macroscopic advection inherent in logistics tides. Consequently, they lack the predictive capability necessary to compensate for service setup latency.

Adapting dynamic resource-scaling strategies from terrestrial mobile networks offers a potential solution. These strategies are generally categorized into robust optimization and learning-based prediction. Robust optimization frameworks, such as chance-constrained programming [17], address workload uncertainty via statistical moments. While robust against random fluctuations in stationary terrestrial clients, these methods typically adopt a discrete time-slotted perspective, which lacks the temporal resolution required to track a rapidly propagating demand wavefront. This leads to inherent hysteresis-induced outages. Similarly, data-driven paradigms, including long short-term memory (LSTM) networks [18], [19], [20] and deep reinforcement learning (DRL) [21], depend on continuous workload observability, a condition unavailable to sleeping nodes. Furthermore, these models often lack the physical generalizability required to handle emergent swarm behaviors on unseen logistics routes [21].

The ineffectiveness of adapting the aforementioned terrestrial strategies to the logistics UAV swarm context is fundamentally rooted in their mathematical modeling of system workload. Standard stochastic geometry analysis [22], [23], [24], [25], widely adopted to derive EE limits in mobile networks, relies heavily on the snapshot assumption, treating network evolution as a sequence of independent static realizations. Although

advances in spatio-temporal stochastic geometry have begun to address temporal interference correlations [26], [27], [28], these frameworks predominantly focus on stationary users or random waypoint mobility. They inherently lack the capability to capture the structured macroscopic advection characteristic of logistics tides. Consequently, such static modeling frameworks fail to provide the kinematic foresight needed to preemptively overcome service setup latency.

To address the limitations of discrete agent modeling, fluid dynamic models were proposed to capture macroscopic mobility patterns [9], [10], [11]. Similarly, mean field games were employed to mathematically formalize the interactions of massive UAV agents as a continuum [29], [30], [31]. However, these works primarily focus on distributed trajectory optimization or power control, leaving the coupled dynamics between macroscopic workload flow and ground resource activation unexplored. Prior studies [9], [10] utilized the fluid principles for capacity estimation or trajectory synthesis. These contributions, however, remain largely descriptive rather than prescriptive. A critical gap persists in mapping continuous kinematic metrics, such as flow velocity and flux, to the discrete switching dynamics of resource activation. Mitigating service setup latency [32] in high-mobility corridors remains an open engineering challenge. This limitation leads to a deceptive reliability-efficiency trade-off, where reactive strategies achieve misleadingly high nominal EE by systematically shedding the energy-intensive demand wavefront. To bridge this gap, we propose a fluid-based spatiotemporal framework that leverages information flux as a kinematic precursor. Unlike reactive strategies limited by snapshot-based inputs, our approach proactively compensates for service setup latency by aligning network activation with the macroscopic momentum of the digital tide.

B. Our Contributions

Our main contributions are summarized as follows.

- We establish a fluid-dynamic traffic model governed by conservation laws to characterize the pulsating dynamics of logistics UAV swarms. By explicitly solving the continuity equation, we derive the closed-form macroscopic velocity field. This continuous formulation overcomes the static limitations of classical stochastic modeling, establishing a fluid-based spatiotemporal framework that captures the cumulative spatiotemporal workload volume for evaluating non-stationary infrastructure sustainability.
- We propose a flux-aware asymmetric activation strategy that bridges the gap between macroscopic flow dynamics and discrete resource states. Identifying the information flux vector as a kinematic precursor for demand momentum, the proposed control logic exploits the intrinsic spatial phase-lead at the wavefront. This mechanism generates a proactive guard ring that effectively compensates for service setup latency, thereby enabling zero-latency tracking of the expanding demand wavefront.
- We develop a tractable analytical framework to quantify system performance, deriving closed-form expressions for the instantaneous service availability and period-average

TABLE I
SUMMARY OF KEY MATHEMATICAL NOTATIONS

Symbol	Description
Φ_{BS}, λ_{BS}	PPP and density of GBSs
H_{UAV}	Operational altitude of logistics UAVs
P_{act}, P_{slp}	Active and sleep power consumption of the LAS module
ν, γ	Path loss exponent and SINR threshold
B, η_{SE}	Bandwidth and target spectral efficiency
$\lambda(r, t)$	Spatio-temporal workload density at distance r and time t
$N(t)$	Total workload demand
$\sigma(t), \dot{\sigma}(t)$	Spatial spread parameter and its expansion rate
N_0, σ_0	Baseline workload intensity and spatial spread
δ_N, δ_σ	Load modulation index and spatial expansion index
$v_r(r, t)$	Macroscopic radial expansion velocity
$\alpha(\mathbf{x}, t)$	Binary activation state of a GBS at location $\mathbf{x}$
$\Phi(\mathbf{x}, t)$	Information flux vector quantifying demand momentum
$R_{wf}(t)$	Logistic wavefront position
$\lambda_{act}, \lambda_{hold}$	Static activation and holding density thresholds
δ_{th}	Flux sensitivity threshold
τ_{boot}	Service setup latency
$R_{active}(t)$	Effective activation radius of the network
$\mathcal{G}_{flux}$	Spatial flux gain
$\mathcal{T}_{total}, E_{total}$	Total served workload volume and energy consumption
η_{EE}, η_{eff}	Raw energy efficiency and effective energy efficiency

EE. To resolve the deceptive reliability-efficiency trade-off suffered by existing strategies, we formulate a quality of service (QoS)-penalized effective EE metric. Guided by this metric, we demonstrate that the proposed strategy can be calibrated to identify the Pareto-optimal operating point, effectively decoupling the conflict between energy conservation and service reliability.

This paper is organized as follows. Section II introduces the fundamental concepts of digital tides and motivates the adoption of the fluid-dynamic paradigm. Section III establishes the system model, formulates the low-altitude digital tide using fluid dynamics and introduces the hardware constraints. Section IV details the flux-aware asymmetric activation strategy. Section V provides the theoretical performance analysis, deriving the resource activation geometry, service availability, and EE, alongside the proof of the flux phase-lead property. Section VI presents numerical results, validating the analytical models and quantifying the performance gains regarding outage elimination and Pareto optimality. Section VII discusses the practical limitations and the generalizability of the proposed framework under non-ideal conditions. Finally, Section VIII concludes the paper.

Throughout this paper, $\mathbb{P}(\cdot)$ and $\mathbb{E}[\cdot]$ denote the probability and expectation operators, respectively. $\mathbb{I}(\cdot)$ denotes the indicator function, which equals 1 if the condition inside is true, and 0 otherwise. $\mathcal{L}_I(s) = \mathbb{E}[e^{-sI}]$ is the Laplace transform of the interference I . The probability density function (PDF) of random variable X is denoted by $f_X(x)$. ${}_2F_1(a, b; c; z) = \sum_{n=0}^{\infty} \frac{(a)_n (b)_n}{(c)_n n!} z^n$ is the standard mathematical notation for the Gaussian hypergeometric function, and $[\cdot]^+ \triangleq \max\{0, \cdot\}$. For easy reference, Table I summarizes the key symbols used in this paper.

II. BACKGROUND AND MOTIVATION

To establish the physical intuition behind the proposed methodology, this section introduces the fundamental concepts

of digital tides and clarifies the motivation for adopting a fluid-dynamic perspective.

A. From Stochastic Mobility to Digital Tides

Traditional mobile networks are primarily designed to accommodate independent users with stochastic mobility patterns, such as random waypoint or Brownian motion. Under such microscopic randomness, the aggregate spatial distribution of users remains statistically stationary over short time scales.

However, the deployment of logistics UAV swarms introduces a distinct macroscopic dynamic. Governed by centralized scheduling and delivery deadlines, these swarms exhibit structured macroscopic flow patterns rather than isolated random walks. When a logistics hub dispatches a wave of delivery UAVs, it creates a pulsating traffic surge propagating outward, which we conceptualize as a digital tide. Serving this tide requires the ground infrastructure to activate and deactivate synchronously with the swarm propagation.

B. Limitations of Reactive Density Detection

Evaluating dynamic infrastructure activation requires an accurate characterization of the spatial workload. Conventional resource provisioning strategies, often modeled via static stochastic geometry, predominantly rely on localized density detection. This approach triggers BSs by independently observing the instantaneous spatial accumulation of users.

While effective for stationary hotspots, the static density metric lacks predictive capability regarding macroscopic mobility. Specifically, a BS relying on local density detection initiates its activation sequence only after the swarm enters its coverage area. Given the service setup latency associated with hardware wake-up and MEC container cold-starts, this reactive detection creates a temporal lag. Consequently, the leading edge of the swarm, namely, the mission-critical wavefront, experiences significant service outages.

C. The Fluid-Dynamic Paradigm and Information Flux

To overcome the hysteresis in reactive control, the infrastructure must anticipate the arrival of the demand. This requires a shift in mathematical abstraction from discrete static agents to a continuous dynamic flow. By modeling the dense UAV swarm as a compressible fluid governed by the continuity equation, we analytically define the macroscopic velocity field of the traffic.

Building upon this fluid-dynamic foundation, we introduce the concept of information flux. Defined as the product of the scalar spatial density and the macroscopic velocity vector, the flux encapsulates the momentum of the service demand. The physical motivation for utilizing this metric lies in its spatial phase-lead property. At the expanding wavefront, the macroscopic velocity increases the flux signal, allowing it to reach detection thresholds at greater distances compared to the static density. By exploiting this kinematic property, the network generates a proactive guard ring, reserving the required time buffer to execute MEC cold-starts before the physical workload arrives.




Fig. 1. System model of the MEC-enabled UAV logistics network.

III. SYSTEM MODEL AND PROBLEM FORMULATION

Fig. 1 illustrates an MEC-enabled urban logistics network driven by a pulsating logistics tide, which is conceptualized as a macroscopic spatiotemporal field of traffic density that undergoes periodic, radial expansion and contraction.

A. Network Architecture

1) *Logistics Hub and UAV Kinematics*: The system serves a centralized logistics area $\mathcal{A} \subset \mathbb{R}^2$, targeting scenarios such as food delivery or medical transport. A logistics hub is located at the origin of $\mathcal{A}$, serving as the source and sink of the UAV swarm. UAVs operate at a constant altitude H_{UAV} . Unlike stochastic mobility models, such as random way-point, the aggregate swarm kinematics exhibit a deterministic macroscopic structure: a radial velocity field driven by the periodic waves of delivery demand.

2) *Ground Infrastructure and Hardware Constraints*: The ground network comprises cellular BSs deployed according to a homogeneous Poisson point process (PPP) Φ_{BS} with density λ_{BS} . To balance ubiquitous coverage reliability with high-capacity throughput requirements, we adopt a dual-layer heterogeneous architecture consisting of: 1) the terrestrial sector layer, where existing down-tilted sectors remain always-active to provide basic command and control links via antenna sidelobes; and 2) the low-altitude sector (LAS) layer, where BSs are equipped with dedicated up-tilted modules utilizing orthogonal frequency resources for high-bandwidth computational task offloading.

In this framework, we adopt an MEC architecture where each GBS is equipped with an edge server to handle latency-sensitive logistics tasks, such as visual navigation and path planning. Given that the radio unit and the edge computing module share a power supply and control interface, their activation is synchronized. Accordingly, the service setup latency τ_{boot} is modeled as the aggregate of the radio frequency circuit wake-up time and the MEC container cold-start delay, which includes the time needed to instantiate virtual machines and load service images.

This work focuses on the energy-efficient activation control of the power-consuming LAS layer. Consistent with analytical models for infrastructure sleep depth [32], we incorporate the activation latency τ_{boot} to account for the non-instantaneous transition from sleep to active mode. This imposes a timing

constraint on the network control: to guarantee service availability at time t , the activation decision $\alpha(\mathbf{x}, t)$ must be asserted at or before $t - \tau_{\text{boot}}$. The specific design of the flux-aware control logic satisfying this constraint will be detailed in Section IV.

B. Fluid-Dynamic Workload Modeling

To capture the macroscopic mobility of the UAV swarm, we model the traffic as a compressible fluid driven by logistics demand.

1) *Density Field Formulation*: We model the instantaneous mobile agent density $\lambda(r, t)$ at distance r and time t as a time-varying, non-separable Gaussian field:

$$\lambda(r, t) = \frac{N(t)}{2\pi\sigma^2(t)} \exp\left(-\frac{r^2}{2\sigma^2(t)}\right). \quad (1)$$

To explicitly capture the periodicity, we model the load $N(t)$ and spread $\sigma(t)$ as sinusoidal functions with period T_{period} :

$$N(t) = N_0 (1 + \delta_N \cos(\omega t)), \quad (2)$$

$$\sigma(t) = \sigma_0 (1 + \delta_\sigma \cos(\omega t + \varphi)), \quad (3)$$

where N_0 and σ_0 are the baseline workload intensity and spatial spread, respectively, δ_N and δ_σ are the workload modulation index and spatial expansion index, while $\omega = 2\pi/T_{\text{period}}$ is the angular frequency, and φ accounts for the phase synchronization offset between traffic generation and spatial expansion.

2) *Macroscopic Velocity and Information Flux*: Governed by the law of conservation of mass, the evolution of the density field $\lambda(\mathbf{x}, t)$ at position $\mathbf{x} \in \mathbb{R}^2$ and time t adheres to the continuity equation:

$$\frac{\partial \lambda}{\partial t} + \nabla \cdot (\lambda \mathbf{v}) = 0, \quad (4)$$

where $\mathbf{v}(\mathbf{x}, t)$ is the macroscopic velocity field. By solving this equation (see Section S-I of the Supplementary File), we derive the radial expansion velocity of the logistics tide as $v_r(r, t) \approx r \frac{\dot{\sigma}(t)}{\sigma(t)}$, where $r = \|\mathbf{x}\|$ is the radial distance from the logistics hub.

Remark 1: The derivation in Section S-I reveals that the macroscopic flow is governed by two orthogonal forces: 1) source dynamics driven by load variation $\dot{N}$, representing a “fictitious velocity” required to accommodate node creation or removal; and 2) transport dynamics driven by spatial reshaping $\dot{\sigma}$. In high-mobility logistics scenarios, the kinematic transport dominates the source variation, i.e., $|\dot{\sigma}/\sigma| \gg |\dot{N}/N|$. This dominance ensures that the flow physically manifests as a uniform expansion rather than a localized fluctuation, thereby validating the use of the linear velocity model $v_r \propto r$ for proactive control.

Based on this kinematic description, we introduce the information flux to quantify the momentum of service demand.

Definition 1: The information flux, denoted by $\Phi(\mathbf{x}, t)$, is a vector field representing the rate of traffic demand flow passing through a unit spatial cross-section. Mathematically, it is defined as the product of the scalar spatio-temporal user density $\lambda(\mathbf{x}, t)$ and the macroscopic velocity vector $\mathbf{v}(\mathbf{x}, t)$:

$$\Phi(\mathbf{x}, t) \triangleq \lambda(\mathbf{x}, t) \mathbf{v}(\mathbf{x}, t) \approx \lambda(r, t) \left(r \frac{\dot{\sigma}(t)}{\sigma(t)} \right) \mathbf{e}_r, \quad (5)$$

where $\mathbf{e}_r = \mathbf{x}/\|\mathbf{x}\|$ is the unit radial vector pointing outwards from the origin.

This vector formulation offers a critical physical insight that cannot be obtained from scalar density alone. While $\lambda(\mathbf{x}, t)$ merely measures the static accumulation of load, $\Phi(\mathbf{x}, t)$ quantifies the dynamic trend of the demand momentum. The magnitude $\|\Phi\|$ is amplified by the expansion velocity v_r , which scales linearly with distance r . At the expanding wavefront where r is large, the flux intensity reaches its detection threshold before the workload density accumulates to the activation level. This kinematic property creates a spatial phase lead, serving as a predictive indicator for proactive activation.

3) *Logistic Wavefront*: We define the logistic wavefront, denoted as $R_{\text{wf}}(t)$, as the spatial boundary where the user density decays to the minimum service threshold λ_{th} :

$$R_{\text{wf}}(t) \triangleq \{r \in \mathbb{R}^+ : \lambda(r, t) = \lambda_{\text{th}}\}. \quad (6)$$

Physically, $R_{\text{wf}}(t)$ delineates the effective edge of the serving swarm. To quantify its propagation dynamics, we define the wavefront phase velocity as the time derivative of the wavefront position:

$$v_{\text{wf}}(t) \triangleq \frac{dR_{\text{wf}}(t)}{dt}. \quad (7)$$

During the expansion phase where $\dot{\sigma} > 0$, this wavefront propagates outward with a positive velocity $v_{\text{wf}}(t) > 0$. A critical implication is that any reactive strategy relying solely on the local density condition $\lambda \geq \lambda_{\text{th}}$ will track this moving boundary with a temporal lag. This necessitates a proactive mechanism to compensate for the activation latency.

C. Adiabatic Link Quality Model and SINR

To bridge the continuous macroscopic mobility and discrete microscopic transmission, we invoke the adiabatic assumption, also known as the time-scale separation principle. Analogous to the quasi-static block-fading model [33], this assumption holds because the task offloading time T_{pkt} is orders of magnitude shorter than the characteristic evolution scale of the logistics tide T_{hydro} . Physically, T_{hydro} characterizes the time required for significant macroscopic reshaping, defined as the inverse of the maximum relative expansion rate: $T_{\text{hydro}} \sim (\max |\dot{\sigma}/\sigma|)^{-1} \propto T_{\text{period}}$. Since the delivery cycle T_{period} operates on hourly scale while T_{pkt} is in milliseconds, $T_{\text{pkt}} \ll T_{\text{hydro}}$ is strictly satisfied. Consequently, the spatial intensity field $\lambda(\mathbf{x}, t)$ and the aggregate interference are treated as time-invariant snapshots during each transmission interval.

Furthermore, the set of potential interferers is determined by the instantaneous binary activation states $\alpha(\mathbf{X}_i, t) \in \{0, 1\}$:

$$\Phi_{\text{active}}(t) = \{\mathbf{X}_i \in \Phi_{\text{BS}} \mid \alpha(\mathbf{X}_i, t) = 1\}, \quad (8)$$

where $\mathbf{X}_i \in \mathbb{R}^2$ denotes the spatial location of the i -th GBS. In contrast to static homogeneous PPP networks, $\Phi_{\text{active}}(t)$ constitutes a time-varying, spatially inhomogeneous point process governed by the pulsating logistics tide.

For a typical UAV located at $\mathbf{x}_0$, the instantaneous signal-to-interference-plus-noise ratio (SINR) is given by:

$$\text{SINR}(\mathbf{x}_0, t) = \frac{P_{\text{LAS}} h_0 L(\|\mathbf{x}_0 - \mathbf{X}^*\|)}{I_{\text{agg}}(\mathbf{x}_0, t) + P_{\text{N}}}, \quad (9)$$

where $\mathbf{X}^* \in \Phi_{\text{active}}(t)$ is the serving BS that provides the maximum received power, and P_{N} is the noise power, while the aggregate interference $I_{\text{agg}}(\mathbf{x}_0, t)$ is the summation of signals from all other active BSs within the active region (AR) $\mathcal{A}_{\text{on}}(t)$:

$$I_{\text{agg}}(\mathbf{x}_0, t) = \sum_{\mathbf{X}_i \in \Phi_{\text{active}}(t) \setminus \{\mathbf{X}^*\}} P_{\text{LAS}} h_i L(\|\mathbf{x}_0 - \mathbf{X}_i\|). \quad (10)$$

In (9) and (10), h_0 and h_i are the respective small-scale fading channel power gains, which are exponentially distributed for Rayleigh fading, $L(d) = d^{-\nu}$ is the power-law path loss function with exponent $\nu > 2$, and P_{LAS} denotes the transmission power of the LAS module.

D. Spatiotemporal Performance Metrics

To quantify the trade-off between sustainability and reliability under dynamic load, we define three key metrics over a logistics cycle T_{period} . Let $\alpha = \{\alpha(\mathbf{x}, t) \in \{0, 1\} \mid \mathbf{x} \in \mathcal{A}, t \in [0, T_{\text{period}}]\}$ be the binary activation policy of the LAS layer.

1) *Total Served Workload Volume*: The primary utility of the logistics network is the computational tasks offloaded by the UAVs. Unlike static networks where the instantaneous sum-rate suffices, the non-stationary logistics tide requires evaluating the cumulative service volume. We define the total served traffic, denoted as $\mathcal{T}_{\text{total}}(\alpha)$, as the spatio-temporal integration of the successful service capacity density:

$$\mathcal{T}_{\text{total}}(\alpha) \triangleq \int_0^{T_{\text{period}}} \int_{\mathcal{A}} \mathcal{R}(\mathbf{x}, t; \alpha) d\mathbf{x} dt, \quad (11)$$

where $\mathcal{R}(\mathbf{x}, t; \alpha)$ represents the successful service capacity density, also known as the effective spatial throughput, formulated as $\mathcal{R}(\mathbf{x}, t; \alpha) = B \eta_{\text{SE}} \lambda(\mathbf{x}, t) P_{\text{cov}}(\mathbf{x}, t \mid \alpha)$. Here, $P_{\text{cov}}(\mathbf{x}, t \mid \alpha) = \mathbb{P}(\text{SINR}(\mathbf{x}, t) > \gamma)$ denotes the service availability which is mathematically quantified by the instantaneous coverage probability, B is the system bandwidth, and η_{SE} represents the target spectral efficiency.

In the context of mission-critical logistics, this link-level coverage probability serves as the fundamental metric quantifying the system-level service availability.

2) *Total Energy Consumption*: The network energy consumption is directly determined by the spatiotemporal activation pattern of the GBSs. The total energy consumption over one period is defined as:

$$E_{\text{total}}(\alpha) \triangleq \int_0^{T_{\text{period}}} \int_{\mathcal{A}} \lambda_{\text{BS}} P_{\text{GBS}}(\alpha(\mathbf{x}, t)) d\mathbf{x} dt, \quad (12)$$

where $P_{\text{GBS}}(\cdot)$ models the power consumption of a single LAS module:

$$P_{\text{GBS}}(\alpha) = P_{\text{act}} \cdot \alpha + P_{\text{slp}} \cdot (1 - \alpha). \quad (13)$$

Here, P_{act} and P_{slp} denote the power consumption in the active and sleep modes, respectively.

3) *Service Unavailability (Wavefront Outage)*: To quantify the reliability risk at the expanding edge, we define the instantaneous service unavailability, $\mathcal{O}(t)$, as the fraction of workload mass located in the “outage zone”—the spatial gap between the physical demand wavefront $R_{\text{wf}}(t)$ and the effective service boundary $R_{\text{active}}(t)$ induced by the control policy α .

Mathematically, assuming that the policy generates a contiguous coverage area of radius $R_{\text{active}}(t)$, $\mathcal{O}(t)$ is calculated as the normalized user mass within the gap $[R_{\text{active}}(t), R_{\text{wf}}(t)]$:

$$\mathcal{O}(t; \alpha) \triangleq \frac{\left[\int_{R_{\text{active}}(t)}^{R_{\text{wf}}(t)} \lambda(r, t) \cdot 2\pi r \, dr \right]^+}{\int_0^{R_{\text{wf}}(t)} \lambda(r, t) \cdot 2\pi r \, dr}. \quad (14)$$

The unavailability is zero if the service boundary exceeds the wavefront, i.e., $R_{\text{active}} \geq R_{\text{wf}}$. This metric specifically captures the percentage of the logistics swarm—most notably the strategic leading edge—that is physically present but disconnected due to infrastructure hysteresis.

E. QoS-Constrained EE Maximization Problem

1) *Objective Function and Constraints*: The primary objective of sustainable mobile computing networks is to maximize the network EE, $\eta_{\text{EE}}(\alpha)$, which is the ratio of the total served workload volume to the total energy consumption over the logistics cycle:

$$\eta_{\text{EE}}(\alpha) \triangleq \frac{\mathcal{T}_{\text{total}}(\alpha)}{E_{\text{total}}(\alpha)}. \quad (15)$$

Mathematically, we formulate the network control problem as a QoS-constrained EE maximization problem:

$$\mathcal{P}_0 : \max_{\alpha} \eta_{\text{EE}}(\alpha), \quad (16a)$$

$$\text{s.t. } \mathcal{O}(t; \alpha) \leq \epsilon_{\text{th}}, \quad \forall t \in \mathbb{T}_{\text{exp}}, \quad (16b)$$

$$\alpha(\mathbf{x}, t) \in \{0, 1\}, \quad \forall \mathbf{x}, t, \quad (16c)$$

where $\mathbb{T}_{\text{exp}}$ denotes the time interval corresponding to the expansion phase of the logistics tide, and ϵ_{th} is the maximum allowed service unavailability threshold, e.g., $\epsilon_{\text{th}} = 0.01$.

2) *Generalized Effective Energy Efficiency*: Problem $\mathcal{P}_0$ is a non-convex fractional programming problem. Furthermore, the constraint (16b) imposes an infinite number of instantaneous QoS requirements over the continuous domain $\mathbb{T}_{\text{exp}}$, rendering direct global optimization intractable.

To obtain a tractable formulation, we adopt a penalty-function method to transform the constrained optimization problem into an unconstrained maximization of a surrogate objective function, termed the generalized effective EE, denoted as η_{eff} .

Definition 2: We define the effective EE as the following QoS-penalized metric:

$$\eta_{\text{eff}} \triangleq \eta_{\text{EE}} \cdot \mathcal{U}_{\beta} \left(\sup_{t \in \mathbb{T}_{\text{exp}}} \mathcal{O}(t) \right). \quad (17)$$

Here, $\mathcal{U}_{\beta}(\cdot)$ serves as a soft barrier function, or a multiplicative penalty function, derived from the reliability constraint. We

adopt the polynomial form:

$$\mathcal{U}_{\beta}(x) = \begin{cases} (1-x)^{\beta}, & \text{if } 0 \leq x < 1, \\ 0, & \text{if } x \geq 1, \end{cases} \quad (18)$$

where $\beta \geq 1$ is the penalty factor.

Remark 2: The formulation of η_{eff} reconciles the conflicting goals of $\mathcal{P}_0$. Physically, the term $(1-x)$ represents the service availability. As the penalty factor β increases, the cost of violating the outage constraint (16b) becomes severe, forcing the optimal solution of the unconstrained proxy problem to asymptotically approach the feasible region of the original problem $\mathcal{P}_0$. This conversion allows us to evaluate the system performance using a single scalar metric that reflects the holistic operational value under QoS constraints.

IV. FLUX-AWARE ASYMMETRIC ACTIVATION STRATEGY

To effectively manage the high-frequency pulsating logistics tide, the network control logic must bridge the gap between macroscopic fluid dynamics and discrete hardware states. We propose a flux-aware asymmetric activation strategy. Unlike symmetric thresholding, this strategy decouples the activation and deactivation dynamics to accommodate the distinct kinematic requirements of the expansion and contraction phases.

A. Principles of Asymmetric Activation

The core idea of the proposed strategy is to align the network response with the distinct risk profiles inherent in the pulsating workload. Rather than applying a uniform control logic, we design tailored state-transition mechanisms specifically adapted to the unique vulnerabilities of the expansion and contraction phases.

1) *Expansion Phase* ($\dot{\sigma} > 0$): The primary risk is the wavefront outage caused by boot-up latency. To mitigate this, the strategy employs a proactive activation mechanism, leveraging the spatial phase-lead of the information flux to trigger GBSs before the workload density physically accumulates.

2) *Contraction Phase* ($\dot{\sigma} \leq 0$): The primary challenge is maintaining connectivity for agents located at the spatial periphery, i.e., the Gaussian tail, as the swarm recedes. To mitigate the risk of premature disconnection, the strategy employs a hysteresis retention mechanism. The flux trigger is deactivated because the inward-pointing macroscopic velocity vector no longer serves as a valid precursor for proactive activation. Instead, the control logic relies exclusively on a lowered density holding threshold, ensuring that the effective service boundary contracts at a slower rate than the physical decay of user density.

B. Triggering Logic Formulation

Let $\alpha(\mathbf{x}, t) \in \{0, 1\}$ denote the binary activation state of a GBS at location $\mathbf{x}$. We construct the state transition logic based on three distinct Boolean triggering conditions: the proactive flux trigger $\mathcal{C}_{\text{flux}}$, the reactive load trigger $\mathcal{C}_{\text{load}}$, and the retention

Algorithm 1: Flux-Aware Asymmetric Cycle Control.

Require: Time t ; Fluid state $\{\sigma, \dot{\sigma}, N\}$; Previous state α_{t-1} .

Ensure: New activation state α_t .

- 1: **Determine Phase:**
- 2: **if** $\dot{\sigma}(t) > 0$ **then** $Mode \leftarrow \text{EXPANSION}$
- 3: **else** $Mode \leftarrow \text{CONTRACTION}$
- 4: **Update GBS States:**
- 5: **for** each GBS i at location $\mathbf{x}_i$ **do**
- 6: Compute the local density $\lambda(\mathbf{x}_i, t)$ using (1).
- 7: **if** $Mode$ is **EXPANSION** **then**
- 8: $\triangleright$ Compute Flux only when needed for proactive trigger
- 9: $\|\Phi(\mathbf{x}_i)\| \leftarrow \lambda(\mathbf{x}_i) \cdot \|\mathbf{x}_i\| \cdot \frac{\dot{\sigma}}{\sigma}$
- 10: **Trigger:** $\mathcal{C}_{\text{flux}} \leftarrow (\|\Phi(\mathbf{x}_i)\| \geq \delta_{\text{th}})$
- 11: $\alpha(\mathbf{x}_i, t) \leftarrow \mathbb{I}(\lambda(\mathbf{x}_i) \geq \lambda_{\text{act}} \vee \mathcal{C}_{\text{flux}})$
- 12: **else** $\triangleright$ Hysteresis Retention Phase
- 13: $\alpha(\mathbf{x}_i, t) \leftarrow \mathbb{I}(\lambda(\mathbf{x}_i) \geq \lambda_{\text{hold}}) \cdot \alpha(\mathbf{x}_i, t-1)$
- 14: **return** α_t

trigger $\mathcal{C}_{\text{hold}}$, which are defined by:

$$\begin{cases} \mathcal{C}_{\text{flux}} : & \|\Phi(\mathbf{x}, t)\| \geq \delta_{\text{th}}, \\ \mathcal{C}_{\text{load}} : & \lambda(\|\mathbf{x}\|, t) \geq \lambda_{\text{act}}, \\ \mathcal{C}_{\text{hold}} : & \lambda(\|\mathbf{x}\|, t) \geq \lambda_{\text{hold}}, \end{cases} \quad (19)$$

where δ_{th} is the flux sensitivity threshold, λ_{act} is the activation density threshold, and we introduce a lower holding threshold $\lambda_{\text{hold}} < \lambda_{\text{act}}$, such as $\lambda_{\text{hold}} \rightarrow 0$, to implement the hysteresis mechanism. The activation protocol is defined as:

$$\alpha(\mathbf{x}, t) = \begin{cases} \mathbb{I}(\mathcal{C}_{\text{load}} \vee \mathcal{C}_{\text{flux}}), & \text{if } \dot{\sigma}(t) > 0, \\ \mathbb{I}(\mathcal{C}_{\text{hold}}) \cdot \alpha(\mathbf{x}, t-1), & \text{if } \dot{\sigma}(t) \leq 0. \end{cases} \quad (20)$$

To facilitate practical deployment in resource-constrained network controllers, we formalize the proposed logic in Algorithm 1.

C. Complexity and Implementation Analysis

The proposed strategy exhibits linear computational complexity $O(M)$, where M is the number of GBSs. Unlike combinatorial optimization or reinforcement learning approaches that typically scale with $O(M^2)$ or higher, our method relies on calculating analytical fluid metrics ($\lambda, \|\Phi\|$) locally for each node. Beyond computational efficiency, the framework ensures minimal signaling overhead. Unlike centralized optimization schemes requiring real-time channel state information (CSI) from all UAVs, our framework operates on macroscopic statistical parameters, such as $\sigma(t), N(t)$, that evolve slowly on the scale of minutes (the logistics cycle) rather than milliseconds. Consequently, the broadcast of these fluid parameters entails negligible bandwidth cost, and the localized flux calculation supports a fully distributed, event-driven activation logic that requires no backhaul coordination between neighboring GBSs. This low overhead is critical for real-time resource provisioning in large-scale mobile systems.

D. Spatiotemporal Mechanism Analysis

The proposed logic results in a dynamic network topology that physically adapts to the logistics tide.

1) *Traveling Wave Activation (Expansion)*: During the expansion phase, the condition $\|\Phi\| \geq \delta_{\text{th}}$ creates a proactive guard ring ahead of the density wavefront. This does not imply simultaneous global activation. Since the flux magnitude $\|\Phi(r, t)\|$ is spatially graded (peaking at the wavefront), the activation condition is satisfied sequentially from the center to the periphery. This results in a “traveling active zone” that propagates ahead of the UAV swarm, ensuring resources are provisioned selectively and just-in-time.

2) *Hysteresis Retention (Contraction)*: As the swarm recedes, the flux intensity diminishes. The protocol automatically switches to the retention mode governed by $\mathcal{C}_{\text{hold}}$. This ensures that the active region shrinks slower than the density decay, maintaining connectivity for tail users until the local load becomes negligible.

Remark 3: A potential concern with kinematic triggers is the vulnerability to false alarms caused by high-velocity but low-density transient outliers. However, under the proposed fluid-dynamic framework, the activation is governed by the information flux ($\Phi = \lambda \mathbf{v}$), which acts as a robust spatiotemporal momentum filter. Mathematically, the exponential spatial decay of the density field strictly dominates the linear growth of the macroscopic velocity, ensuring that the flux magnitude naturally vanishes in sparse far-field regions. Physically, constrained by the maximum aerodynamic speed of UAVs, a sparse outlier with negligible density cannot generate sufficient momentum to cross the macroscopic threshold δ_{th} . Consequently, such transient anomalies fail to trigger $\mathcal{C}_{\text{flux}}$, thereby preventing premature infrastructure activation and energy waste.

V. PERFORMANCE ANALYSIS

Given the spatial inhomogeneity of the pulsating tide, for tractability, we employ a local homogeneity approximation to derive the CP and EE metrics.

A. Spatio-Temporal Activation Geometry

The proposed asymmetric strategy yields a phase-dependent network topology. First, we formally derive the effective activation radius, denoted as $R_{\text{active}}(t)$.

Proposition 1: The effective activation radius $R_{\text{active}}(t)$ that adapts to the expansion-contraction cycle of the logistics tide is expressed as

$$R_{\text{active}}(t) = \begin{cases} \max\{R_{\text{act}}(t), R_{\text{flux}}(t)\}, & \text{if } \dot{\sigma}(t) > 0, \\ R_{\text{hold}}(t), & \text{if } \dot{\sigma}(t) \leq 0, \end{cases} \quad (21)$$

where the component radii are given by:

$$R_{\text{act}}(t) = \sigma(t) \sqrt{2 \left[\ln \left(\frac{N(t)}{2\pi\sigma^2(t)\lambda_{\text{act}}} \right) \right]^+}, \quad (22a)$$

$$R_{\text{hold}}(t) = \sigma(t) \sqrt{2 \left[\ln \left(\frac{N(t)}{2\pi\sigma^2(t)\lambda_{\text{hold}}} \right) \right]^+}, \quad (22b)$$

$$R_{\text{flux}}(t) = \sup \left\{ r \in \mathbb{R}^+ : \frac{r}{\sigma^3(t)} e^{-\frac{r^2}{2\sigma^2(t)}} = \mathcal{K}_{\text{th}}(t) \right\}, \quad (22c)$$

with $\mathcal{K}_{\text{th}}(t) = \frac{2\pi\delta_{\text{th}}}{N(t)|\dot{\sigma}(t)|}$ representing the normalized flux threshold.

Proof: The detailed derivation is provided in Section S-II of the Supplementary File. $\square$

This derived activation radius reveals distinct physical behaviors across the expansion-contraction cycle.

1) *Expansion (Proactive Guard Ring):* During the expansion phase, i.e., $\dot{\sigma} > 0$, the macroscopic velocity $v_r \propto r$ amplifies the flux intensity at the periphery. This kinematic gain pushes the flux boundary $R_{\text{flux}}(t)$ beyond the static density boundary $R_{\text{act}}(t)$, creating a proactive guard ring of width $\Delta R_{\text{lead}} = R_{\text{flux}} - R_{\text{act}}$ that spatially compensates for the boot-up latency.

2) *Contraction (Hysteresis Retention):* During the contraction phase, i.e., $\dot{\sigma} \leq 0$, the flux trigger is disabled. The activation boundary is governed strictly by the holding threshold λ_{hold} . Since $\lambda_{\text{hold}} < \lambda_{\text{act}}$, the condition $R_{\text{hold}}(t) > R_{\text{act}}(t)$ holds. This creates a retention zone of width $\Delta R_{\text{lag}} = R_{\text{hold}} - R_{\text{act}}$, preventing the premature disconnection of users at the spatial periphery as the swarm recedes.

B. Local Service Availability Analysis

The dynamic activation of LAS modules creates a spatially finite and inhomogeneous interference field. Unlike standard stochastic geometry approaches that assume an infinite field of interferers, we explicitly model the truncation of interference at the activation boundary $R_{\text{active}}(t)$ to derive an accurate CP.

Theorem 1: Consider the finite active region bounded by $R_{\text{active}}(t)$. The instantaneous service availability CP for a typical UAV within the active zone is expressed as:

$$P_{\text{cov}}(r, t) \approx \int_0^{R_{\text{active}}(t)} 2\pi\lambda_{\text{BS}} y \exp \left(-\pi\lambda_{\text{BS}} y^2 \left(1 + \mathcal{F}_{\text{finite}}(y, R_{\text{active}}) \right) \right) dy, \quad (23)$$

where $\mathcal{F}_{\text{finite}}(y, R_{\text{active}})$ is the finite interference factor explicitly capturing the boundary effect:

$$\mathcal{F}_{\text{finite}}(y, R_{\text{active}}) = {}_2F_1(1, \kappa; 1 + \kappa; -\gamma) - \left(\frac{R_{\text{active}}}{y} \right)^2 {}_2F_1 \left(1, \kappa; 1 + \kappa; -\gamma \left(\frac{y}{R_{\text{active}}} \right)^\nu \right), \quad (24)$$

with $\kappa = 2/\nu$ and γ being the SINR threshold.

Proof: The detailed derivation is provided in Section S-III of the Supplementary File. $\square$

Remark 4: The second term in (24) represents the boundary correction. As $R_{\text{active}} \rightarrow \infty$, this term vanishes, reducing to the classical infinite plane solution. However, near the network edge where $y \rightarrow R_{\text{active}}$, this term significantly reduces the estimated interference.

C. Network Energy Efficiency

By substituting the phase-dependent activation radius derived in Proposition 1 into the metric definitions, we obtain the closed-form expression for EE.

Theorem 2: The period-average EE of the network, defined as the ratio of total served traffic to total energy consumption over a cycle T_{period} , is given by:

$$\eta_{\text{EE}} = \frac{\mathcal{T}_{\text{total}}}{E_{\text{total}}} = \frac{B \eta_{\text{SE}} \int_0^{T_{\text{period}}} \mathcal{M}(t) \bar{P}_{\text{cov}}(t) dt}{\int_0^{T_{\text{period}}} P(t) dt}, \quad (25)$$

where $\bar{P}_{\text{cov}}(t)$ denotes the spatially averaged CP over the dynamic active region $R_{\text{active}}(t)$, approximated as:

$$\bar{P}_{\text{cov}}(t) \approx \frac{1}{1 + \mathcal{F}_{\text{finite}}(R_{\text{active}}(t))}, \quad (26)$$

with $\mathcal{F}_{\text{finite}}$ given in (24), while $\mathcal{M}(t)$ and $P(t)$ are the served workload mass and network power consumption defined below.

1) *Served User Mass $\mathcal{M}(t)$:* This term represents the aggregate workload demand covered by the active network, derived from the Gaussian integration:

$$\mathcal{M}(t) = N(t) \left(1 - \exp \left(-\frac{R_{\text{active}}^2(t)}{2\sigma^2(t)} \right) \right). \quad (27)$$

2) *Network Power Consumption $P(t)$:* This term captures the instantaneous power cost of the infrastructure:

$$P(t) = \pi R_{\text{active}}^2(t) \lambda_{\text{BS}} \Delta P + P_{\text{base}}, \quad (28)$$

where $\Delta P = P_{\text{act}} - P_{\text{slp}}$ represents the activation power penalty, and $P_{\text{base}} = |\mathcal{A}| \lambda_{\text{BS}} P_{\text{slp}}$ denotes the baseline sleep power of the entire service area $\mathcal{A}$.

Proof: The detailed derivation is provided in Section S-IV of the Supplementary File. $\square$

This analytical form facilitates the numerical optimization of the flux threshold δ_{th} . The objective is to balance the size of the proactive guard ring, which dictates coverage reliability, against the incremental energy cost, thereby locating the Pareto-optimal operating point.

D. Theoretical Analysis of Flux Phase-Lead

The effectiveness of the proposed strategy depends on whether the information flux Φ provides an earlier detection signal than the scalar density λ at the expanding wavefront. We formally prove this property by analyzing the spatial monotonicity of the flux-to-load ratio.

Lemma 1: During the expansion phase where $\dot{\sigma}(t) > 0$, the ratio of the flux magnitude to the workload density, denoted as $\mathcal{R}_{\Phi/\lambda}(r, t)$, is a strictly monotonically increasing function of the radial distance r .

Proof: Recall the definition of information flux from Definition 1: $\|\Phi(r, t)\| = \lambda(r, t) \cdot |v_r(r, t)|$. Substituting the derived macroscopic velocity $v_r(r, t) \approx r \frac{\dot{\sigma}(t)}{\sigma(t)}$, the ratio is given by:

$$\mathcal{R}_{\Phi/\lambda}(r, t) \triangleq \frac{\|\Phi(r, t)\|}{\lambda(r, t)} = |v_r(r, t)| = r \frac{\dot{\sigma}(t)}{\sigma(t)}. \quad (29)$$

For any fixed time instant t in the expansion phase where $\dot{\sigma} > 0$, the term $\frac{\dot{\sigma}}{\sigma}$ is a positive constant with respect to space. Taking the partial derivative with respect to r :

$$\frac{\partial}{\partial r} \mathcal{R}_{\Phi/\lambda}(r, t) = \frac{\dot{\sigma}(t)}{\sigma(t)} > 0. \quad (30)$$

Thus, the relative strength of the flux signal compared to the density signal increases linearly with distance. This completes the proof. $\square$

Remark 5: Lemma 1 provides the theoretical basis for the kinematic phase lead. At the swarm center where $r \rightarrow 0$, density dominates. However, at the far-field wavefront where $r \gg \sigma$, the high expansion velocity amplifies the flux signal. Consequently, for a decaying Gaussian tail, the flux threshold condition $\|\Phi\| \geq \delta_{th}$ is satisfied at a strictly larger radius than the density condition $\lambda \geq \lambda_{th}$ (assuming normalized thresholds). This ensures that the flux-triggered proactive guard ring is generated ahead of the traffic load.

1) *Quantification of Spatial Flux Gain:* Based on the monotonicity proven above, we quantify the spatial gain afforded by the flux-aware strategy.

Definition 3: During the expansion phase where $\dot{\sigma} > 0$, the flux gain $\mathcal{G}_{flux}(t)$ is defined as the spatial difference between the proactive flux boundary and the reactive density boundary:

$$\mathcal{G}_{flux}(t) \triangleq R_{flux}(t) - R_{act}(t). \quad (31)$$

Since the macroscopic velocity $v_r(r) \approx r\dot{\sigma}/\sigma$ increases linearly with distance, the flux intensity $\|\Phi\| = \lambda v_r$ decays slower than the density λ at the wavefront. Consequently, for a properly calibrated δ_{th} , the condition $\mathcal{G}_{flux}(t) > 0$ holds.

Remark 6: Physically, $\mathcal{G}_{flux}$ represents the width of the proactive guard ring. From a control theory perspective, this metric represents a space-time conversion: the strategy trades spatial redundancy in the form of an extra active ring of width $\mathcal{G}_{flux}$ for temporal margin. This spatial advance compensates for the inevitable time delay τ_{boot} inherent in hardware state transitions.

To ensure uninterrupted connectivity for the leading UAVs, the spatial gain must outpace the physical movement of the swarm during the wake-up process. This leads to the following fundamental condition for reliability.

Theorem 3: The wavefront outage is theoretically eliminated if and only if the flux gain satisfies:

$$\mathcal{G}_{flux}(t) \geq \int_t^{t+\tau_{boot}} v_{wf}(\tau) d\tau \approx v_{wf}(t) \cdot \tau_{boot}, \quad (32)$$

where $v_{wf}(t)$ is the wavefront phase velocity defined in (7). This inequality establishes the fundamental geometric constraint for the flux-aware activation boundary.

Remark 7: It is crucial to note that satisfying the kinematic zero-latency condition in Theorem 3 is a necessary but not sufficient condition for successful task offloading. The holistic system reliability is governed by a two-tier hierarchy. 1) *Hardware Availability:* The GBS must be active upon the UAV's arrival, which is addressed by the flux-aware guard ring $\mathcal{G}_{flux}$; and 2) *Link Reliability:* The instantaneous SINR must exceed the target threshold γ . In our framework, this second tier is rigorously captured by the finite network coverage probability

TABLE II
DEFAULT SIMULATION PARAMETERS

Parameter	Symbol	Value
GBS Density (LAS Layer)	λ_{BS}	5 BSs/km ²
UAV Flight Altitude	H_{UAV}	150 m
Bandwidth per LAS	B	20 MHz
Path Loss Exponent	ν	2.5
Target SINR Threshold	γ	-5 dB
Target Spectral Efficiency	η_{SE}	2 bps/Hz
Baseline Workload Intensity	N_0	25,000 UAVs
Baseline Spatial Spread	σ_0	4.5 km
Load Modulation Index	δ_N	0.8
Spatial Expansion Index	δ_σ	0.7
Active Power (LAS)	P_{act}	400 W
Sleep Power (LAS)	P_{slp}	50 W
Service Setup Latency	τ_{boot}	300 s (5 min)
Flux Sensitivity Threshold	δ_{th}	Variable
Static Activation Threshold	λ_{act}	50 agents/km ²
Holding Threshold	λ_{hold}	2 agents/km ²

P_{cov} derived in Theorem 1. Consequently, the overall served traffic $\mathcal{T}_{total}$ defined in (11) explicitly integrates both the kinematic availability and the SINR-based coverage probability, ensuring that offloading failures due to signal degradation are inherently penalized.

The condition established in Theorem 3 imposes a fundamental geometric constraint on the activation boundary. To translate this theoretical kinematic requirement into practical system design, we derive the upper bound for the flux sensitivity threshold.

Proposition 2: To ensure zero-latency coverage given a service setup constraint τ_{boot} , the flux sensitivity threshold δ_{th} must be calibrated to satisfy:

$$\delta_{th} \leq \inf_{t \in \mathbb{T}_{exp}} \|\Phi(R_{act}(t) + v_{wf}(t)\tau_{boot}, t)\|. \quad (33)$$

This inequality serves as a design criterion: it quantifies the minimum sensitivity required to detect the precursor momentum at the specific look-ahead distance $d_{lead} = v_{wf}\tau_{boot}$.

Proof: From Theorem 3, the zero-latency requirement is $R_{flux} \geq R_{act} + v_{wf}\tau_{boot}$. Let the right-hand side be the required target radius $R_{req} \triangleq R_{act} + v_{wf}\tau_{boot}$. Recall that R_{flux} is the outer root of $\|\Phi(r)\| = \delta_{th}$. In the wavefront region where $r > \sigma$, the flux magnitude $\|\Phi(r)\|$ is a monotonically decreasing function of r . Therefore, to ensure the trigger radius R_{flux} extends beyond R_{req} , the triggering threshold δ_{th} must be lower than the flux intensity at R_{req} :

$$\delta_{th} = \|\Phi(R_{flux})\| \leq \|\Phi(R_{req})\|. \quad (34)$$

Substituting the definition of Φ yields the bound in (33). This completes the proof. $\square$

VI. NUMERICAL RESULTS AND ANALYSIS

We validate the analytical framework using Monte Carlo simulations averaged over 50,000 iterations within a service area $\mathcal{A}$ of 20×20 km² governed by the pulsating workload model. The default system parameters are listed in Table II. To ensure the realism of the evaluation, these parameters are selected based on a combination of standardization reports and practical hardware specifications. Specifically, the communication parameters, including the UAV operational altitude (150 m) and path loss

exponent ($\nu = 2.5$), are aligned with the 3GPP technical report for low-altitude aerial vehicles (TR 36.777) [34]. The hardware profiles, such as the active/sleep power consumption and the service setup latency ($\tau_{\text{boot}} = 5$ min), reflect the typical specifications of MEC-enabled micro BSs. This service setup latency accounts for the aggregate delay of the hardware transitioning from a deep sleep mode [32] and the software overhead of MEC container cold-starts [35]. Furthermore, the macroscopic traffic variables (N_0 and σ_0) are configured as design assumptions to represent a city-scale logistics hub serving a metropolitan area.

To assess the system-level gains, we compare the proposed flux-aware asymmetric strategy against the following four baseline strategies.

- 1) *Static Always-On*: All infrastructure modules remain active continuously. This provides an upper theoretical bound on service reliability but incurs the maximum possible energy cost.
- 2) *Reactive Strategy*: BSs activate strictly upon detecting local user density $\lambda \geq \lambda_{\text{act}}$. This serves as the fundamental baseline for analyzing boot-up latency and resultant wavefront outages.
- 3) *Robust Margin Strategy (RMS)*: Adapted from robust optimization methods based on statistical moments and chance constraints [17], this baseline adds a conservative spatial safety margin outside the density boundary to hedge against mobility fluctuations. In our implementation, this margin is instantiated as a fixed 8.5 km worst-case guard ring.
- 4) *Traffic Snapshot Prediction (TSP) Baseline*: Inspired by learning-based traffic prediction models (e.g., LSTM [18], [19], [20] and DRL [21]), this serves as an oracle baseline. It assumes perfect foresight of the future static density boundary to fully cancel the temporal boot-up latency, activating BSs ahead of time based strictly on the predicted spatial density snapshots.

The analysis proceeds in four steps: validating the analytical coverage model and visualizing flow dynamics (Section VI-A), validating the spatial phase-lead of the information flux (Section VI-B), quantifying the reliability gain in eliminating outages (Section VI-C), and evaluating the overall network EE (Section VI-D).

A. Model Visualization and Analytical Validation

1) *Macroscopic Dynamics Visualization*: To verify the kinematic basis of the proposed framework, Fig. 2 illustrates the spatio-temporal evolution of the workload density field $\lambda(r, t)$ alongside the dynamic logistic wavefront $R_{\text{wf}}(t)$ represented by a white contour. By superimposing the Lagrangian characteristic trajectories of hypothetical fluid particles (plotted as thin grey lines), we observe distinct phase-dependent behaviors: during the expansion phase where $t < 1$ h, the grey trajectories diverge rapidly, corroborating that the wavefront propagation is physically driven by the macroscopic advection velocity $v_r > 0$; conversely, during the contraction phase where $t > 1$ h, the




Fig. 2. Spatiotemporal evolution of the pulsating logistics tide.




Fig. 3. Validation of the analytical local coverage probability.

grey trajectories converge, reflecting the inward momentum as the swarm returns to the hub. This confirms that the derived information flux metric provides an accurate kinematic representation of the swarm's momentum.

2) *Validation of Analytical Expressions*: We validate the finite network coverage model established in Theorem 1 against Monte Carlo simulations. Fig. 3 presents the spatial coverage profile at a representative snapshot of $R_{\text{active}} = 10$ km, corresponding to the fully expanded service boundary, with path loss exponent $\nu = 3.0$ and SINR threshold $\gamma = -3$ dB. Notably, the distinct boundary-induced coverage gain observed near the periphery ($r \approx 10$ km) is precisely predicted by the proposed analysis. This validation confirms that the finite interference factor derived in (24) correctly quantifies the reduction in aggregate interference due to the sleep mode of GBSs outside the active region, thereby eliminating the estimation error at the wavefront.

Fig. 4 extends this validation to the temporal domain over a full expansion-contraction cycle. The heatmap reveals a traveling service plateau, visualized as the green region, which is strictly confined within the analytical activation boundary traced by the white curve. The precise alignment between the analytical boundary and the effective service zone confirms the robustness




Fig. 4. Spatiotemporal evolution of the coverage probability over a full expansion-contraction cycle.




Fig. 5. Spatial mechanism of the flux-aware strategy at $t_0 = 0.6$ h, showing the proactive guard ring.

of the derived model under time-varying topologies. Even as the network size undergoes substantial fluctuations, the analytical framework accurately tracks the spatiotemporal evolution of the SINR field, demonstrating that the proposed fluid-based stochastic geometry offers a rigorous performance baseline for dynamic logistics corridors.

B. Validation of Kinematic Mechanism: Flux Phase-Lead

We first validate the kinematic basis of the proposed framework, specifically the spatial and temporal phase-lead of the information flux Φ . The simulation focuses on the rapid expansion phase, as the resulting macroscopic advection velocity is the physical driver of the flux signal.

1) *Spatial Mechanism. Proactive Guard Ring:* Fig. 5 elucidates the spatial interplay between the normalized agent density $\lambda(r, t_0)$ and the flux magnitude $\|\Phi(r, t_0)\|$ at a representative snapshot $t_0 = 0.6$ h. The results reveal a fundamental limitation of conventional detection methods. As shown by the blue solid curve, the agent density decays rapidly at the network periphery. This metric is strictly coupled to the physical presence of agents: the density threshold λ_{act} is crossed only upon the swarm's arrival. Consequently, any strategy relying solely on density fails




Fig. 6. Temporal validation of zero-latency triggering at the network edge ($r = 14.5$ km).

to anticipate the load, making the service setup latency τ_{boot} an unavoidable penalty.

In contrast, the flux magnitude represented by the red dashed curve exhibits a distinct wavefront enveloping effect. Driven by the high macroscopic velocity $v_r \propto r$ in the far field, the flux signal remains significant even in regions where the agent density is negligible, specifically at distances exceeding 10 km. This kinematic property effectively decouples the trigger signal from the physical load. As a result, the GBSs are activated by the flux condition $\|\Phi\| \geq \delta_{\text{th}}$ well before the density condition is met. This predictive mechanism generates a substantial proactive guard ring of width $\Delta R \approx 8.25$ km, buying the necessary time for the infrastructure to wake up before the agents actually arrive.

2) *Temporal Validation. Zero-Latency Triggering:* The translation of this spatial gain into an effective temporal lead is verified in Fig. 6, which traces the metric evolution at a representative far-field location located at a radial distance of 14.5 km. As the logistics tide propagates outward, the flux curve exhibits a sharp rise well in advance of the density curve. The calibrated trigger activates at time t_{flux} , which precedes the density arrival time t_{den} by approximately 5 minutes. This interval closely approximates the mandatory service setup latency τ_{boot} , confirming that the information flux functions as a robust precursor signal to ensure the GBS completes its activation sequence just as the physical load arrives.

C. Reliability Analysis: Elimination of Wavefront Outage

Having verified the kinematic phase-lead mechanism, we now quantify the system-level reliability gain during the critical expansion phase spanning the interval $t \in [0.2, 0.8]$ h. The primary objective is to demonstrate how the proposed strategy mitigates service outages caused by a service setup latency of $\tau_{\text{boot}} = 5$ min. We evaluate reliability using the service unavailability metric $\mathcal{O}(t)$ defined in (14).

The effective service boundary is physically constrained by the activation latency, formulated as $R_{\text{active}}(t) = R_{\text{trig}}(t - \tau_{\text{boot}})$. For the reactive strategy, the trigger follows the static




Fig. 7. Reliability analysis during the expansion phase. (a) Service-boundary trajectory versus the physical wavefront. (b) Instantaneous service unavailability.

density wavefront where $R_{trig} = R_{act}$. Conversely, for the flux-aware strategy, we employ the robust calibration derived in Proposition 2 by setting $\delta_{th} \approx 0.2 \cdot \delta_{max}$ to create a sufficient safety margin against the expansion velocity. Here, δ_{max} denotes the maximum flux magnitude over the expansion phase.

Fig. 7 visualizes the dynamic interaction between the traffic wavefront and the service boundary. The upper panel reveals a distinct tracking disparity where the reactive strategy, plotted as a blue dashed line, consistently lags behind the physical wavefront. This hysteresis creates a widening outage zone of approximately 1–2 km, confirming that static density detection is fundamentally too slow for high-mobility logistics corridors. In sharp contrast, the flux-aware strategy represented by the red dashed line maintains a service boundary that strictly envelopes the physical wavefront, effectively predicting the expansion trajectory.

The consequences are quantified in the lower panel. The reactive scheme suffers from a persistent service unavailability peaking at approximately 15%. This value represents the complete disconnection of the leading swarm, the most strategic assets executing long-range missions. The proposed strategy reduces this outage metric to near zero, validating the effectiveness of flux-based proactive control in eliminating wavefront outages.

D. System-Level Energy Efficiency and Pareto Optimization

Finally, we evaluate the comprehensive system performance to validate the economic viability of the proposed framework.

1) *Comparison of Raw Versus Effective Energy Efficiency:* Fig. 8 presents the normalized performance comparison across the five evaluated strategies. An initial observation is the counter intuitive behavior of the raw EE (η_{EE}). The reactive strategy and the TSP achieve the highest nominal η_{EE} . However, this metric is deceptive in high-mobility logistics scenarios. As indicated by the served traffic ratio, the reactive strategy captures only 79.6% of the computational demand, while TSP, despite perfectly




Fig. 8. System-level performance comparison across benchmark strategies.

predicting the density boundary, only marginally improves this to 82.8%. This implies that both strategies fundamentally fail to serve the critical wavefront of the digital tide—specifically, the high-velocity leading UAVs entering new coverage areas before the local density reaches the activation threshold. Consequently, their high raw energy efficiencies are merely artifacts of severe service denial, which is operationally unacceptable for mission-critical logistics.

To resolve this efficiency-reliability trade-off, we analyze the effective EE (η_{eff}), which explicitly penalizes service outages. Under this QoS-constrained metric, the performance of both the reactive and TSP strategies collapses. To guarantee service availability, the RMS enforces a static worst-case spatial guard ring, successfully elevating the served traffic to approximately 99.0%. However, this rigid geometric expansion incurs massive energy waste through severe over-activation, causing its η_{eff} to degrade to a level comparable with the static always-on upper bound. In contrast, the proposed flux-aware strategy leverages the fluid-dynamic velocity field to dynamically adjust the activation boundary. It seamlessly captures the wavefront momentum, achieving a 99.1% traffic service rate that is numerically comparable to the always-on benchmark, while promptly shrinking the boundary during swarm contraction to reduce trailing energy waste. By inherently balancing kinematic reliability and dynamic power consumption, the flux-aware framework achieves the maximum effective EE among all the evaluated approaches.

2) *Optimal Threshold Calibration:* To validate the existence of a unique optimal operating point, Fig. 9 analyzes the sensitivity of system performance to the flux threshold δ_{th} . The effective EE and service reliability curves depicted in Fig. 9 elucidate the fundamental trade-off between energy consumption and service reliability. Specifically, the optimization landscape reveals how flux sensitivity dictates the balance between energy waste and outage risks.

a) *Over-Sensitive Regime ($\delta_{th} < 10^{-1}$):* Here, the low threshold results in high sensitivity. Both the effective EE and service reliability improve as δ_{th} increases. As δ_{th} approaches 10^{-1} , the service reliability reaches approximately 100% but this comes at a severe energy cost with the effective EE suppressed below 0.7. This is because the system maintains an excessively wide proactive guard ring, activating GBSs far earlier than necessary.

b) *Transition Regime ($10^{-1} \leq \delta_{th} < 10^2$):* This interval marks the efficiency-gaining phase. As δ_{th} increases, the system progressively sheds spatial redundancy, and the effective EE climbs




Fig. 9. Impact of the flux sensitivity threshold δ_{th} on system performance and the resulting Pareto-optimal operating point.

steadily from 0.7 to its peak, while reliability remains largely robust above 95%. This confirms that the flux trigger still provides sufficient lead time to compensate for τ_{boot} while eliminating the waste of an oversized guard ring.

c) Pareto-Optimal Point ($\delta_{th} \approx 10^2$): The global maximum of the effective EE curve defines the optimal equilibrium. Crucially, this peak efficiency does not correspond to perfect reliability: at the optimal δ_{th} , service reliability drops from its peak of 100% to approximately 93%.

d) Under-Sensitive Regime ($\delta_{th} > 10^2$): Beyond the optimal point, δ_{th} becomes excessively high. Physically, the proactive guard ring contracts until its width is insufficient to compensate for the service setup latency τ_{boot} . In the asymptotic limit, the flux trigger is effectively disabled, causing the strategy to degenerate into the reactive baseline. Consequently, the network detects the wavefront too late, causing service reliability to plummet below 75%. The heavy wavefront outage triggers the steep slope of the QoS penalty function $\mathcal{U}_\beta$, driving the effective EE to collapse toward the negligible performance floor of the reactive strategy.

This result demonstrates a strategic trade-off: to maximize the effective EE, the system must tolerate a marginal degradation in reliability, corresponding to an approximately 7% outage risk at the wavefront. This calibration tightens the proactive guard ring to the minimum size necessary to offset hardware latency, thereby locating the true Pareto-optimal frontier that is unattainable by strictly reactive mechanisms.

E. System Sensitivity and Robustness Analysis

1) Sensitivity to Service Setup Latency: We evaluate the system robustness against service setup latency $\tau_{boot} \in [1, 10]$ min, representing varying MEC instantiation and RF wake-up overheads. We fix the flux sensitivity threshold to the Pareto-optimal $\delta_{th} = 10^2$. As shown in Fig. 10(a), the reactive baseline's service unavailability scales linearly with τ_{boot} , approaching 37% at a 10-minute latency. Conversely, the flux-aware strategy exhibits graceful degradation, restricting the outage to 6.2% at 5 minutes and below 15% at 10 minutes. This confirms that




Fig. 10. Sensitivity of system performance to the service setup latency under the Pareto-optimal flux threshold: (a) impact on service reliability, and (b) impact on effective energy efficiency.




Fig. 11. Normalized effective EE versus the flux sensitivity threshold under varying SINR thresholds and spectral efficiency requirements.

the proactive guard ring effectively absorbs the temporal lag. Consequently, as depicted in Fig. 10(b), the reactive strategy's effective EE plummets due to severe QoS penalties, whereas the proposed framework maintains superior energy efficiency across the latency spectrum, validating its stability for practical deployments.

2) Robustness Against RF Constraints: We further investigate the system sensitivity to underlying RF parameters under five scenarios using the control variable method, as illustrated in Fig. 11. First, fixing $\eta_{SE} = 2.0$ bps/Hz and increasing γ from -5 dB to 5 dB nonlinearly degrades the coverage probability, compressing the effective EE curves downward. Second, fixing $\gamma = -5$ dB and increasing η_{SE} from 2.0 to 4.0 bps/Hz linearly multiplies the spatial capacity, stretching the effective EE curves upward.

Despite these severe vertical fluctuations, the Pareto-optimal flux sensitivity threshold remains stationary at $\delta_{th} \approx 10^2$ across all scenarios. This invariant optimal point suggests that the required proactive guard ring is primarily governed by the macroscopic expansion velocity and hardware setup latency, while being largely insensitive to the tested microscopic RF link conditions. Consequently, network operators may be able to upgrade modulation and coding schemes without materially recalibrating the kinematic infrastructure triggering logic, although the absolute performance levels may still change.

VII. DISCUSSION AND PRACTICAL LIMITATIONS

While the proposed fluid-dynamic modeling and flux-aware activation framework establish a tractable methodology for macroscopic infrastructure provisioning, certain idealized mathematical abstractions warrant further discussion regarding their practical limitations under non-ideal conditions.

1) *Impact of Spatial Deployment Deviations:* The theoretical framework evaluates instantaneous coverage probability and energy efficiency under the assumption of a homogeneous PPP, denoted by Φ_{BS} . In complex urban environments, base station deployments are constrained by topography and zoning laws. Stochastic geometry literature establishes that the complete spatial randomness of a PPP typically yields a conservative lower bound on coverage probability compared to planned repulsive networks. Consequently, if the actual cellular deployment exhibits a more regular grid-like pattern, the variance of the nearest-serving distance decreases, mitigating severe local interference. In such scenarios, the effective service availability improves beyond our analytical predictions. Conversely, extreme clustering of infrastructure may introduce localized coverage holes. However, the core mechanism proposed in this paper, namely, the kinematic phase-lead of the macroscopic information flux over the scalar density, is expected to remain qualitatively robust. This predictive capability is governed by the continuity equation of the fluid dynamics and is decoupled from the microscopic spatial distribution of the ground infrastructure.

2) *Fluid Approximation Boundaries and Abrupt Reconfiguration:* Treating the UAV swarm as a continuous compressible fluid is effective under macroscopic high-density conditions with smooth workload flows. However, the validity of this continuous fluid model relies on the cohesive movement of the logistics fleet. Under non-ideal conditions, such as abrupt swarm reconfiguration or chaotic individual trajectories, the macroscopic continuity equation loses its predictive accuracy. In these regimes, the macroscopic flux vector fails to resolve microscopic discontinuities, causing the proactive system performance to degrade toward the reactive baseline. Addressing this limitation necessitates a hybrid methodology combining macroscopic fluid dynamics with discrete multi-agent tracking in future research.

3) *Impact of Estimation Noise:* The current analysis assumes the accurate acquisition of macroscopic state variables. In practical deployments, sensory measurements of the swarm density and velocity fields are subject to estimation noise. Such inaccuracies in the flux estimation can induce localized false alarms or delayed infrastructure activations. Mitigating this vulnerability requires the integration of temporal smoothing techniques or predictive filters, such as Kalman filtering, prior to the flux evaluation to ensure robust activation triggers.

4) *Heterogeneous Service Startup Delays:* The theoretical and simulation analyses evaluate system robustness under uniform service setup latencies. Real-world infrastructure often exhibits heterogeneous startup times due to varying hardware capabilities and backhaul states across different cells. Managing such spatial heterogeneity requires extending the current framework by transforming the global scalar flux sensitivity threshold, denoted by δ_{th} , into a spatially varying field. This

adaptation allows the system to dynamically adjust the proactive guard ring on a per-cell basis according to localized hardware constraints.

5) *Applicability to Geometry-Constrained Corridors:* The current analytical evaluation assumes a radially symmetric Gaussian tide. In practice, UAV logistics often follow geometry-constrained flight corridors due to urban topology or airway regulations. Nevertheless, the proposed flux-induced phase-lead mechanism is expected to remain valid under such asymmetric or one-dimensional constrained trajectories. This resilience is fundamentally rooted in the continuity equation. For a 1D delivery corridor, the spatial distribution reduces to a longitudinal profile $\lambda(x, t)$. Driven by the conservation of mass during the expansion phase, the macroscopic advection velocity v_x still scales proportionally with the propagation distance x . Consequently, the information flux magnitude $\|\Phi\| = \lambda \cdot v_x$ inherently decays slower than the scalar density λ at the wavefront. Thus, the spatial monotonicity as proven in Lemma 1 mathematically holds, ensuring that the flux trigger continues to generate a proactive temporal margin to compensate for service setup latency in corridor-based deployments.

VIII. CONCLUSION

This paper has investigated the critical conflict between EE and service reliability in urban UAV logistics networks. We have shown that conventional reactive sleeping strategies yield misleadingly high nominal efficiency, a metric inadvertently resulting from the systematic shedding of the energy-intensive traffic wavefront. To resolve this misalignment, we have transcended the static snapshot limitations of classical stochastic geometry by establishing a fluid-dynamic traffic model and proposing a flux-aware asymmetric activation strategy. By exploiting the intrinsic spatial phase-lead of the information flux, our control logic generates a proactive guard ring that effectively compensates for service setup latency. Numerical results have validated that this framework achieves a zero-latency response, reducing wavefront outages from approximately 20% to near zero compared to reactive baselines. Under our derived QoS-penalized metric, the proposed strategy achieves Pareto optimality, demonstrating that integrating macroscopic kinematic states into network control loops offers a resilient pathway for aerial mobile computing systems.

REFERENCES

- [1] M. M. Azari et al., "Evolution of non-terrestrial networks from 5G to 6G: A survey," *IEEE Commun. Surv. Tut.*, vol. 24, no. 4, pp. 2633–2672, Fourthquarter 2022.
- [2] G. Geraci et al., "What will the future of UAV cellular communications be? A flight from 5G to 6G," *IEEE Commun. Surv. Tut.*, vol. 24, no. 3, pp. 1304–1335, thirdquarter 2022.
- [3] W.-Y. Dong, S. Yang, and S. Chen, "Uplink performance analysis of heterogeneous non-terrestrial networks in harsh environments: A novel stochastic geometry model," *IEEE Trans. Commun.*, vol. 73, no. 8, pp. 6734–6747, Aug. 2025.
- [4] W.-Y. Dong, S. Yang, P. Zhang, and S. Chen, "Stochastic geometry based modeling and analysis of uplink cooperative satellite-aerial-terrestrial networks for nomadic communications with weak satellite coverage," *IEEE J. Sel. Areas Commun.*, vol. 42, no. 12, pp. 3428–3444, Dec. 2024.

- [5] G. Wu, C. Yang, S. Li, and G. Y. Li, "Recent advances in energy-efficient networks and their application in 5G systems," *IEEE Wireless Commun.*, vol. 22, no. 2, pp. 145–151, Apr. 2015.
- [6] W.-Y. Dong, S. Yang, P. Zhang, and S. Chen, "Modeling and performance analysis of IoT-over-LEO satellite systems under realistic operational constraints: A stochastic geometry approach," *IEEE Internet Things J.*, vol. 12, no. 15, pp. 30576–30593, Aug. 2025.
- [7] W.-Y. Dong, S. Yang, W. Lin, W. Zhao, J.-X. Gui, and S. Chen, "Outage probability analysis of uplink heterogeneous non-terrestrial networks: A novel stochastic geometry model," in *Proc. GLOBECOM 2024-2024 IEEE Glob. Commun. Conf.*, 2024, pp. 2588–2593.
- [8] W.-Y. Dong et al., "Stochastic geometry based performance analysis of terrestrial-to-aerial networks for nomadic communications," in *Proc. GLOBECOM 2024 - 2024 IEEE Glob. Commun. Conf.*, 2024, pp. 2731–2736.
- [9] C. Sergiou and V. Vassiliou, "Estimating maximum traffic volume in wireless sensor networks using fluid dynamics principles," *IEEE Commun. Lett.*, vol. 17, no. 2, pp. 257–260, Feb. 2013.
- [10] R. J. Skehill and S. McGrath, "The application of fluid mobility modelling in wireless cellular networks," *Mobile Inf. Syst.*, vol. 3, no. 2, pp. 89–106, Jan. 2007.
- [11] C.-F. Chiasserini et al., "Fluid models for large-scale wireless sensor networks," *Perform. Eval.*, vol. 64, nos. 7–8, pp. 715–736, Aug. 2007.
- [12] M. Zhou, H. Chen, L. Shu, and Y. Liu, "UAV-assisted sleep scheduling algorithm for energy-efficient data collection in agricultural Internet of Things," *IEEE Internet Things J.*, vol. 9, no. 13, pp. 11043–11056, Jul. 2022.
- [13] L. Zhang, A. Celik, S. Dang, and B. Shihada, "Energy-efficient trajectory optimization for UAV-assisted IoT networks," *IEEE Trans. Mobile Comput.*, vol. 21, no. 12, pp. 4323–4337, Dec. 2022.
- [14] H. Li et al., "UAV assisted BS sleep strategy for green communication," *IEEE Trans. Netw. Sci. Eng.*, vol. 12, no. 5, pp. 3770–3783, Sep./Oct. 2025.
- [15] Z. Liu, J. Qi, Y. Shen, K. Ma, and X. Guan, "Maximizing energy efficiency in UAV-assisted NOMA-MEC networks," *IEEE Internet Things J.*, vol. 10, no. 24, pp. 22208–22222, Dec. 2023.
- [16] A. Chowdary, Y. Ramamoorthi, A. Kumar, and L. R. Cenkeramaddi, "Joint resource allocation and UAV scheduling with ground radio station sleeping," *IEEE Access*, vol. 9, pp. 124505–124518, 2021.
- [17] W. Teng, M. Sheng, X. Chu, K. Guo, J. Wen, and Z. Qiu, "Joint optimization of base station activation and user association in ultra dense networks under traffic uncertainty," *IEEE Trans. Commun.*, vol. 69, no. 9, pp. 6079–6092, Sep. 2021.
- [18] X. Wang, B. Lyu, C. Guo, J. Xu, and M. Zukerman, "A base station sleeping strategy in heterogeneous cellular networks based on user traffic prediction," *IEEE Trans. Green Commun. Netw.*, vol. 8, no. 1, pp. 134–149, Mar. 2024.
- [19] G. Jang, N. Kim, T. Ha, C. Lee, and S. Cho, "Base station switching and sleep mode optimization with LSTM-based user prediction," *IEEE Access*, vol. 8, pp. 222711–222723, 2020.
- [20] A. Azari, F. Salehi, P. Papapetrou, and C. Cavdar, "Energy and resource efficiency by user traffic prediction and classification in cellular networks," *IEEE Trans. Green Commun. Netw.*, vol. 6, no. 2, pp. 1082–1095, Jun. 2022.
- [21] Q. Wu, X. Chen, Z. Zhou, L. Chen, and J. Zhang, "Deep reinforcement learning with spatio-temporal traffic forecasting for data-driven base station sleep control," *IEEE/ACM Trans. Netw.*, vol. 29, no. 2, pp. 935–948, Apr. 2021.
- [22] F. Baccelli and B. Błaszczyszyn, *Stochastic Geometry and Wireless Networks, Part I: Theory*. New York, NY, USA: Now Publishers Inc., 2009.
- [23] J. Peng, H. Tang, P. Hong, and K. Xue, "Stochastic geometry analysis of energy efficiency in heterogeneous network with sleep control," *IEEE Wireless Commun. Lett.*, vol. 2, no. 6, pp. 615–618, Dec. 2013.
- [24] A. Shabbir et al., "Optimizing energy efficiency in heterogeneous networks: An integrated stochastic geometry approach with novel sleep mode strategies and QoS framework," *PLoS One*, vol. 19, no. 2, Feb. 2024, Art. no. e0296392.
- [25] G. Zhao, S. Chen, and L. Hanzo, "Energy-spectral-efficient heterogeneous cellular networks: Joint optimization of cross-tier inter-BS cooperation and BS deployment," *IEEE Trans. Veh. Technol.*, vol. 73, no. 4, pp. 5659–5673, Apr. 2024.
- [26] M. Z. Win, P. C. Pinto, and L. A. Shepp, "A mathematical theory of network interference and its applications," *Proc. IEEE*, vol. 97, no. 2, pp. 205–230, Feb. 2009.
- [27] M. Haenggi and R. Smarandache, "Diversity polynomials for the analysis of temporal correlations in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 11, pp. 5940–5951, Nov. 2013.
- [28] X. Lu, M. Salehi, M. Haenggi, E. Hossain, and H. Jiang, "Stochastic geometry analysis of spatial-temporal performance in wireless networks: A tutorial," *IEEE Commun. Surv. Tut.*, vol. 23, no. 4, pp. 2753–2801, Fourthquarter 2021.
- [29] H. Shiri, J. Park, and M. Bennis, "Communication-efficient massive UAV online path control: Federated learning meets mean-field game theory," *IEEE Trans. Commun.*, vol. 68, no. 11, pp. 6840–6857, Nov. 2020.
- [30] D. Chen, Q. Qi, Z. Zhuang, J. Wang, J. Liao, and Z. Han, "Mean field deep reinforcement learning for fair and efficient UAV control," *IEEE Internet Things J.*, vol. 8, no. 2, pp. 813–828, Jan. 2021.
- [31] Y. Wang, C. Yang, T. Li, X. Mi, L. Li, and Z. Han, "A survey on mean-field game for dynamic management and control in space-air-ground network," *IEEE Commun. Surv. Tut.*, vol. 26, no. 4, pp. 2798–2835, Fourthquarter 2024.
- [32] O. Onireti, A. Mohamed, H. Pervaiz, and M. Imran, "Analytical approach to base station sleep mode power consumption and sleep depth," in *Proc. IEEE 28th Annu. Int. Symp. Pers., Indoor, Mobile Radio Commun.*, Montreal, QC, Canada, 2017, pp. 1–7.
- [33] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [34] 3rd Generation Partnership Project (3GPP), "Enhanced LTE support for aerial vehicles," 3GPP TR 36.777, V15.0.0, Release 15, Jan. 2018.
- [35] L. Pan, L. Wang, S. Chen, and F. Liu, "Retention-aware container caching for serverless edge computing," in *Proc. IEEE INFOCOM 2022 - IEEE Conf. Comput. Commun.*, London, U.K., 2022, pp. 1069–1078.



Wen-Yu Dong received the BS degree in electronic and information engineering from Sichuan University, Chengdu, China, in 2019, and the PhD degree in information and communication engineering from the School of Information and Communication Engineering, Beijing University of Posts and Telecommunications, Beijing, China, in 2025. He is currently a researcher with Future Technology Research Center, China Telecom Research Institute, Beijing. His research interests include space-air-ground integrated networks, fluid-spatiotemporal stochastic geometry, and wireless communications.



Song Zhao received the PhD degree in telecommunications and information systems from the Beijing University of Posts and Telecommunications, Beijing, China. He is currently a senior engineer with Future Technology Center, China Telecom Research Institute, Beijing. He is also a delegate with China Telecom in the 3GPP SA2 and SA5 working groups, and rapporteur for multiple 3GPP work items related to network automation and intelligence. His research interests include wireless channel modeling, radio resource management, proximity services, and network AI.



Rui-Si Han received the BS degree in e-commerce and law and the MEng degree in information and communication engineering from the Beijing University of Posts and Telecommunications, Beijing, China, in 2021 and 2024, respectively. She is currently a researcher with China Telecom Cloud Network Operating System R&D Center, Beijing. Her research interests include mobile ad hoc networks and research project management.



Qi Bi (Fellow, IEEE) received the MS degree from Shanghai Jiao Tong University and the PhD degree from Pennsylvania State University. He is currently the chief scientist with China Telecom and CTO with its Research Institute, specializing in 5G and 6G. He was awarded Bell Labs fellow in 2002, Bell Labs President's Gold Awards in 2000 and 2002, Asian American Engineer of the Year in 2005, Beijing Outstanding Contribution Award for Innovation and Entrepreneurship for Overseas Scholars in 2019, and three Patent Silver Prizes from the China National Intellectual Property Administration from 2022 to 2024.



Sheng Chen (Life Fellow, IEEE) received the BEng degree in control engineering from East China Petroleum Institute, Dongying, China, in 1982, the PhD degree in control engineering from City University, London, U.K., in 1986, and the higher doctoral (DSc) degree from the University of Southampton, Southampton, U.K., in 2005. From 1986 to 1999, he held research and academic appointments with the University of Sheffield, University of Edinburgh, and University of Portsmouth. In 1999, he joined the School of Electronics and Computer Science, University of Southampton. He is currently an emeritus professor with the University of Southampton. Dr. Chen has authored or coauthored more than 750 research papers. He has 23,000+ Web of Science citations with h-index 65, and 43,000+ Google Scholar citations with h-index 87. Professor Chen has established more than 40-year distinguished academic career in the fields of adaptive signal processing, wireless communications, neural network, and machine learning. He was the recipient of IEEE ComSoc Signal Processing and Computing for Communications 2025 Technical Recognition Award. Professor Chen was elected to a fellow of the United Kingdom Royal Academy of Engineering in 2014. He is a fellow of Asia-Pacific Artificial Intelligence Association, a fellow of IET, and one of the original 200 ISI Highly Cited Researchers in engineering (March 2004).