

On Energy Efficiency for Multi-Flow Semantic Routing in Multi-UAV Relay Networks[★]

Tung Son Do^a, Thanh Phung Truong^a, Kiet Nguyen Tuan Tran^a, Huy Dang Mac^a, Jaemin Kim^a, Tri-Hai Nguyen^b, Nhu-Ngoc Dao^c and Sungrae Cho^{a,*}

^a*School of Computer Science and Engineering, Chung-Ang University, Seoul, 06974, Republic of Korea*

^b*Faculty of Information Technology, Van Lang School of Technology, Van Lang University, Ho Chi Minh City, Vietnam*

^c*Department of Computer Science and Engineering, Sejong University, Seoul, Republic of Korea*

ARTICLE INFO

Keywords:

Ant colony optimization
energy efficiency
multi-flow routing
semantic communication
UAV networks

ABSTRACT

Unmanned aerial vehicles (UAVs) relay concurrent semantic flows in rapid-response networks, where multi-hop fidelity degrades multiplicatively whereas energy accumulates additively. Meanwhile, medium access control (MAC)-layer contention couples relay-sharing flows, making per-flow optimization inconsistent with this coupling and the joint problem computationally prohibitive. We formulate the Multi-Flow Semantic Routing Energy Efficiency (MF-SREE) problem, NP-hard at unit service availability and scored over communication-layer rather than total mission energy. Accordingly, we propose the multi-flow Semantic Ant Colony System (S-ACS MF) that leverages load-aware sequential construction to jointly optimize fidelity, inter-flow coupling, and energy efficiency. Under a load-count contention surrogate, a measured codec-calibrated fidelity curve, and non-binding capacity, S-ACS MF leads on MF-SREE in 11 of 12 settings. The sole exception is a rounded linear programming router. Paired margins are 0.024%–0.648% over Sequential Dijkstra, 1.04%–4.40% over traditional ant colony optimization (ACO), and 8.84%–46.43% over Greedy Walk, which Holm–Bonferroni correction separates from zero at 5, 9 and 12 settings, respectively. However, under a worst-case full-reuse spatial evaluator, S-ACS MF leads on re-scored MF-SREE in none of eight settings. A scheduling-aware variant, scored under the schedule it optimizes, ranks first in six of twelve schedule-evaluated settings. At the dense 50-node, 49-flow setting it is third of four in both capacity regimes, trailing the ACO and greedy baselines by Holm-corrected margins under binding capacity. Matched on the reduced ant budget with compilation disabled for both methods, S-ACS MF is 1.44 to 1.66 times slower than Traditional ACO MF; the earlier wall-clock advantage reflected configuration, not algorithm design.

1. Introduction

Unmanned aerial vehicles (UAVs) have been increasingly deployed as relay infrastructure for time-critical missions, including disaster response, surveillance, and remote sensing. In multi-UAV flying ad-hoc networks (FANETs), UAV relay chains need to deliver semantically rich payloads under stringent energy and latency constraints over bandwidth-limited aerial links. Semantic communication [43] transmits task-relevant meaning rather than raw bits, which suits this setting well. However, multi-hop relay chains introduce a fundamental challenge. A single weak hop can dominate the whole path and sharply reduce the semantic quality at the destination [44, 46].

This sensitivity becomes especially problematic when K source–destination pairs share the relay network simultaneously. At a shared relay, medium access control (MAC)-layer contention grows with the number of flows that the relay forwards. This contention lowers the effective link quality of every flow traversing it. Each admitted flow therefore changes the relay-load state faced by the flows already routed. Consequently, a route that is efficient in isolation can degrade the

fidelity and the energy cost of its neighbors once it is committed. Notably, the model adopted in this paper captures this effect through the relay load count alone. Transmitter geometry and the resulting spatial signal-to-interference-plus-noise ratio (SINR) lie outside its scope. These interactions introduce three intertwined difficulties. First, the shared relay-load coupling prevents independent per-flow optimization. Second, the Multi-Flow Semantic Routing Energy Efficiency (MF-SREE) maximization is challenging to optimize even in the relaxed constrained case. Third, the multiplicative fidelity structure of decode-and-forward semantic relaying breaks the additive path-cost assumptions in classical multi-commodity flow (MCF) formulations [3].

Prior work has addressed these three aspects largely in isolation. Single-flow semantic relaying [44, 18, 26] reports per-stage fidelity loss along an isolated chain without deriving a closed-form composition law. Furthermore, it does not model shared-relay load coupling across flows. Classical MCF [3] handles flow coupling on shared edges yet assumes additive path costs that are incompatible with multiplicative semantic fidelity. Deep reinforcement learning (DRL)-based multi-flow routing [29] adapts to dynamic load conditions but omits semantic fidelity in its reward design and offers no convergence guarantee. Hence, to the best of our knowledge, no existing framework jointly optimizes multi-hop semantic fidelity, MAC-layer flow coupling, and energy efficiency.

[★] Code availability: The source code, simulation engine, and experimental configurations for S-ACS MF will be made publicly available at <https://github.com/tung491/mf-sree> upon acceptance.

*Corresponding author.

✉ srcho@cau.ac.kr (S. Cho)

To address these limitations, we propose the multi-flow Semantic Ant Colony System (S-ACS MF), which leverages multi-flow-aware sequential construction to jointly optimize multi-hop semantic fidelity, multi-flow routing, and MAC-layer energy efficiency. Each design choice is forced by a structural feature of MF-SREE. Sequential per-flow construction on the evolving relay-load state accounts for inter-flow load-count contention coupling through the load-dependent heuristic. A marginal coupled-objective selection step then prices each candidate by its effect on the coupled total. Consequently, route construction and the subsequent coordinated reroute local search optimize the same quantity. A log-domain ratio surrogate addresses the non-additivity of multiplicative fidelity by acting as a candidate-generation prior. Candidate paths are subsequently re-scored under the true MF-SREE objective. Because the load-count surrogate abstracts away transmitter geometry and channel access, we additionally construct a scheduling-aware variant, S-ACS MF-SINR. Its routing objective is evaluated under the same explicit time-division multiple access (TDMA) schedule that scores it. In contrast, flow admission and candidate generation remain surrogate-based. The main contributions of this paper are as follows.

- To handle multi-flow routing under shared-relay contention coupling, we formulate the MF-SREE problem for concurrent flows under the load-count contention surrogate and per-node service-availability constraints, scored over communication-layer energy alone. Its decision version is proved NP-hard at unit service availability by reduction from directed vertex-disjoint paths [16], and per-edge monotonicity is derived.
- To tackle the coupled multi-flow optimization, we propose S-ACS MF, a metaheuristic framework that constructs routes via a load-aware energy-per-fidelity rule, selects candidates by marginal coupled-objective evaluation, and refines assignments through coordinated reroute local search. To align the routing objective with channel access, we further develop the scheduling-aware variant S-ACS MF-SINR under an explicit TDMA schedule with pairwise SINR feasibility.
- To establish rigorous performance bounds, we certify solution quality from two independent directions. We compute the per-flow zero-load MF-SREE ceiling exactly with a Pareto label-setting dynamic program that returns an optimality certificate on every solve, and obtain certified global optima by exhaustive enumeration on small-scale instances.
- Under measured codec calibration, S-ACS MF ranks first on MF-SREE in 11 of 12 grid settings and 20 of 24 measurably binding quality-of-service levels. Its paired margins span 0.024% to 0.648% over Sequential Dijkstra, 1.04% to 4.40% over Traditional ACO MF, and 8.84% to 46.43%

over Greedy Walk MF, which Holm–Bonferroni correction separates from zero at 5, 9 and 12 settings, respectively. However, S-ACS MF-SINR ranks third of four at the densest scheduled setting, and S-ACS MF leads in none of eight worst-case full-reuse spatial cells. Matched on the reduced ant budget with compilation disabled for both methods, S-ACS MF is 1.44 to 1.66 times slower than Traditional ACO MF, not faster.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 introduces the network setting and formulates the MF-SREE problem under the load-count contention surrogate. Section 4 details the proposed framework, the certified per-flow zero-load ceiling, and the fixed-load linear programming relaxation. Section 5 presents the numerical results, and Section 6 concludes the paper, with every formal proof given in Appendix A.

2. Related Work

2.1. Semantic Multi-Hop Communication and Routing

To position our contribution, we review prior efforts along two complementary axes, namely semantic multi-hop communication and multi-flow routing under wireless interference. Gündüz *et al.* [19] consolidated the semantic communication paradigm, in which transmission conveys the meaning of data instead of raw bits. Xie *et al.* [43] realized that paradigm for text through a Transformer-based semantic transceiver. When extending such paradigms to multi-hop networks, end-to-end semantic quality degrades progressively along the relay chain. Zhang *et al.* [46] demonstrated empirical semantic distortion accumulation along multi-hop decode-and-forward relays and proposed recursive retraining. However, analytical modeling of multi-hop fidelity cascades under shared-relay load coupling remains unaddressed. In this work, we model this degradation via a multiplicative fidelity cascade $\prod_{h=1}^{H_k} S(\gamma_h^{\text{eff}})$, where a single bottleneck link severely degrades end-to-end task performance. To alleviate this loss in the single-flow regime, Yin *et al.* [44] investigated cooperative semantic relaying. Similarly, Gao *et al.* [18] considered semantic-aware routing for low Earth orbit (LEO) satellite networks. Bian *et al.* [7] proposed a hybrid joint source-channel coding (JSCC) forwarding scheme tailored to mobile scenarios. However, each of these efforts is restricted to a single source–destination pair and does not account for the coupling introduced by simultaneous flows.

Guo *et al.* [20] proposed an integer linear programming formulation for semantic-aware routing in LEO satellite networks with knowledge-base matching. Their formulation targets single-flow optimization without incorporating multi-flow routing or contention coupling. Tang *et al.* [40] defined the Semantic Energy Efficiency (SEE) metric for single-relay cooperative links. Our Semantic Routing Energy Efficiency (SREE) generalizes this notion to multi-hop

routing. Under SREE, fidelity accumulates multiplicatively along the selected path, yielding an end-to-end routing efficiency ratio rather than a per-link indicator. More recently, Huang *et al.* [26] investigated a multi-hop semantic communication system that prioritizes transmission according to semantic importance. However, this line of work does not address multi-flow contention coupling when several semantic streams share spectrum.

In a parallel direction, Wu *et al.* [42] applied sleeping multi-armed bandit (MAB) policies to semantic path selection in space-ground networks. Meanwhile, Hu *et al.* [25] addressed multi-modal semantic resource allocation in UAV collaborative networks. Neither work formulates multi-flow routing under contention coupling, leaving the joint treatment of semantic fidelity and contention-coupled routing open.

2.2. Multi-Flow and Interference-Aware Routing

From the physical-layer routing perspective, Alam *et al.* [4] studied interference-aware UAV swarm routing by jointly optimizing trajectory, frequency allocation, and physical-layer metrics. Their formulation focuses exclusively on raw bit delivery without accounting for task-oriented semantic fidelity or multi-flow relay contention. Similarly, classical ACO [13] and its FANET variants remain agnostic to semantic awareness. Liu *et al.* [33] proposed a harmonic-function-based potential field (HF-PF) for interference-aware multi-hop routing in UAV networks, and He *et al.* [22] introduced DGMR for dynamic aerial ad-hoc networks. Both efforts target physical-layer routing without semantic fidelity metrics, rendering direct experimental comparison infeasible.

At the network-game level, Aggarwal *et al.* [2] applied mean-field game (MFG) theory to competitive channel scheduling among semantic transmitters. Their formulation addresses medium-access-control (MAC)-layer resource allocation rather than multi-hop routing, omitting the interplay between path selection and interference. On the multi-flow routing side, integer multi-flow routing is known to be NP-hard [15]. Kim *et al.* [29] proposed a DRL-based multi-flow routing scheme; semantic fidelity metrics, however, are omitted from their design.

In a related direction, Erfaniantaghvayi *et al.* [14] combined ACO with backpressure routing for multi-flow wireless traffic, optimizing physical-layer throughput without semantic-fidelity considerations. Our S-ACS MF framework instead optimizes the multiplicative SREE metric with cross-flow load tracking. Li *et al.* [31] introduced FRRA-ACO for capacitated multi-flow backup routing in wired data-center networks. In comparison, S-ACS MF extends this paradigm to wireless UAV networks with explicit semantic-fidelity awareness and load-dependent heuristics. Paul *et al.* [35] proposed DIAMOND, a graph-neural-network-aided reinforcement-learning approach for multi-flow transmission. DIAMOND targets physical-layer throughput, whereas S-ACS MF optimizes multiplicative SREE under unsplittable single-path routing constraints.

Unsplittable routing per flow inherits additional hardness beyond the fractional case [15].

Across both threads, none of the works reviewed above jointly SREE-optimizes multi-hop semantic fidelity, unsplittable multi-flow routing, and load-count contention coupling in UAV networks. Motivated by these limitations, we propose a unified framework that jointly optimizes multi-hop semantic fidelity, unsplittable multi-flow routing, and load-count contention coupling under a single MF-SREE objective. Solution quality is evaluated against a certified per-flow zero-load ceiling. The contention representation adopted here is deliberately a load-count contention surrogate. It charges a relay for the number of flows it forwards. As a result, this surrogate represents neither transmitter-geometry-based spatial interference nor medium-access scheduling. Section 3 delimits this scope, and every comparison on the surrogate evaluation track inherits it. The scheduled and full-reuse evaluations are identified separately where they are reported.

2.3. Learning-Based and Graph-Based Routing

Recently, deep reinforcement learning (DRL) and graph neural networks (GNNs) have been widely adopted for intelligent routing. DRL-based schemes jointly optimize multi-flow transmission [29] and dynamic aerial topologies [22], while GNN-assisted DRL models enhance topological generalization across unseen networks [41, 5]. The learning-based routers reviewed above, however, neither optimize fidelity-normalized energy metrics nor evaluate solution quality against certified optimality bounds.

In cloud computing, load-aware energy management has been explored for resource allocation [27] and dynamic voltage scaling [28]. While sharing the principle of load-dependent energy modeling, these cloud-centric methods omit wireless channel dynamics, multi-hop relay coupling, and semantic fidelity metrics.

Prior works have reported successful DRL and GNN applications in routing and traffic engineering [29, 22, 41, 4, 5, 6]. This paper does not attempt to tune such architectures for the non-additive, capacity-coupled MF-SREE objective. It instead reports untuned diagnostic implementations of MAPPO [45], QMIX [37], and a GNN policy. Under the state-action representation, hyperparameters and training budget adopted here, these implementations did not outperform a parameter-free greedy walk. That outcome is a property of these particular configurations. It is not evidence about what the corresponding method classes can achieve once tuned. Accordingly, no such conclusion is drawn in this paper. One candidate contributing factor, which the present experiments do not isolate, is that QMIX's monotonic value factorization cannot represent inter-agent action dependencies of the kind that shared relay capacity constraints induce [36].

The evaluation protocol enforces family-wise error control using Holm's procedure [24] and follows the statistical power analysis of Colas *et al.* [10]. These safeguards apply to the diagnostic comparisons described above. They do not

extend the resulting conclusions to learning architectures, state encodings or training budgets outside the tested set.

2.4. Certified Bounds for Non-Additive Path Objectives

Establishing exact performance bounds for fractional and non-additive path objectives is challenging. In particular, non-additivity invalidates standard shortest-path algorithms [17] and renders standard Dinkelbach-type parameterizations insufficient [12]. To address this, multicriteria label-setting methods [34] provide exact dynamic programming solutions. Building upon these principles, we instantiate a Pareto label-setting dynamic program to compute the exact per-flow zero-load ceiling. To the best of our knowledge, this is the first study to evaluate coupled semantic routing objectives against a certified per-flow zero-load upper bound.

For a concise overview, Table 1 contrasts representative prior works with the proposed approach.

The preceding subsections have positioned the proposed method against prior work on semantic relaying, interference-aware multi-flow routing, learned routing, and exact path bounding. Next, Section 3 formalizes the network, channel and fidelity models on which the SREE objective is defined.

3. System Model

In this section, we present the system model of a homogeneous multi-UAV network employed for semantic message relaying. We model the network as a time-varying directed graph $\mathcal{G}(t) = (\mathcal{V}, \mathcal{E}(t))$ indexed by discrete time slots $t \in \{0, 1, \dots\}$, where $\mathcal{V} = \{1, \dots, N\}$ denotes the set of N UAVs with $N \in \mathbb{Z}_+$. A directed edge $(i, j) \in \mathcal{E}(t)$ is established whenever the inter-UAV distance satisfies $\|\mathbf{p}_i(t) - \mathbf{p}_j(t)\| \leq d_{\max}$. Here, $\mathbf{p}_i(t) \in \mathbb{R}^3$ (m) is the 3D position of UAV i , and $d_{\max}$ (m) denotes the maximum air-to-air communication range.

Furthermore, the UAV positions follow a Gauss–Markov mobility model [32], expressed by the velocity-position recursion:

$$\mathbf{v}_i(t+1) = \alpha_m \mathbf{v}_i(t) + (1 - \alpha_m) \bar{\mathbf{v}} + \sqrt{1 - \alpha_m^2} \mathbf{w}_i(t), \quad (1a)$$

$$\mathbf{p}_i(t+1) = \mathbf{p}_i(t) + \mathbf{v}_i(t) \Delta t, \quad (1b)$$

where $\mathbf{v}_i(t), \bar{\mathbf{v}} \in \mathbb{R}^3$ (m/s) are the instantaneous and the mean UAV velocity vectors (mean speed $\bar{v} \triangleq \|\bar{\mathbf{v}}\|$), respectively, and $\alpha_m \in [0, 1]$ is the dimensionless memory coefficient. The driving noise $\mathbf{w}_i(t) \in \mathbb{R}^3$ is independent and identically distributed (i.i.d.) across nodes and time slots, drawn from the zero-mean Gaussian $\mathcal{N}(\mathbf{0}, \sigma_v^2 \mathbf{I})$. Here, $\mathbf{0} \in \mathbb{R}^3$ and $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ are the zero vector and the identity matrix, and σ_v (m/s) is the per-component velocity-perturbation standard deviation. Finally, Δt (s) is the slot duration, i.e., the position update interval. Notably, this recursion captures

Table 1

Comparison with related works.

Approach	Multi-hop	Sem-antic	MCF	Load coup.	Spat. SINR	MAC sch.	En-ergy	Bound
Coop. relay [44]	✓	✓	×	×	×	×	×	×
SC-aware LEO [18]	✓	✓	×	×	×	×	×	×
Multi-flow DRL [29]	✓	×	✓	×	✓	×	×	×
Ant-BP [14]	✓	×	✓	×	×	×	×	×
HF-PF [33]	✓	×	×	×	✓	×	×	×
DGMR [22]	✓	×	✓	×	×	×	×	×
FRRA-ACO [31]	✓	×	✓	×	×	×	×	×
DIAMOND [35]	✓	×	✓	×	✓	×	×	×
DRL+GNN [5]	✓	×	✓	×	×	×	×	×
MARL+GNN TE [6]	✓	×	✓	×	×	×	×	×
MAB path sel. [42]	✓	✓	×	×	×	×	×	×
UAV res. alloc. [25]	✓	✓	×	×	×	×	✓	×
Semantic MFG [2]	×	✓	×	×	×	✓	×	×
MFG routing [30]	✓	×	×	×	×	×	Partial	×
S-ACS MF	✓	✓	✓	✓	×	×	✓	✓
S-ACS MF-SINR	✓	✓	✓	✓	✓	✓	✓	×

Note: MCF = multi-flow routing ($K \geq 2$); Load coup. = per-relay forwarding-load (flow-count) contention represented inside the optimizer; Spat. SINR = co-channel interference derived from transmitter geometry and channel gains; MAC sch. = explicit slot or airtime allocation; Bound = scoring against certified zero-load ceiling. The three contention columns denote distinct modeling choices rather than successive levels of a single capability. Both proposed rows are variants of one framework: S-ACS MF prices contention by relay load alone, whereas S-ACS MF-SINR additionally routes under an explicit TDMA schedule with pairwise SINR feasibility.

the short-term correlation typical of rotary-wing trajectories at low altitude.

We consider a set of K source-destination pairs $\{(s_k, d_k)\}_{k=1}^K$ with $K \in \mathbb{Z}_+$ and $s_k, d_k \in \mathcal{V}$ that simultaneously transmit semantic messages through the shared relay network. Each flow k requires an unsplittable path $\mathcal{P}_k = (n_0^{(k)}, n_1^{(k)}, \dots, n_{H_k}^{(k)})$ with $n_h^{(k)} \in \mathcal{V}$ for $h = 0, \dots, H_k$, where $H_k \in \mathbb{Z}_+$ denotes the number of edges (hops) on path $\mathcal{P}_k$. For consistency, we adopt the convention $n_0^{(k)} \triangleq s_k$ and $n_{H_k}^{(k)} \triangleq d_k$. Let $V(\mathcal{P}_k) \triangleq \{n_h^{(k)} : h \in \{0, \dots, H_k\}\}$ denote the constituent node set of $\mathcal{P}_k$. The path $\mathcal{P}_k$ thus traverses $H_k - 1$ intermediate relays. Unless stated otherwise, all routing decisions and the routing-derived quantities introduced below are evaluated at the current slot t . Their explicit dependence on t is suppressed for brevity. The overall system architecture is illustrated in Fig. 1.




Figure 1: System model of the multi-UAV semantic relay network.

All UAVs share identical hardware specifications. Specifically, each UAV is equipped with a full-charge battery of capacity $B_{\text{bat}} = 300 \text{ Wh} = 1.08 \times 10^6 \text{ J}$. Each UAV also carries an on-board compute platform yielding the per-hop inference cost $F\epsilon_{\text{comp}} + P_{\text{proc}}T_{\text{proc}}$. Here, $F \in \mathbb{Z}_+$ is the floating-point operation (FLOP) count of the deployed semantic codec, ϵ_{comp} (J/FLOP) is the compute energy efficiency, and P_{proc} (W), T_{proc} (s) denote the baseband processing power and processing time, respectively. Additionally, all UAVs employ identical air-to-air radio modules detailed in Section 3.1. Hence, the per-hop inference cost is identical across relays. However, the residual energy state $E_j^{\text{rem}}(t) \in [0, B_{\text{bat}}]$ evolves dynamically as each relay expends energy on its assigned flows. Notably, this state upper-bounds the per-node cumulative energy expenditure via a budget constraint introduced later in Section 3.4. To facilitate readability, all primary notation and symbols are summarized in Table 2.

3.1. Channel Model and Load-Count Contention Surrogate

We adopt the 3rd Generation Partnership Project (3GPP) probabilistic line-of-sight (LoS)/non-line-of-sight (NLoS) path loss model specified for aerial user equipment [1]. That specification models air-to-ground links between an aerial user equipment and a terrestrial base station. We repurpose it as an air-to-air (A2A) surrogate for UAV-to-UAV links, evaluated at the inter-UAV separation and altitude difference between the two communicating UAVs. The channel operates at carrier frequency $f_c = 2 \text{ GHz}$ with bandwidth $W = 10 \text{ MHz}$ and transmit power $P_{\text{tx}} = 23 \text{ dBm}$. Hereafter, P_{tx} denotes the corresponding linear-scale value (i.e., $P_{\text{tx}} \approx 0.2 \text{ W}$) whenever it appears inside the energy expression (10). This configuration yields a nominal signal-to-noise ratio (SNR) of $\gamma^{(0)} = 6.6 \text{ dB}$ at the worst-case 3D range $d_{\text{max}} = 500 \text{ m}$ with altitude separation between communicating UAVs $\Delta_{\text{alt}} = 50 \text{ m}$ ($\theta \approx 5.74^\circ$). Under purely horizontal separation at the same maximum range ($\Delta_{\text{alt}} = 0 \text{ m}$, $\theta = 0^\circ$), the 3GPP line-of-sight probability attains its geometric minimum, yielding a nominal SNR of 5.9738 dB as the weakest admissible link margin in the topology (serving as the scheduled admission threshold γ_{th}

Table 2
Notation and Symbol Definitions.

Symbol	Description
$\mathcal{V}, N$	Set of UAV relay nodes and total fleet size ($ \mathcal{V} = N$)
$\mathcal{E}(t), d_{\text{max}}$	Directed edge set at slot t and maximum range d_{max}
$K, (s_k, d_k)$	Number of semantic flows and source-destination pairs
P_k, H_k	Unsplittable routing path and hop count of flow k
$B_{\text{bat}}, E_j^{\text{rem}}(t)$	Battery capacity ($300 \text{ Wh} = 1.08 \times 10^6 \text{ J}$) and residual energy
P_{tx}, W, f_c	Transmit power (23 dBm), bandwidth (10 MHz), frequency (2 GHz)
L_j, C_{max}	Relay load at node j and maximum concurrent session cap
$\gamma_{ij}^{(0)}, \gamma_{ij}^{\text{eff}}$	Nominal SNR and load-degraded effective SNR
C_{ij}^{eff}	Effective Shannon link capacity under load-count contention
$n_b, F, \epsilon_{\text{comp}}$	Semantic payload size, FLOP count, compute efficiency
$S(\gamma), S_{\text{e2e}}$	Per-hop semantic fidelity and end-to-end cascade fidelity
$k_s, \gamma_{0.5}, S_{\text{min}}, S_{\text{max}}$	Per-hop fidelity-curve steepness, midpoint, and task-score bounds
q_h, ϕ	Cumulative delivered task score after h hops and codec state-protection coefficient
$E_{ij}^{(k)}, E_{\text{total}}(P_k)$	Per-hop energy drain and total path energy expenditure
$E_j^{\text{node}}, E_j^{\text{sig}}$	Per-node epoch drain (J/epoch) and network signaling energy (J/epoch)
$S_k^{\text{req}}, E_{\text{max}}^{(k)}$	Per-flow fidelity floor and per-flow route-energy cap (J/task)
$P_{\text{prop}, j}, T_{\text{mis}}, T_{\text{epoch}}$	Mission propulsion power (route-independent), mission horizon, epoch duration
$\text{AoSI}_{p_k}, \text{AoSI}_{\text{max}}$	Age of Semantic Information and maximum staleness bound
$\text{SREE}(P_k)$	Semantic Routing Energy Efficiency ($S_{\text{e2e}}/E_{\text{total}}$)

in Section 4.2). Since all methods share identical channel conditions, the moderate link margin is sufficient for valid routing comparisons. In particular, the margin follows from the worst-case low-altitude geometry stated above. It is not calibrated against A2A propagation measurements. Each link (i, j) is associated with a nominal SNR $\gamma_{ij}^{(0)}$ (dB) and Shannon capacity $C_{ij}^{(0)} = W \log_2(1 + 10^{\gamma_{ij}^{(0)}/10})$ (bit/s).

When multiple concurrent flows traverse a common relay j , medium-access contention degrades the effective link quality. We model this degradation via a load-count contention surrogate that prices interference by the number of flows forwarded by a relay, rather than by the explicit medium-access schedule. Accordingly, for each $j \in \mathcal{V}$, we let $L_j \in \{0, 1, \dots, K\}$ denote the dimensionless relay load, i.e., the number of flows passing through j as an intermediate

node, which is expressed as:

$$L_j = \sum_{k=1}^K \mathbf{1}[j \in V(\mathcal{P}_k) \setminus \{s_k, d_k\}], \quad (2)$$

where $\mathbf{1}[\cdot]$ is the indicator function that returns 1 if its predicate is true and 0 otherwise. By construction, L_j excludes flow k 's own source s_k and destination d_k , since endpoints are not subject to forwarding contention from flow k itself. Nevertheless, s_k and d_k can still incur load when other flows traverse them as relays. The effective SNR $\gamma_{ij}^{\text{eff}} \in \mathbb{R}$ (dB) on each link $(i, j) \in \mathcal{E}(t)$ then decreases logarithmically with the load and is given by:

$$\gamma_{ij}^{\text{eff}} = \gamma_{ij}^{(0)} - 10 \log_{10}(1 + L_j) \quad [\text{dB}], \quad (3)$$

Specifically, a single flow traversing a relay ($L_j = 1$) already incurs a 3 dB half-duplex penalty. Each additional flow forwarded by the same relay then incurs an incremental load-count penalty. The effective capacity under load (bit/s) is given by:

$$C_{ij}^{\text{eff}} = W \log_2(1 + 10^{\gamma_{ij}^{\text{eff}}/10}). \quad (4)$$

Alternative contention laws dependent solely on L_j (such as linear $3L_j$ dB or quadratic $10 \log_{10}(1 + L_j^2)$ dB) also increase monotonically in L_j . Since L_j depends on the routing decisions of all K semantic flows, we write $\gamma_{ij}^{\text{eff},(k)}$ and $C_{ij}^{\text{eff},(k)}$ when the flow index must be explicit.

Notably, (3) represents a *load-count contention surrogate* that prices contention by the scalar relay load L_j rather than transmitter geometry, concurrent transmission schedules, or spatial co-channel interference. Consequently, the formulated objective and primary routing comparisons inherit this scope. A complementary spatial evaluator assuming simultaneous transmission on shared resources is analyzed in Supplementary Material Section S-D.1 as a worst-case full-reuse stress test, with ranking differences discussed in Section 5.6. Furthermore, a scheduled-objective variant evaluated under explicit time-division multiple access (TDMA) is formulated in Section 4. Meanwhile, the load-count contention surrogate remains the operational model for all other comparisons.

Importantly, the effective SNR reduction (3) simultaneously affects both the available transmission rate and the semantic fidelity. On the one hand, the reduced capacity C_{ij}^{eff} increases the transmission time n_b/C_{ij}^{eff} and hence the per-hop energy (10). On the other hand, the lower SNR degrades the per-hop semantic fidelity $S(\gamma_{ij}^{\text{eff}})$ through the per-hop fidelity map (6). The routing formulation captures both effects through a single coupling variable γ_{ij}^{eff} . The MF-SREE objective (17a) then jointly balances the two. The Shannon capacity governs energy accounting. In contrast, the per-hop fidelity map captures the task-performance envelope of a learned semantic codec operating under channel noise [43]. This dual-model design is consistent with prior multi-hop semantic communication studies [18, 26].

3.2. Semantic Communication Pipeline

In this subsection, we present the semantic communication pipeline operating along each relay route. Each flow k carries a semantic task instance. The source UAV observes raw data $x_k \in \mathcal{X}$ from a task-specific input space $\mathcal{X}$ (e.g., images or sensor frames). It applies a learned encoder $f_\theta : \mathcal{X} \rightarrow \{0, 1\}^{n_b}$ with trainable parameters $\theta \in \mathbb{R}^{p_\theta}$ to extract a compact semantic representation $z_k = f_\theta(x_k) \in \{0, 1\}^{n_b}$ of $n_b \in \mathbb{Z}_+$ bits. Learned semantic encoders such as DeepSC [43] emit real-valued symbol vectors, so the binary representation adopted here is an abstraction of this paper. It expresses the transmitted payload in bits, so that transmission time and per-hop energy follow from the link capacity. The per-hop semantic fidelity $S(\gamma) \in [0, 1]$ quantifies the expected normalized task utility conditioned on link SNR:

$$S(\gamma) \triangleq \mathbb{E}[u(y_k, \hat{y}_k) \mid \gamma]. \quad (5)$$

Here, $\mathbb{E}[\cdot]$ averages over task data, channel noise, and codec stochasticity. The ground-truth task output $y_k \in \mathcal{Y}$ is drawn from output space $\mathcal{Y}$, $\hat{y}_k \in \mathcal{Y}$ is the receiver-side decoded prediction, and $u : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, 1]$ is a normalized task-specific score (e.g., classification accuracy, mAP, or SSIM).

3.3. Semantic Fidelity, Energy, and Staleness

In this subsection, we instantiate the conditional utility (5) for the two codec profiles used in this paper, which share one per-hop functional family. Deep JSCC image transmission [8] and learned semantic text codecs [43] both report task scores that increase monotonically with channel SNR and saturate at a codec-dependent ceiling. Building on that qualitative behavior, we adopt a bounded logistic function of the effective SNR as the per-hop fidelity family. This functional form is a modeling choice of this paper:

$$S(\gamma) = S_{\min} + (S_{\max} - S_{\min}) \text{sig}(k_s(\gamma - \gamma_{0.5})), \quad (6)$$

where $\text{sig}(x) \triangleq 1/(1 + e^{-x})$, $x \in \mathbb{R}$, is the standard logistic function and $0 \leq S_{\min} \leq S_{\max} \leq 1$ are dimensionless task-score bounds. Here, $S_{\max}$ is the high-SNR task-score ceiling and $S_{\min}$ the low-SNR floor. The midpoint $\gamma_{0.5}$ (dB) is the SNR at which the codec reaches half its usable range, below which fidelity falls sharply toward $S_{\min}$ (cliff effect) and above which it saturates toward $S_{\max}$. The slope $k_s > 0$ (dB⁻¹) sets the steepness of that transition.

Two codec profiles carry the results of this paper, and a third parameter set is disclosed below as a calibration record. All three are distinguished explicitly wherever a result depends on the choice. The *synthetic* profile evaluates (6) at ($S_{\min} = 0.05$, $S_{\max} = 0.95$, $\gamma_{0.5} = 10.0$ dB, $k_s = 0.30$ dB⁻¹) and composes it across hops as the per-hop product (8) introduced below. It is a normalized reference profile selected for interpretability. It is fitted to no codec. Its 10.0 dB midpoint deliberately places the tested link margins on the steep portion of the curve. It serves as a stress profile. Absolute fidelity levels reported under it carry no claim of codec realism.

The *measured* profile is a distinct empirical artifact that carries the primary codec evidence of this paper. It is the

curve under which the principal grid, the binding per-flow QoS sweep of Section 5.5, and the measured-codec arm of the scheduled crossed evaluation are all produced. It composes delivered task score along a route by a measured state-and-SNR recursion:

$$q_h = q_{h-1} \left[1 - (1 - S(\gamma_h))(1 - \varphi q_{h-1}) \right], \quad (7)$$

where $q_h \in [0, 1]$ is the cumulative one-gram bilingual evaluation understudy (BLEU) score delivered after h hops, initialized at $q_0 = 1$ for a fresh source observation, and γ_h (dB) is the effective SNR of hop h . The coefficient $\varphi \in [0, 1]$ is a state-protection term that couples the loss incurred on a hop to the quality already carried into it. The curve parameters and φ were fitted by least squares to the 60 hop-to-hop BLEU transitions of 21 training SNR sequences, at a transition root-mean-square residual of 0.0250. Those transitions were decoded from 46,080 chain records over 512 test-split Europarl sentences, three stochastic seeds, and one- to four-hop chains of the public DeepSC implementation [43] under an additive white Gaussian noise channel. On nine held-out sequences the recursion attains a final-hop mean absolute error of 0.0908 (95% bootstrap confidence interval [0.0873, 0.0948]), against 0.5826 ([0.5765, 0.5891]) for the synthetic profile composed as a product of single-hop values. The fitted state protection is numerically zero, falling below 10^{-31} . Consequently, (7) reduces exactly to the per-hop product (8) with (6) evaluated at ($S_{\min} = 0, S_{\max} = 0.9744, \gamma_{0.5} = -0.8934$ dB, $k_s = 0.7205$ dB $^{-1}$). The two profiles share the per-hop family and the composition law, differing in codec parameters alone. A contrast between them therefore isolates the codec curve.

The measured midpoint lies 10.89 dB below the synthetic one. This shift determines where the tested links fall on the fidelity curve. At the nominal margin $\gamma^{(0)} = 6.6$ dB of Section 3.1, the synthetic profile places such a link 3.4 dB below its midpoint, on the steep portion of (6). The load-induced degradation (3) then moves it further down that portion. The measured profile places the same link 7.5 dB above its midpoint, in the saturating region. A given difference in effective SNR therefore maps to a much larger fidelity difference under the synthetic profile than under the measured one. Consequently, most contrasts reported under it are wider than their measured counterparts. Unless a result is explicitly reported under the measured DeepSC codec model, it is produced under the synthetic profile. Each subsection of Section 5 states which profile is in force.

A third set, the *calibrated* single-hop fit, is reported for calibration transparency. The underlying metric is one-gram BLEU. Its nine operating points were digitized off the published DeepSC evaluation figure of [43] for an additive white Gaussian noise channel, spanning -6 to 18 dB. The fit is to published results rather than to a codec rerun here. It is unweighted because no per-point uncertainty is published. Least-squares fitting of (6) to those points yields ($S_{\min} = 0.0567, S_{\max} = 0.9235, \gamma_{0.5} = -2.9219$ dB, $k_s = 0.3083$ dB $^{-1}$). The fit attains a root-mean-square residual of 0.0033 and a coefficient of determination $R^2 = 0.9998$. Its

midpoint is 12.92 dB, that is, approximately 13 dB, below the synthetic one.

Several diagnostics characterize the confidence attached to that fit. On five degrees of freedom, the 95% Wald parameter confidence intervals are $[-0.0270, 0.1403]$ for $S_{\min}$, $[0.9162, 0.9307]$ for $S_{\max}$, $[0.2767, 0.3398]$ dB $^{-1}$ for k_s , and $[-3.631, -2.213]$ dB for $\gamma_{0.5}$. Because its interval covers zero, $S_{\min}$ is the least determined parameter. The low-SNR floor is not separable from zero at that confidence level. With only nine points, these asymptotic linearized Wald intervals are indicative rather than exact small-sample intervals. Across the digitized points, the widest 95% prediction interval occurs at -6 dB and spans $[0.2826, 0.3146]$. The residuals have a standard deviation of 0.0035 and a largest magnitude of 0.0070 at 0 dB. A Shapiro–Wilk test on them returns $W = 0.9652$ at $p = 0.8513$, which reads as no evidence against normality. Outside -6 to 18 dB the curve is an extrapolation of a nine-point fit. It is not quoted here as a measurement. This single-hop fit is the codec-calibration record of this paper. It produces no result reported in the main manuscript. The principal grid is produced under the measured recursion rather than under this fit, and so is the binding per-flow QoS sweep of Section 5.5. Its one live use is the codec-parameter sensitivity sweep of Supplementary Material Section S-D.2, which perturbs $\gamma_{0.5}$ and k_s about the parameter set above.

Assuming conditionally independent channel noise across decode-and-forward relays, the end-to-end fidelity for flow k is given by the product of per-hop fidelities:

$$S_{e2e}(\mathcal{P}_k) = \prod_{h=1}^{H_k} S(\gamma_{n_{h-1}, n_h}^{\text{eff}, (k)}), \quad (8)$$

which admits an exact log-domain additive representation:

$$\ln S_{e2e}(\mathcal{P}_k) = \sum_{h=1}^{H_k} \ln S(\gamma_{n_{h-1}, n_h}^{\text{eff}, (k)}), \quad (9)$$

where $\mathcal{P}_k$ is the path, H_k the hop count, and $\gamma_{n_{h-1}, n_h}^{\text{eff}, (k)}$ the effective SNR on hop h . The relay load $L_{n_h}^{(k)}$ aggregates every flow traversing $n_h^{(k)}$. Sweeping $k = 1, \dots, K$ thus yields K flow-specific fidelity cascades. This multiplicative structure directly results from the decode-and-forward (DF) relay cascade. Suppressing the flow index k for clarity, each relay h decodes the upstream representation $\hat{z}_{h-1}$ to a task estimate $\hat{y}_{h-1}$ and re-encodes it to $\hat{z}_h \in \{0, 1\}^{n_b}$ (with $\hat{z}_0 \triangleq z_k$). Denoting cumulative fidelities at relay h input and output as $S_{\text{in}, h}$ and $S_{\text{out}, h}$, we have $S_{\text{out}, h} = S(\gamma_h^{\text{eff}}) \cdot S_{\text{in}, h}$. Iterating from $h = 1$ to H_k produces $S_{e2e} = \prod_{h=1}^{H_k} S(\gamma_h^{\text{eff}})$. The cliff effect of (6), triggered whenever any $\gamma_h^{\text{eff}} < \gamma_{0.5}$, sharply degrades path fidelity and makes route selection consequential [44, 46]. The synthetic midpoint places the tested links on that cliff, whereas the measured midpoint places them above it.

Alongside fidelity, each relay hop incurs energy costs from payload transmission, decoding, and processing. The

transmitter i pays over-the-air transmission energy, whereas receiver j pays neural-inference and baseband-processing costs. The per-hop link energy (J) is:

$$E_{ij}^{(k)} = \underbrace{P_{\text{tx}} \frac{n_b}{C_{ij}^{\text{eff},(k)}}}_{e_{ij}^{\text{tx},(k)} \text{ (paid by } i)} + \underbrace{F\epsilon_{\text{comp}} + P_{\text{proc}}T_{\text{proc}}}_{e^{\text{rx}} \text{ (paid by } j)}, \quad (10)$$

where $e_{ij}^{\text{tx},(k)} \triangleq P_{\text{tx}} n_b / C_{ij}^{\text{eff},(k)}$ and $e^{\text{rx}} \triangleq F\epsilon_{\text{comp}} + P_{\text{proc}}T_{\text{proc}}$ denote per-hop transmit and receive energies. Here, n_b (bits) is the encoded semantic payload, and $C_{ij}^{\text{eff},(k)}$ (bit/s) is the load-dependent effective capacity (4).

Summing over all hops yields the total path energy (J):

$$E_{\text{total}}(\mathcal{P}_k) = \sum_{h=1}^{H_k} E_{n_{h-1}, n_h}^{(k)}. \quad (11)$$

Similarly, the cumulative energy drawn at UAV node j (J) aggregates transmission, decode-and-forward, and epoch-level control signaling across traversing flows:

$$\begin{aligned} E_j^{\text{node}}(\{\mathcal{P}_k\}_{k \in \mathcal{K}_{\text{adm}}}) &= \sum_{k \in \mathcal{K}_{\text{adm}}} \mathbf{1}[j \in V(\mathcal{P}_k)] \\ &\times \left(e_{j, \text{succ}_k(j)}^{\text{tx},(k)} \mathbf{1}[j \neq d_k] + e^{\text{rx}} \mathbf{1}[j \neq s_k] \right) \\ &+ E_j^{\text{sig}}, \end{aligned} \quad (12)$$

where $\mathcal{K}_{\text{adm}}$ is the admitted-flow index set with $K_{\text{adm}} \triangleq |\mathcal{K}_{\text{adm}}|$, and $\text{succ}_k(j) \in \mathcal{V}$ denotes the successor of node j on path $\mathcal{P}_k$ for $j \in V(\mathcal{P}_k) \setminus \{d_k\}$, with $\text{succ}_k(d_k) \triangleq d_k$ and $e_{d_k, d_k}^{\text{tx},(k)} \triangleq 0$. Indicators $\mathbf{1}[j \neq d_k]$ and $\mathbf{1}[j \neq s_k]$ ensure that the transmitter pays the communication cost and the receiver pays the decoding and processing cost. The term E_j^{sig} is every node's share of the epoch-level control signaling defined in (13) below, because each UAV transmits a beacon irrespective of whether it lies on a delivered data route. Thus, E_j^{node} couples the routing decisions of all admitted flows through node j . Constraint (17c) requires that aggregate drain respect residual battery $E_j^{\text{rem}}(t)$.

Control signaling is charged at the epoch scope through a single auditable radio model. Over one decision epoch of duration T_{epoch} , the centralized controller aggregates one beacon of b_{bc} bits from each of the N UAVs. The controller then disseminates one forwarding-table entry of b_{fe} bits to each on-route node of each of the K_{adm} admitted flows. With mean relay depth $\bar{H}$ over those routes, the network-wide control energy (J per epoch) is:

$$E^{\text{sig}} = \left(\underbrace{Nb_{\text{bc}}\bar{H}}_{\text{aggregation}} + \underbrace{K_{\text{adm}}b_{\text{fe}}\bar{H}^2}_{\text{dissemination}} \right) \epsilon_{\text{bit}}, \quad \epsilon_{\text{bit}} = \frac{P_{\text{tx}}}{C(\bar{\gamma})}, \quad (13)$$

where ϵ_{bit} (J/bit) is the radio energy per bit at the mean link SNR $\bar{\gamma}$, consistent with the transmit term of (10). The per-node share entering (12) is $E_j^{\text{sig}} \triangleq E^{\text{sig}}/N$. Notably, (13) is

the only signaling model used in this paper, and it is stated network-wide per epoch rather than per node.

Building on the terms introduced above, we now fix a single energy-accounting boundary for the remainder of the paper. Every energy quantity reported belongs to exactly one of three scopes, and quantities from different scopes are never summed or compared directly. The first scope is the *per-flow route energy* $E_{\text{total}}(\mathcal{P}_k)$ [J per delivered task], which covers over-the-air transmission and receive-side inference plus baseband processing, accumulated over the hops of one route via (10)–(11). It is the only scope entering the SREE ratio (16). The second is the *per-node epoch energy* E_j^{node} [J per epoch], which aggregates the route-energy contributions of every flow traversing node j and adds that node's signaling share E_j^{sig} via (12). It is the only scope entering the battery constraint (17c). The third is the *mission-level propulsion energy* $E_{\text{prop},j} = P_{\text{prop},j}T_{\text{mis}}$ [J per mission], where T_{mis} is the mission horizon. The fleet follows a trajectory that is fixed for the whole mission and is not a decision variable of the routing problem. Hence, $P_{\text{prop},j}$ is taken to be identical across candidate routing assignments. Propulsion is therefore a route-independent additive constant that lies outside the routing boundary. It is excluded from (10), from (12), and from the SREE ratio (16). Route-independence is a structural modeling premise. Varying $P_{\text{prop},j}$ thus illustrates this accounting convention rather than validating it, as shown in the propulsion sweep of Section 5.4. Accordingly, the SREE ratio (16) and MF-SREE objective (17a) strictly quantify communication-layer energy efficiency (transmission, semantic inference, and baseband processing) rather than total UAV mission energy. Signaling illustrates why these scopes must be separated. At $N = 50$, requested $K = 50$ and $K_{\text{adm}} = 49$, $b_{\text{bc}} = 256$ bits, $b_{\text{fe}} = 32$ bits, and $T_{\text{epoch}} = 1.0$ s, (13) yields $E^{\text{sig}} = 1.19 \times 10^{-4}$ J per epoch for the entire network. In contrast, the same 49-flow assignment draws 30.6 J of route energy (roughly 0.62 J per delivered flow). Control overhead is thus over five decimal orders below aggregate data-plane energy. Hence, signaling is incorporated into the per-node battery constraint (17c) and omitted from the per-flow ratio (16).

To characterize semantic staleness, we define an Age of Semantic Information (AoSI) for each flow. AoSI is a latency-based staleness quantity defined in this paper. It is distinct from the Age of Semantic Importance of Chen and Gong [9]. That metric weights information freshness by the semantic importance of each update. The two formulations therefore share an acronym but not a definition. First, end-to-end path latency [s] is introduced as:

$$D_{\mathcal{P}_k} = \sum_{h=1}^{H_k} \left(\frac{n_b}{C_{n_{h-1}, n_h}^{\text{eff},(k)}} + T_{\text{proc}} \right), \quad (14)$$

which combines transmission time and processing latency. Letting $\Delta_c(\mathcal{P}_k) = \sum_{h=1}^{H_k} \delta_c^{(h)} \geq 0$ denote accumulated context drift over the path (in dimensionless context units),

AoSI is defined as:

$$\text{AoSI}_{P_k} = D_{P_k} \exp(\lambda_{\text{AoSI}} \Delta_c(P_k)) \quad [\text{s}], \quad (15)$$

where $\lambda_{\text{AoSI}} > 0$ controls AoSI sensitivity to context drift. For stationary ergodic source processes, $\Delta_c = \kappa D_{P_k}$ with $\kappa > 0$ (s^{-1}) a source-correlation rate. Larger λ_{AoSI} penalizes context change exponentially more aggressively. In the limit $\lambda_{\text{AoSI}} \rightarrow 0$, AoSI reduces to classical latency D_{P_k} . Notably, (15) depends strictly on path latency D_{P_k} and context drift $\Delta_c(P_k)$, omitting mobility speed, topology-update intervals, and TDMA queueing delays. The staleness constraint (17d) inherits this focused scope.

3.4. MF-SREE Problem Formulation

Having established the fidelity and energy models, we now formulate the routing design problem. Let $\text{SREE}(P_k)$ denote the per-flow Semantic Routing Energy Efficiency (SREE), which quantifies delivered semantic quality per unit of energy:

$$\text{SREE}(P_k) \triangleq \frac{S_{\text{e2e}}(P_k)}{E_{\text{total}}(P_k)} \quad [\text{semantic units/J}]. \quad (16)$$

Explicit load dependence is written as $\text{SREE}(P_k; \mathbf{L})$, where effective SNRs and capacities in (8) and (10) are evaluated under load profile $\mathbf{L}$. The SREE metric generalizes the Semantic Energy Efficiency (SEE) of Tang *et al.* [40] from a single-relay topology to a route-level multi-hop setting. Rejection is a first-class outcome of this design problem, where flows failing budget constraints on all candidate routes are refused rather than delivered. Accordingly, letting $\mathcal{K}_{\text{adm}} \subseteq \{1, \dots, K\}$ denote the admitted-flow set introduced in (12) with $K_{\text{adm}} = |\mathcal{K}_{\text{adm}}|$, candidate solutions are ordered lexicographically over $(K_{\text{adm}}, \sum_{k \in \mathcal{K}_{\text{adm}}} \text{SREE}(P_k))$. Under energy, staleness, service-availability, and connectivity constraints, the optimization problem is formulated as:

$$\max_{\mathcal{K}_{\text{adm}}, \{P_k\}}^{\text{lex}} \left(K_{\text{adm}}, \sum_{k \in \mathcal{K}_{\text{adm}}} \text{SREE}(P_k) \right) \quad (17a)$$

$$\text{s.t. } (n_{h-1}^{(k)}, n_h^{(k)}) \in \mathcal{E}(t), \quad \forall h \leq H_k, \quad \forall k \in \mathcal{K}_{\text{adm}}, \quad (17b)$$

$$E_j^{\text{node}}(\{P_k\}_{k \in \mathcal{K}_{\text{adm}}}) \leq E_j^{\text{rem}}(t), \quad \forall j \in \mathcal{V}, \quad (17c)$$

$$\text{AoSI}_{P_k} \leq \text{AoSI}_{\text{max}}, \quad \forall k \in \mathcal{K}_{\text{adm}}, \quad (17d)$$

$$H_k \leq H_{\text{max}}, \quad \forall k \in \mathcal{K}_{\text{adm}}, \quad (17e)$$

$$L_j \leq C_{\text{max}}, \quad \forall j \in \mathcal{V}, \quad (17f)$$

$$S_{\text{e2e}}(P_k) \geq S_k^{\text{req}}, \quad \forall k \in \mathcal{K}_{\text{adm}}, \quad (17g)$$

$$E_{\text{total}}(P_k) \leq E_{\text{max}}^{(k)}, \quad \forall k \in \mathcal{K}_{\text{adm}}, \quad (17h)$$

where $\max^{\text{lex}}$ orders candidate solutions by admitted flows first, then by summed SREE. Flows $k \notin \mathcal{K}_{\text{adm}}$ contribute zero to both keys; per-flow constraints (17b) and (17d)–(17h) span $\mathcal{K}_{\text{adm}}$ alone, whereas per-node constraints (17c) and (17f) aggregate drain and relay load over

all nodes in $\mathcal{V}$. Here, $\text{AoSI}_{\text{max}} > 0$ (s) is the maximum admissible AoSI, $S_k^{\text{req}} \in [0, 1]$ is the required semantic quality floor, and $E_{\text{max}}^{(k)} > 0$ (J/task) is the per-flow route-energy cap. Dimensionless limits $H_{\text{max}}, C_{\text{max}} \in \mathbb{Z}_+$ bound hop count and concurrent relay sessions, respectively. Specifically, (17b) enforces link connectivity, (17c) guards battery $E_j^{\text{rem}}(t)$ via cumulative drain (12), (17d) bounds staleness, (17e) caps path length, and (17f) enforces relay service availability $L_j \leq C_{\text{max}}$ [3], where each traversing flow consumes one relay session per (2).

The three budget constraints operate at distinct scopes: (17c) couples concurrent flows at the per-node battery level, (17h) caps per-flow task energy, and (17g) specifies per-flow traffic fidelity requirement S_k^{req} (distinct from codec floor S_{min}). Default evaluations in Section 5 set $E_{\text{max}}^{(k)} = \infty$ and $S_k^{\text{req}} = 0$, whereas Section 5.5 sweeps both under measured DeepSC models. Notably, in binding regimes, efficiency rankings must be read beside K_{adm} because the router prioritizes flow admissions. Alternative objectives (e.g., max-min SREE fairness $\max \min_k \text{SREE}(P_k)$ or log-utility fairness $\sum_k \log \text{SREE}(P_k)$) can be incorporated by adjusting candidate selection in Section 4.1.

4. Proposed Solution

To address the computational difficulty of the MF-SREE problem (17), in this section, we propose S-ACS MF, a meta-heuristic framework tailored for multi-UAV relay networks. Solving Problem (17) is challenging because its decision version is NP-hard at unit service availability ($C_{\text{max}} = 1$), by a polynomial-time reduction from the directed 2-vertex-disjoint paths (2-VDP) problem [16] (Theorem 9, Appendix A.2). At $C_{\text{max}} = 1$, feasibility forces internally disjoint paths. The decision target is the lexicographic pair itself, namely $(K_{\text{adm}}, \sum_k \text{SREE}_k) \geq_{\text{lex}} (K, \Lambda_{\text{target}})$ for a target aggregate SREE $\Lambda_{\text{target}} > 0$. Meeting it thus requires solving 2-VDP. The first key must be named here rather than left implicit, because rejection is admissible in (17): a threshold on the efficiency key alone could be met by admitting one short flow and refusing the rest, which forces no disjointness. Hardness therefore covers exactly the regime in which the first key is at its maximum $K_{\text{adm}} = K$, which is the regime the router itself pursues first. Hardness is not claimed for the non-binding regime $C_{\text{max}} = N$ used in Section 5, where the flows remain coupled through load-count contention and the objective remains non-decomposable (Remark 1). Meta-heuristics trade an optimality guarantee for tractability on such coupled problems. We thus propose S-ACS MF, a meta-heuristic framework tailored to jointly optimize routing and flow ordering. Specifically, it leverages a load-aware energy-per-fidelity construction rule, a coupled-marginal selection step, and a coordinated reroute local search on the true coupled total MF-SREE. A fixed-load linear programming (LP) relaxation bound provides a diagnostic gap analysis, complemented by formal convergence and complexity analyses. We further define a scheduled routing objective, and

the S-ACS MF-SINR variant that optimizes it, alongside the retained load-count contention surrogate.

4.1. S-ACS MF: Multi-Flow Semantic Ant Colony System

To apply ant colony optimization (ACO) to the MF-SREE problem (17), in this subsection, we extend the classical ACO framework [13] to the multi-flow semantic routing setting. The extension introduces the three complementary mechanisms named above: load-aware construction, coupled-marginal selection, and coordinated reroute. Ants preferentially select relays that preserve more task utility $S(\gamma)$ per joule of neural inference and transmission energy, rather than merely following high-SNR links.

Before presenting the proposed algorithm, we summarize the notation introduced for the metaheuristic framework. Let $\mathcal{N}_i \triangleq \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\}$ denote the feasible-neighbor set of node i on the current routing snapshot. Let $\Pi(s_k, d_k)$ denote the set of simple s_k -to- d_k paths in $\mathcal{G}$. Among the algorithmic parameters, M denotes the number of ants per iteration, $I_{\max}$ the iteration budget, and τ_0 the initial pheromone level. Additionally, $Q > 0$ is the pheromone-deposit scaling constant, and $\rho_{\text{base}} \in (0, 1)$ is the baseline evaporation rate. Furthermore, R_{LS} is the cap on coordinated-reroute sweeps, and $\zeta \geq 0$ is the congestion sensitivity coefficient. Without loss of generality, the routing problem is solved on a frozen snapshot of $\mathcal{G}(t)$ at a decision epoch t . The time argument is therefore suppressed throughout Section 4.

4.1.1. Routing Heuristic and Transition Probability

To capture the dual trade-off of semantic quality and energy consumption, the S-ACS MF heuristic encodes two attributes per edge—load-modulated semantic fidelity $S(\gamma_{ij}^{\text{eff},(k)})$ and load-modulated per-hop energy $E_{ij}^{(k)}$ —combined as their ratio. All subsequent algorithmic components (transition probability, pheromone update, load-aware candidate generation) operate on this two-attribute edge cost.

At each iteration, the K semantic flows are routed sequentially in a randomly permuted order. When constructing paths for a given flow k , the algorithm dynamically accounts for the partial relay load $L_j^{(<k)}$, representing the load (2) induced by previously routed flows ($k' < k$). This partial load updates both the effective SNR (3) and the per-hop energy (10).

During hop-by-hop construction by ant m located at node i , the candidate next-hop set is restricted to the feasibility-constrained subset $\mathcal{N}_{\text{feas}}(i, k)$, defined as:

$$\mathcal{N}_{\text{feas}}(i, k) \triangleq \{j \in \mathcal{N}_i : L_j < C_{\max}, |\mathcal{P}_m^{(k)}| < H_{\max}, E_j^{\text{node}} + E_{ij}^{(k)} \leq E_j^{\text{rem}}\}. \quad (18)$$

This subset jointly enforces the service-availability (17f), hop-limit (17e), and energy-budget (17c) constraints, where $|\mathcal{P}_m^{(k)}|$ denotes the current partial path length.

Ants select the next relay $j \in \mathcal{N}_{\text{feas}}(i, k)$ according to the transition probability:

$$p_{ij}^{(k)} = \frac{[\tau_{ij}^{(k)}]^{\alpha_p} [\eta_{ij}^{(k)}]^{\beta_p}}{\sum_{\ell \in \mathcal{N}_{\text{feas}}(i, k)} [\tau_{i\ell}^{(k)}]^{\alpha_p} [\eta_{i\ell}^{(k)}]^{\beta_p}}, \quad (19)$$

where $\tau_{ij}^{(k)}$ denotes the pheromone level that commodity k maintains on edge (i, j) , while α_p and β_p are the pheromone and heuristic exponents, respectively. The K pheromone matrices are independent in storage, one per commodity. Cross-flow information is carried instead by the load-dependent heuristic, the congestion-dependent evaporation rate, and the coupled-marginal selection rule (25). This evaporation rate is computed from the load of all K flows and applied to every per-flow matrix. To quantify the desirability of a link, the per-hop heuristic is evaluated in the log-domain as:

$$\eta_{ij}^{(k)} = \exp(\ln S(\gamma_{ij}^{\text{eff},(k)}) - \ln E_{ij}^{(k)}). \quad (20)$$

Here, $\gamma_{ij}^{\text{eff},(k)}$ and $E_{ij}^{(k)}$ are the effective SNR and edge energy, evaluated under the accumulated partial load. The log-form parameterization (20) is algebraically equivalent to the product-form $S(\gamma_{ij}^{\text{eff},(k)})/E_{ij}^{(k)}$ at the per-hop level. However, this parameterization is essential for numerical stability. Specifically, it prevents the underflow that otherwise arises when raising small semantic fidelities to the power β_p in (19).

Remark 1 (Heuristic Design Rationale). *The exact MF-SREE objective combines a multiplicative fidelity numerator $\prod_h S_h$ with an additive energy denominator $\sum_h E_h$. This violates Bellman's principle of optimality: the SREE-optimal s-v subpath does not strictly extend to the SREE-optimal s-d path [17]. To enable a greedy per-hop construction, both the log-domain heuristic (20) and the Dijkstra ratio weight E_{ij}/S in (24) adopt the same additive surrogate $w_{ij} = E_{ij}/S(\gamma_{ij}^{\text{eff}})$. Although not strictly equivalent at the path level, this surrogate preserves the per-edge fidelity-energy ranking that drives candidate generation. All final paths are scored under the true SREE objective (16), ensuring the mismatch is confined entirely to the candidate-generation stage. The size of that mismatch is quantified in Proposition 10 and Appendix A.3: the empirical gap over the sampled source-destination pairs stays below 2.6% at each of the four tested sizes. That figure bounds the average and not the worst pair, whose largest sampled gap is 30.5%. It provides evidence that the surrogate suffices for candidate generation, rather than a per-instance guarantee.*

4.1.2. Pheromone Mechanisms

To guide ant exploration toward high-quality paths, the routing memory is maintained per commodity: each flow k owns a pheromone matrix $\{\tau_{ij}^{(k)}\}$, initialized at $\tau_{ij}^{(k)} \leftarrow \tau_0$ on every edge and for every k . After routing all flows in an iteration, the iteration-best ant m_k^* of flow k deposits on that flow's own matrix. To maintain tight MAX-MIN Ant System (MMAS) bounds, this deposit is normalized by K in

the high-load mode, which is updated as:

$$\tau_{ij}^{(k)} \leftarrow \tau_{ij}^{(k)} + \frac{Q \cdot \text{SREE}(\mathcal{P}_{m_k^*})}{K_{\text{norm}}}, \quad \forall (i, j) \in \mathcal{P}_{m_k^*}, \quad (21)$$

where $K_{\text{norm}} = K$ when $K/N > 0.4$ (high- K mode) and $K_{\text{norm}} = 1$ otherwise. This normalization anchors the reference pheromone level to the per-flow quality rather than the aggregate total. Consequently, the MMAS bounds remain tight regardless of the number of concurrent flows.

Between iterations, every per-flow matrix undergoes congestion-dependent evaporation to induce load-balancing pressure:

$$\tau_{ij}^{(k)} \leftarrow (1 - \rho_{ij})\tau_{ij}^{(k)}, \quad \forall k, \quad (22)$$

where the common per-edge evaporation rate is dynamically adjusted as:

$$\rho_{ij} = \min(\rho_{\text{base}} + \zeta \max(L_i - 1, L_j - 1, 0), 0.9). \quad (23)$$

Pheromone values are then clipped to MMAS bounds $\tau_{ij}^{(k)} \in [\tau_{\min}, \tau_{\max}]$ once per iteration, after applying evaporation and deposits. The adaptive bounds are computed as:

$$\tau_{\max} = \frac{Q \cdot \text{SREE}_{\text{best}}^{\Sigma}}{K_{\text{norm}} \cdot \rho_{\text{base}}}, \quad \tau_{\min} = \frac{\tau_{\max}}{2N},$$

where $\text{SREE}_{\text{best}}^{\Sigma} = \sum_k \text{SREE}_k^*$ denotes the best-so-far total MF-SREE. The division by K_{norm} aligns $\tau_{\max}$ with the deposit normalization, maintaining standard MMAS convergence guarantees [39]. The two bounds are uniform across commodities as both anchor on the shared best-so-far total. Equations (21)–(23) together with this clipping rule constitute the complete pheromone update.

4.1.3. Candidate Generation and Local Refinement

To enhance candidate quality alongside M stochastic ants, each flow k receives one deterministic candidate $\mathcal{P}_D^{(k)}$ per iteration via a shortest-path search under ratio surrogate weights:

$$w_{ij}^{(k)} = \frac{E_{ij}^{(k)}}{S(\gamma_{ij}^{\text{eff},(k)})}. \quad (24)$$

These weights are rebuilt inside every iteration using the current accumulated load $\mathbf{L} \triangleq (L_1, \dots, L_N) \in \{0, 1, \dots, K\}^N$. The effective SNRs in (24) therefore incorporate the load-count contention penalties of previously committed flows (Algorithm 1). The resulting load-aware candidate directly competes with stochastic ant paths. This mechanism differs from single-pass zero-load seeding by executing dynamically every iteration under load-aware weights and adding candidates directly to Π_k^{feas} without mutating pheromones.

Adaptive Destination Exploration To prevent premature convergence to a suboptimal direct-to-destination shortcut, an adaptive destination exploration rule is enforced. Whenever the destination d_k is a neighbor of the current node i

($d_k \in \mathcal{N}_i$), the ant commits to d_k only if the direct edge is feasible under the physical constraints (17c)–(17e). Otherwise, d_k is masked from the sampling set, forcing the selection of an intermediate relay via (19).

Local-Search Post-Optimization To refine path assignments, a joint local-search step concludes each iteration. For each flow k , all candidate paths $\mathcal{P} \in \{\mathcal{P}_1^{(k)}, \dots, \mathcal{P}_M^{(k)}, \mathcal{P}_D^{(k)}\}$ retained at the current iteration are re-evaluated under the full joint load induced by the remaining $K-1$ flows. The algorithm updates the best path $\mathcal{P}_k^*$ whenever this evaluation strictly improves the global MF-SREE and preserves all constraints (17b)–(17f).

The two per-flow QoS budgets are treated as membership conditions of Π_k^{feas} rather than as a post-hoc screen. A candidate is removed from the pool before the selection step if its cascade fidelity falls below S_k^{req} (17g) or its route energy exceeds $E_{\max}^{(k)}$ (17h). Thus, the argmax cannot return an inadmissible route. Both stages of refinement apply the same test to the whole assignment, since adding contention at a shared relay can push a previously admissible peer below its floor. The final scoring stage also re-applies both budgets to every method under comparison. Any route that violates either budget is not counted as delivered. Consequently, feasibility is a property of every reported route rather than of the proposed algorithm alone.

A second refinement stage, the coordinated reroute refinement, runs once on the incumbent assignment after the iteration loop terminates. It is a worst-flow-first block-coordinate ascent on the true coupled total MF-SREE. Each sweep visits the flows in increasing order of their current per-flow SREE. For the visited flow it regenerates a single candidate by a shortest-path search under (24), evaluated against the full residual load of the other $K-1$ flows. Relays already saturated by those flows are masked out of that search. The candidate is thus capacity-feasible by construction (17f). The candidate is accepted only if it is energy- and capacity-feasible and strictly increases SREE^{Σ} once its own load is folded back in. Sweeps repeat until no flow improves or the cap R_{LS} is reached. The total is therefore non-decreasing, and the stage terminates.

4.1.4. Coupled-Marginal Selection

To address the coupling between sequential flow assignments, the selection step must align flow choices with the joint system objective. The construction pass commits one flow at a time, so the rule that picks a flow's path determines whether construction and refinement pursue the same objective. Let S denote the commodities already committed in the current permutation, with $\mathbf{L}^{(<k)}$ representing the load profile they induce. Additionally, let $\mathbf{I}(\mathcal{P})$ be the unit load a

candidate $\mathcal{P}$ deposits at its own interior nodes. The coupled-marginal selection step is then formulated as:

$$\mathcal{P}_k^* = \arg \max_{\mathcal{P} \in \Pi_k^{\text{feas}}} \left\{ \text{SREE}_k(\mathcal{P}; \mathbf{L}^{(<k)}) + \mathbf{I}(\mathcal{P}) \right. \\ \left. - \sum_{j \in S} [\text{SREE}_j(\mathcal{P}_j; \mathbf{L}^{(<k)}) - \text{SREE}_j(\mathcal{P}_j; \mathbf{L}^{(<k)}) + \mathbf{I}(\mathcal{P})] \right\}. \quad (25)$$

The rule maximizes the marginal contribution to the coupled objective rather than the flow's own utility. A candidate is credited with its own SREE under the load it itself creates. It is debited the contention-induced SREE loss it imposes on the commodities already sharing its relays. Because the debited term does not depend on the flows not yet routed, the argmax in (25) is exactly the greedy step on $\sum_j \text{SREE}_j$. Consequently, construction and the coordinated reroute refinement optimize the same quantity. A myopic argmax on $\text{SREE}_k(\mathcal{P}; \mathbf{L}^{(<k)})$ instead prices neither the load-count contention externality inflicted on already-committed flows nor the candidate's own self-load. Two properties bound the change. The debited term vanishes at $C_{\max} = 1$, where no relay reuse is feasible. Moreover, it is non-negative under any monotone contention law, so the priced score lower-bounds the myopic one. The rule is not one of the component-removal arms of the ablation study in Section 5.3. Its measured cost and benefit are evaluated by replaying both selection rules on identical seeds and topologies. This comparison is conducted at a non-binding design point and at a binding-capacity one. The rule re-prices relay sharing rather than suppressing it. No reduction in contention is claimed for it.

4.1.5. Adaptive Scaling for High-Density Flow Regimes

To handle dense multi-flow scenarios uniformly, S-ACS MF activates a high-load mode whenever the relative congestion satisfies $K/N > 0.4$. In this mode, the pheromone deposit and bounds in (21) are normalized by $K_{\text{norm}} = K$. Furthermore, the heuristic exponent in (19) is annealed to $\beta_{\text{p,eff}} = \max(0.5, \beta_p - 2.0(K/N - 0.4))$ to promote wider exploration. An optional load-diversity multiplier $(1 + L_j^{(<k)})^{-p}$ on the heuristic (20) is exposed by the implementation. However, it is switched off in every configuration reported here, with exponent $p = 0$ by default. Hence, congestion enters the construction process exclusively through the load-count contention penalty carried by the effective SNR (3), avoiding additional heuristic-side over-penalization. The activation threshold 0.4, attenuation slope 2.0, and exponent floor 0.5 are empirical constants calibrated on an $N = 50$ tuning grid. Outside the high-load regime ($K/N \leq 0.4$), $K_{\text{norm}} = 1$ and $\beta_{\text{p,eff}} = \beta_p$.

4.1.6. Algorithm Summary

The overall structure of S-ACS MF is presented in Algorithm 1. The outer loop controls the ACO iterations. Each iteration performs a complete K -flow sequential construction

pass, followed by a joint local-search refinement and a global pheromone update. The coordinated reroute refinement then runs once on the incumbent assignment after the loop terminates (Algorithm 2). Congestion information propagates across flows through three channels rather than through a shared pheromone matrix: the load-dependent heuristic (20), the congestion-dependent evaporation rate (23), and the coupled-marginal selection rule (25). The evaporation rate is computed from the load of all K flows and applied to every per-flow matrix. Thus, paths committed in iteration t still shape the landscape that guides every flow in iteration $t + 1$. Additionally, the commodity routing order is reshuffled at each iteration, so that no single flow is permanently routed last. Notably, a flow is not guaranteed a route. Whenever the per-flow filter leaves the candidate pool Π_k^{feas} empty, the flow is refused, it is excluded from $\mathcal{K}_{\text{adm}}$, and it loads no relay. Each refusal is recorded against one of six mutually exclusive reasons: no reachable route, hop cap, AoSI cap, fidelity floor, per-flow energy cap, or relay service-availability cap. Consequently, the incumbent test compares assignments lexicographically, first on admitted-flow count and then on summed MF-SREE, which is the order of (17a).

4.2. Scheduled Routing Objective and S-ACS MF-SINR

The routing model of Section 4.1 prices contention solely by relay load count, omitting transmitter geometry from the optimization objective. To incorporate an explicit transmission schedule, we formulate a scheduled routing objective that evaluates complete path assignments rather than isolated paths. Specifically, it tests all active link pairs for joint spatial feasibility. It then colors the resulting conflict graph into TDMA slots via a deterministic degree-of-saturation (DSATUR) rule. Any slot whose aggregate interference leaves a member link below the required SINR threshold is split, which fixes the slot count S .

We now state the evaluator in full. Let $\mathcal{T}_s$ denote the set of nodes transmitting in slot s , let $g_{u,v}$ be the linear power gain of link (u, v) under the A2A path-loss model of Section 3.1, and let σ^2 be the receiver noise power. The scheduled quantities are then:

$$\Gamma_{u \rightarrow v}(\mathcal{T}_s) = \frac{P_{\text{tx}} g_{u,v}}{\sigma^2 + I_v(\mathcal{T}_s \setminus \{u\})}, \quad (26a)$$

$$I_v(\mathcal{X}) = \sum_{w \in \mathcal{X} \setminus \{v\}} P_{\text{tx}} g_{w,v}, \quad (26b)$$

$$\alpha = 1/S, \quad (26c)$$

$$R_{u \rightarrow v} = \alpha W \log_2(1 + \Gamma_{u \rightarrow v}), \quad (26d)$$

$$\gamma_{u \rightarrow v}^{\text{srv}} = 10 \log_{10}(2^{\max(R_{u \rightarrow v}, R_{\min})/W} - 1), \quad (26e)$$

$$E^{\text{srv}}(\mathcal{P}_k) = \sum_{h=1}^{H_k} E_{n_{h-1}^{(k)}, n_h^{(k)}}(\gamma_{u \rightarrow v}^{\text{srv}}). \quad (26f)$$

Here, (26a) is the spatial SINR at receiver v , and (26b) aggregates the co-slot interference reaching that receiver. The wanted transmitter u and the receiver itself are both

Algorithm 1 S-ACS MF.

Require: Graph $\mathcal{G}(t)$, K source-destination pairs $\{(s_k, d_k)\}_{k=1}^K$, service availability $C_{\max}$, per-flow budgets $S_k^{\text{req}}, E_{\max}^{(k)}$, params $\alpha_p, \beta_p, \rho_{\text{base}}, Q, M, I_{\max}, R_{\text{LS}}$

Ensure: Admitted set $\mathcal{K}_{\text{adm}}$, path assignment $\{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}$, and total MF-SREE

- 1: $\tau_{ij}^{(k)} \leftarrow \tau_0$ for all edges and all k ; $(K_{\text{best}}, \text{SREE}_{\text{best}}^\Sigma) \leftarrow (-\infty, -\infty)$
- 2: **for** iter = 1 to $I_{\max}$ **do**
- 3: $\pi \leftarrow$ random permutation of $\{1, \dots, K\}$
- 4: $L_j \leftarrow 0$ for all nodes; $\mathcal{K}_{\text{adm}} \leftarrow \emptyset$
- 5: **for** $k \in \pi$ **do**
- 6: Compute $\gamma_{ij}^{\text{eff},(k)}$ using **L** via (3)
- 7: Compute heuristic $\eta_{ij}^{(k)}$ via (20)
- 8: **for** $m = 1$ to M **do**
- 9: Construct $\mathcal{P}_m^{(k)}$ hop-by-hop via (19), drawing each next hop from $\mathcal{N}_{\text{feas}}(i, k)$ (18) re-evaluated at the current node i
- 10: **end for**
- 11: $\mathcal{P}_D^{(k)} \leftarrow$ Dijkstra under load-aware $w_{ij}^{(k)}$ (24)
- 12: $\Pi_k^{\text{feas}} \leftarrow \{\mathcal{P} \in \{\mathcal{P}_1^{(k)}, \dots, \mathcal{P}_M^{(k)}, \mathcal{P}_D^{(k)}\} : \mathcal{P} \text{ satisfies (17c)–(17e) and (17g)–(17h)}\}$
- 13: **if** $\Pi_k^{\text{feas}} = \emptyset$ **then**
- 14: Reject flow k : $\mathcal{P}_k^* \leftarrow \emptyset$, record reason NO-FEASIBLE-ROUTE, and **continue** (load **L** unchanged)
- 15: **end if**
- 16: $\mathcal{P}_k^* \leftarrow$ coupled-marginal selection over Π_k^{feas} via (25)
- 17: $\mathcal{K}_{\text{adm}} \leftarrow \mathcal{K}_{\text{adm}} \cup \{k\}$
- 18: $L_j \leftarrow L_j + \mathbf{1}[j \in V(\mathcal{P}_k^*) \setminus \{s_k, d_k\}]$
- 19: **end for**
- 20: **Local search:**
- 21: **for** $k \in \mathcal{K}_{\text{adm}}$ **do**
- 22: $\mathcal{P}'_k \leftarrow \arg \max_{\mathcal{P} \in \Pi_k^{\text{feas}}} \text{SREE}^\Sigma(\{\mathcal{P}_{k'}^*\}_{k' \neq k} \cup \{\mathcal{P}\})$ s.t. (17c), (17f), and (17g)–(17h)
- 23: **if** $\text{SREE}^\Sigma(\mathcal{P}'_k) > \text{SREE}^\Sigma(\mathcal{P}_k^*)$ **then**
- 24: $\mathcal{P}_k^* \leftarrow \mathcal{P}'_k$; refresh **L**
- 25: **end if**
- 26: **end for**
- 27: Evaluate $\text{SREE}^\Sigma \leftarrow \sum_{k \in \mathcal{K}_{\text{adm}}} S_{\text{e2e}}(\mathcal{P}_k^*, \mathbf{L}) / E_{\text{total}}(\mathcal{P}_k^*, \mathbf{L})$ via (8), (11) and (16), with **L** the joint load of the whole assignment
- 28: **if** $(|\mathcal{K}_{\text{adm}}|, \text{SREE}^\Sigma) >_{\text{lex}} (K_{\text{best}}, \text{SREE}_{\text{best}}^\Sigma)$ **then**
- 29: Update best assignment $\{\mathcal{P}_k^*\}$ and $\mathcal{K}_{\text{adm}}$; $(K_{\text{best}}, \text{SREE}_{\text{best}}^\Sigma) \leftarrow (|\mathcal{K}_{\text{adm}}|, \text{SREE}^\Sigma)$
- 30: **end if**
- 31: Evaporate via (22) at the joint rate (23); deposit via (21) for the iteration-best ant of each flow
- 32: Clip $\tau_{ij}^{(k)} \in [\tau_{\min}, \tau_{\max}]$
- 33: **end for**
- 34: $\{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}, \text{SREE}_{\text{best}}^\Sigma \leftarrow$ reroute refinement via Algorithm 2
- 35: **return** $\mathcal{K}_{\text{adm}}, \{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}, \text{SREE}_{\text{best}}^\Sigma$

excluded from the interfering set, since one radio per node cannot transmit while receiving. Next, (26c) gives the equal airtime share of a link holding one of S slots, and (26d) derates its Shannon rate by that share. Then, (26e) inverts Shannon capacity on the derated rate to recover an effective service SNR in dB, and (26f) evaluates the per-hop energy model (10) at that service SNR. The floor $R_{\min} = 1$ bit/s in (26e) keeps the inversion finite when a link receives negligible airtime. A flow is admitted under this objective only if every hop clears the SINR threshold γ_{th} (dB):

$$\min_{1 \leq h \leq H_k} 10 \log_{10} \Gamma_{n_{h-1}^{(k)} \rightarrow n_h^{(k)}} \geq \gamma_{\text{th}}. \quad (27)$$

A flow failing (27) contributes zero to summed SREE. Notably, path fidelity is composed across hops from the raw

Algorithm 2 Coordinated Reroute Refinement.

Require: Graph $\mathcal{G}(t)$, incumbent $\{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}$, admitted set $\mathcal{K}_{\text{adm}}$, sweeps R_{LS}

Ensure: Refined assignment $\{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}$, total MF-SREE $\text{SREE}_{\text{best}}^\Sigma$

- 1: **for** $r = 1$ to R_{LS} **do**
- 2: **for** $k \in \mathcal{K}_{\text{adm}}$ in increasing order of $\text{SREE}(\mathcal{P}_k^*; \mathbf{L})$ **do**
- 3: $\tilde{\mathcal{P}}_k \leftarrow$ Dijkstra under $w_{ij}^{(k)}$ (24) on residual load of other $K-1$ flows, with relays at $C_{\max}$ masked out
- 4: **if** $\tilde{\mathcal{P}}_k$ satisfies (17c)–(17h) and strictly increases SREE^Σ **then**
- 5: $\mathcal{P}_k^* \leftarrow \tilde{\mathcal{P}}_k$; refresh **L**
- 6: **end if**
- 7: **end for**
- 8: **end for**
- 9: $\text{SREE}_{\text{best}}^\Sigma \leftarrow \sum_{k \in \mathcal{K}_{\text{adm}}} \text{SREE}(\mathcal{P}_k^*; \mathbf{L})$
- 10: **return** $\{\mathcal{P}_k^*\}_{k \in \mathcal{K}_{\text{adm}}}, \text{SREE}_{\text{best}}^\Sigma$

per-hop SINRs of (26a), not from the service SNR of (26e), because the service SNR is an energy-accounting construct rather than a channel observation. Accordingly, assignments are ordered lexicographically by admitted-flow count and subsequently by summed per-flow SREE (16), matching the two keys of (17a).

To optimize this objective, S-ACS MF-SINR adapts Algorithm 1 by redirecting every objective-consulting decision to the scheduled score. Coupled-marginal selection (25) ranks each candidate by the scheduled score of the committed set plus that candidate, penalizing forced slots and inter-flow interference. The construction heuristic (20) replaces the load-count penalty in (3) with the spatial link quality against committed transmitters and the matching airtime-derated service energy. Furthermore, joint local search accepts swaps only when the assignment score strictly improves. Coordinated rerouting performs monotone ascent on that score. Pheromone deposit scales with the scheduled per-flow SREE. Conversely, the feasible-neighbor set (18), the hop and per-node energy budgets, the service-availability cap (17f), the MMAS bounds, and the congestion-dependent evaporation remain unchanged. This ensures that the variant differs from S-ACS MF solely in its objective.

Notably, only the routing objective is scheduled. Both the admission boundary scoring all methods and the candidate-quality filter retain the load-count contention surrogate. The scheduled objective is evaluated alongside the retained surrogate rather than replacing it. Computationally, pairwise feasibility testing incurs $\mathcal{O}(|\mathcal{E}|^2 N)$ complexity per evaluation, where $|\mathcal{E}|$ and N denote active links and nodes. DSATUR coloring adds an additional factor. Precomputed link-gain matrices and assignment memoization maintain tractable runtime. Finally, the service-normalized model of (26e)–(26f) is distinct from on-air per-task accounting. Both remain strictly within the communication-layer boundary.

4.3. Analytical Bounds and Gap Analysis

To quantify how far the heuristic sits from an analytical reference, in this subsection, we develop two complementary bounds with distinct roles. The first is a per-flow zero-load ceiling, serving as a valid global upper limit on the true

MF-SREE (17). Its per-flow terms are computed exactly and carry an optimality certificate on every solve. The second is a fixed-load LP relaxation operating on a linearized log-fidelity-minus-energy surrogate. It separates the shortfall of the heuristic into an integrality component and an algorithmic component. The LP bounds that surrogate and not the MF-SREE itself. Thus, the global upper bound used throughout the paper remains the per-flow zero-load ceiling (28). The LP serves strictly as an offline diagnostic. Neither bound certifies the coupled multi-flow optimum. We therefore reserve the term optimality gap for instances whose exact optimum is available. The distance to the ceiling is reported elsewhere as a zero-load-ceiling gap.

4.3.1. Per-Flow Dinkelbach Upper Bound

The MF-SREE objective (17a) is a sum of K per-flow ratios $\sum_k S_{e2e}(\mathcal{P}_k)/E_{\text{total}}(\mathcal{P}_k)$. Because the sum-of-ratios structure precludes direct application of the classical Dinkelbach fractional iteration [12], we construct an upper bound by theoretically decoupling the flows. Specifically, each per-flow SREE is bounded independently at a zero-load ceiling, and the resulting bounds are summed across flows.

Proposition 1 (Per-Flow SREE Upper Bound). *For each flow $k \in \{1, \dots, K\}$, let $\text{SREE}_k^{(0)} = \max_{\mathcal{P} \in \Pi(s_k, d_k)} S_{e2e}^{(0)}(\mathcal{P})/E_{\text{total}}^{(0)}(\mathcal{P})$ denote the optimal single-flow SREE at zero relay load ($L_j = 0$ for all j), where superscript (0) indicates evaluation in the absence of relay load. Let MF-SREE* denote the second key of the lexicographic optimum (17a), that is, the summed per-flow SREE over the admitted set of an optimal assignment. Then:*

$$\text{MF-SREE}^* \leq \sum_{k=1}^K \text{SREE}_k^{(0)}. \quad (28)$$

The bound holds whether or not flows are refused, since a refused flow contributes zero to the left-hand side while its zero-load term is retained on the right. Each $\text{SREE}_k^{(0)}$ is computed exactly, without candidate-set truncation, by a Pareto label-setting dynamic program that returns an optimality certificate with every solve (Appendix A.5, Part 3).

The validity of this bound stems from the monotonicity of (3): increasing relay load L_j strictly decreases effective SNR, degrading fidelity and increasing energy drain (proof in Appendix A.5). The dynamic program certifies each per-flow zero-load subproblem, providing a valid upper bound on the efficiency key of MF-SREE. That scope is worth stating plainly: the bound constrains summed efficiency alone and certifies nothing about the admitted-flow count, so it is a ceiling on the second key rather than on the lexicographic pair. At large K , the relaxation gap between $\sum_k \text{SREE}_k^{(0)}$ and the coupled optimum widens as $K(K-1)$ contention penalties are discarded. Exact optimality certification is therefore verified on small-scale benchmark instances ($N \leq 10$, $K \leq 3$) via full enumeration. Comparisons at larger scales evaluate distance to $\sum_k \text{SREE}_k^{(0)}$ as a zero-load-ceiling gap.

4.3.2. Fixed-Load Multi-Commodity LP Relaxation

To obtain a tighter diagnostic bound, a fixed-load multi-commodity LP relaxation is formulated on a linearized surrogate of the MF-SREE objective. We apply two linearization steps to obtain a linear program. First, a log-fidelity transformation converts the multiplicative fidelity into additive per-edge contributions. Second, a conditioning step evaluates the formulation on the observed load profile. Let $\pi = (\pi_1, \dots, \pi_K)$ denote a fixed flow ordering.

Definition 2 (Fixed-Load Multi-Commodity Linearized Coefficients). *Under a fixed load profile $\hat{\mathbf{L}} = (\hat{L}_1, \dots, \hat{L}_N)$ and effective SNR $\gamma_{ij}(\hat{L}_j) = \gamma_{ij}^{(0)} - 10 \log_{10}(1 + \hat{L}_j)$, the edge linear coefficient $c_{ij}(\hat{L}_j)$ is:*

$$c_{ij}(\hat{L}_j) = \ln S(\gamma_{ij}(\hat{L}_j)) - \lambda E_{ij}^{(\hat{L}_j)}, \quad (29)$$

where $\lambda > 0$ is a normalization parameter set to $\lambda = 1/\bar{E}$, with $\bar{E} = (1/K) \sum_k E_{\text{total}}(\mathcal{P}_k^{\text{ACO}})$ representing the average per-flow energy of the heuristic solution. The choice $\lambda = 1/\bar{E}$ ensures that the log-fidelity and energy terms operate at comparable magnitudes.

The energy term takes the linear λE_{ij} form rather than the logarithmic $\ln E_{ij}$ form used in the heuristic (20). Specifically, $\sum_{(i,j) \in \mathcal{P}} E_{ij}$ mathematically matches the physical total path energy constrained by (17c). In contrast, $\sum_{(i,j) \in \mathcal{P}} \ln E_{ij}$ would optimize a log-product without operational meaning. The fidelity term remains log-domain since end-to-end fidelity is inherently multiplicative.

Definition 3 (Fixed-Load Multi-Commodity LP Relaxation). *For a flow ordering π and fixed load profile $\hat{\mathbf{L}}$, the LP relaxation is formulated as:*

$$Z_{\text{LP}}(\hat{\mathbf{L}}) = \max_{\mathbf{x} \in [0,1]^{|\mathcal{E}| \times K}} \sum_{k=1}^K \sum_{(i,j) \in \mathcal{E}} x_{ij}^{\pi_k} c_{ij}(\hat{L}_j) \quad (30a)$$

$$\text{s.t.} \quad \sum_{j:(i,j) \in \mathcal{E}} x_{ij}^{\pi_k} - \sum_{j:(j,i) \in \mathcal{E}} x_{ji}^{\pi_k} = b_i^{\pi_k}, \quad \forall i, k, \quad (30b)$$

$$\sum_{k'=1}^K r_j^{\pi_{k'}} \leq C_{\text{max}}, \quad \forall j \in \mathcal{V}, \quad (30c)$$

$$0 \leq x_{ij}^{\pi_k} \leq 1, \quad \forall (i,j) \in \mathcal{E}, k, \quad (30d)$$

where $x_{ij}^{\pi_k} \in [0,1]$ is the fractional routing variable for flow π_k on edge (i,j) , and $b_i^{\pi_k} \in \{+1, 0, -1\}$ is the flow-conservation demand (+1 at source, -1 at destination, 0 otherwise). The fractional relay usage of flow π_k at node j is $r_j^{\pi_k} \triangleq \sum_{i:(i,j) \in \mathcal{E}} x_{ij}^{\pi_k} \cdot \mathbf{1}[j \neq s_{\pi_k}] \cdot \mathbf{1}[j \neq d_{\pi_k}]$, consistent with (2). Here, $\mathcal{D} = \{s_k, d_k : k = 1, \dots, K\}$ is the set of all source-destination nodes. The program incorporates flow conservation (30b), service availability (30c), and box constraints (30d). Conditioned on $\hat{\mathbf{L}}$, the LP serves as an offline diagnostic at the heuristic operating point.

Constraint (30c) mirrors (2) by excluding flow endpoints while bounding interior relay usage across all $j \in \mathcal{V}$. Because $x_{ij}^{\pi_k} \in [0, 1]$ relaxes the integrality requirement, the optimal value Z_{LP} rigorously upper-bounds the surrogate objective for all integer-feasible routings.

4.3.3. Surrogate Gap Decomposition

The LP relaxation (30) yields an upper bound on the surrogate objective Z , not on MF-SREE. That upper bound decomposes into two interpretable components. Together, these components characterize the distance between the LP optimum and the heuristic value under the same surrogate.

Theorem 4 (Gap Decomposition). *Fix the commodity set $\mathcal{K}_{\text{adm}}^{\text{ACO}}$ admitted by S-ACS MF, and pose (30) on that same set. Let $Z_{LP}(\hat{\mathbf{L}})$ denote the optimal value of the resulting fixed-load LP under the load profile $\hat{\mathbf{L}}$, and let Z_{ACO} denote the objective evaluated at the integer S-ACS MF solution. The surrogate gap decomposes as:*

$$\underbrace{Z_{LP}(\hat{\mathbf{L}}) - Z_{\text{ACO}}}_{\text{total gap}} = \underbrace{Z_{LP}(\hat{\mathbf{L}}) - Z_{IP}(\hat{\mathbf{L}})}_{\text{integrality gap}} + \underbrace{Z_{IP}(\hat{\mathbf{L}}) - Z_{\text{ACO}}}_{\text{heuristic gap}} \quad (31)$$

where $Z_{IP}(\hat{\mathbf{L}})$ is the optimal value of the integer version of (30) (enforcing $x_{ij}^{\pi_k} \in \{0, 1\}$). Both components are strictly non-negative.

The proof (Appendix A.6) follows from the continuous relaxation ($Z_{LP} \geq Z_{IP}$) and the fact that S-ACS MF produces a valid integer solution ($Z_{IP} \geq Z_{\text{ACO}}$). The common-admitted-set stipulation is what makes the second step available. A refused flow carries no path and therefore no integer point for its flow-conservation rows, so a decomposition posed over all K commodities would compare programs whose commodity sets differ. Where every flow is admitted, the stipulation is vacuous and the theorem reads as before. The integrality gap measures the structural penalty inherent to combinatorial routing. In contrast, the heuristic gap isolates the suboptimality of path selection relative to the best integer routing under the fixed-load model.

Remark 2 (Practical Significance). *When the integrality gap dominates, the LP relaxation is loose (exhibiting many fractional splits at high K). When the heuristic gap dominates, S-ACS MF has room for path-selection improvement. For the default non-binding capacity ($C_{\max} = N$), the integrality gap is empirically negligible because the multi-flow structure admits near-integer LP solutions. Since both $\hat{\mathbf{L}}$ and λ are derived from the S-ACS MF solution, this gap measures the surrogate integrality gap strictly at the heuristic's operating point.*

We further characterize the integrality gap by examining the single-commodity boundary case:

Proposition 5 (Single-Commodity LP Exactness). *For $K = 1$ and integer service availability $C_{\max} \in \mathbb{Z}_+$, the fixed-load LP relaxation (30) possesses an integer optimum. Thus, the integrality gap vanishes entirely: $Z_{LP}(\hat{\mathbf{L}}) = Z_{IP}(\hat{\mathbf{L}})$.*

The proof (Appendix A.7) applies a standard node-splitting reduction [3] to convert node constraints (30c) into edge capacities. This exposes a totally unimodular node-edge incidence matrix [23]. Although multi-flow coupling ($K \geq 2$) destroys total unimodularity [15], the measurements of Section 5.7 indicate that the gap remains small when capacity is non-binding and K is moderate.

4.4. Convergence and Complexity

We conclude by establishing the formal convergence properties and computational complexity of our proposed algorithm. By strictly anchoring pheromone bounds and evaporation dynamics to the MMAS framework, Algorithm 1 converges in probability to the conditional per-flow optimum under a static load profile. That guarantee does not extend to the load-coupled multi-flow process, as Corollary 7 makes explicit.

Theorem 6 (S-ACS MF Per-Flow Convergence). *Consider the S-ACS inner loop applied to a single flow k with MMAS-style pheromone bounds $0 < \tau_{\min} \leq \tau_{\max} < \infty$, strictly positive heuristics $\eta_{ij}^{(k)} > 0$, and $M \geq 1$ ants. Conditioned on a frozen relay load vector $\mathbf{L}$ (i.e., assuming the loads imposed by preceding flows $1, \dots, k-1$ are constant), the best-so-far path $\mathcal{P}_k^{\text{bsf}}$ converges in probability to the conditional optimum $\mathcal{P}_k^{\text{opt}} \in \arg \max_{\mathcal{P} \in \Pi(s_k, d_k; \mathbf{L})} \text{SREE}(\mathcal{P}; \mathbf{L})$. Specifically, for any $\delta \in (0, 1)$, there exists an iteration threshold $I_0 \in \mathbb{N}$ ensuring:*

$$\Pr[\text{SREE}(\mathcal{P}_k^{\text{bsf}}; \mathbf{L}) < \text{SREE}(\mathcal{P}_k^{\text{opt}}; \mathbf{L})] < \delta, \quad (32)$$

for all $I_{\max} \geq I_0$. This guarantee is per-flow and strictly conditioned on $\mathbf{L}$.

The proof adapts the MMAS convergence result of [39] and is given in Appendix A.4.

Corollary 7 (Non-Establishment of Joint Convergence). *The convergence guarantee in Theorem 6 is strictly per-flow and conditioned on $\mathbf{L}$. A formal guarantee of joint K -flow convergence to the global MF-SREE optimum is not established, as the sequential routing order induces interdependent path optimization. The theorem is stated for one commodity's own pheromone matrix. The K matrices are coupled only through the joint evaporation rate and the load-dependent heuristic, neither of which the MMAS analysis covers. Evidence for the coupled process is therefore empirical, and it is bounded by what can be verified against an exact optimum.*

We now analyze the computational complexity of Algorithm 1. The total computational complexity evaluates to:

$$\mathcal{O}(I_{\max} \cdot [K(MH_{\max}N + N^2 \log N) + K^2MH_{\max}] + R_{\text{LS}} \cdot K \cdot (N^2 \log N + KH_{\max})). \quad (33)$$

In (33), the first bracketed term accounts for hop-by-hop M -ant construction (Algorithm 1, Line 9). The second represents the load-aware Dijkstra candidate step (Line 11),

executed $K \cdot I_{\max}$ times. The third arises jointly from local search (Line 21) and coupled-marginal selection (Line 16), each contributing $\mathcal{O}(K^2 M H_{\max})$ by re-scoring candidates against perturbed relays. The outer term represents coordinated reroute refinement (Algorithm 2, Lines 1–8), running for at most $R_{LS} = 10$ sweeps (with two to four sweeps observed, incurring $\mathcal{O}(N^2 \log N)$ per flow). Memory footprint is $\mathcal{O}(K N^2 + K M H_{\max})$, storing K pheromone matrices ($\mathcal{O}(K N^2)$), network channel and energy matrices ($\mathcal{O}(N^2)$), and candidate paths ($\mathcal{O}(K M H_{\max})$). To counter super-linear complexity in K , we contract $(M, I_{\max})$ from (10, 30) at $K \leq 5$ through (8, 20) at $5 < K \leq 20$ to (8, 15) at $K > 20$. In Table 5 ($N = 50$), matched wall-clock time scales $3.6\times$ from 0.368 ± 0.021 s ($K = 10$) to 1.309 ± 0.031 s ($K = 50$), containing complexity growth across a fivefold load increase. Under periodic topology re-computation in dynamic UAV routing [41, 11], receding-horizon validity requires $T_{\text{solve}} \leq \eta \Delta T_{\text{topo}} = 1.0$ s ($\Delta T_{\text{topo}} = 2.0$ s, $\eta = 0.5$). Matched timings satisfy this deadline at $K = 10$ and $K = 30$, while exceeding it at $K = 50$ under shared-server ambient load. Thus, the margin against 1.0 s provides an operational tuning constraint rather than an unconditional guarantee.

5. Performance Evaluation

5.1. Simulation Setup

We evaluate the proposed S-ACS MF algorithm by numerical simulation under the parameters of Table 3, where path loss repurposes the 3GPP specification [1] as an A2A surrogate per Section 3.1. For S-ACS MF, the ant budget is tiered across $K \leq 5$, $5 < K \leq 20$, and $K > 20$, the local search module is enabled, and the transition heuristic adopts (20).

All paired comparisons are evaluated on identical topology seeds, with 95% confidence intervals and multiplicity controlled by the Holm–Bonferroni correction [24]. Statistical power analysis and equivalence testing are used to evaluate non-detectable performance differences [10]. The full protocol and the effect-size summary appear in Supplementary Material Section S-D.7 (Table S-IV).

During path construction, the link availability (17b), hop-limit (17e), and service-availability (17f) constraints are enforced dynamically by the feasible-neighbor filtering step that forms $\mathcal{N}_{\text{feas}}(i, k)$ in Algorithm 1. That step also applies the prospective per-edge residual-energy test $E_j^{\text{node}} + E_{ij}^{(k)} \leq E_j^{\text{rem}}$. Completed candidates are then verified against the aggregate energy budget (17c) in the feasibility-filter step that forms Π_k^{feas} . The end-to-end AoSI constraint (17d) is evaluated post-construction. The two per-flow QoS budgets are enforced as membership conditions of Π_k^{feas} , in both refinement stages, and again by the common final admission-and-scoring stage used for every routed method under comparison. That stage also enforces connectivity, hop, AoSI, relay-capacity, and per-node battery limits. Thus, no route violating a hard constraint is counted as delivered. Throughout this section the two per-flow budgets are nevertheless held at their inactive defaults, namely $E_{\max}^{(k)} = \infty$ for

the route-energy cap (17h) and $S_k^{\text{req}} = 0$ for the fidelity floor (17g). Every reported route satisfies both vacuously. Consequently, no result below is evidence about binding per-flow energy or fidelity requirements. The single exception is Section 5.5, where both are swept to measured binding levels under a measured DeepSC codec model and reported in Table 10. The per-node battery budget (17c) is likewise non-binding at the default $B_{\text{bat}} = 1.08 \times 10^6$ J, which exceeds the largest observed per-epoch drain by several decimal orders.

To benchmark the proposed scheme, we compare S-ACS MF against four principal routing baselines. Sequential Dijkstra is a greedy shortest-path baseline using surrogate weights $E_{ij}/S(\gamma_{ij})$ with sequential load updates after each flow. Traditional ACO MF [13] is an independent per-flow ACO algorithm without load-count coupling during path construction. Random is an independent per-flow random walk, and Greedy Walk MF is a parameter-free greedy multi-flow walk that decides locally without pheromone memory.

Three learning-based entries are carried alongside these, not as principal baselines but as untuned diagnostic implementations. MAPPO MF [45] is a multi-agent proximal policy optimization (MAPPO) scheme trained across K agents under the centralized training with decentralized execution (CTDE) paradigm. QMIX MF [37] is a value-decomposition multi-agent reinforcement learning (MARL) approach trained under the same paradigm. GNN MF [21] is a GraphSAGE-based graph neural network combined with proximal policy optimization and pre-trained in the single-flow setting. Their architectures, training budgets, and measured outcomes are reported in Supplementary Material Section S-C. No conclusion in this section rests on them. Segregated computational profiling is provided in Supplementary Material Section S-B (Table S-I).

The simulation environment is executed in Python 3.11 on an Intel Core i7-14700 processor with 64 GB RAM. Wall-clock measurements are reported from a separate environment, because they are the only timing-sensitive results here. Every runtime number in this paper was measured on a 128-core Linux server under Python 3.12.3 with per-process thread counts pinned to one. Two campaigns ran there and are reported separately. The matched comparison of Section 5.2 ran with interleaved trials at a measured background load of 14.8%. The deployed re-solve latencies quoted in Section 5.6 ran serially on the same server while otherwise idle. Routing is computed by a centralized controller with global topology knowledge via periodic beaconing, whereas the MARL entries operate under the CTDE paradigm. Control signaling is accounted by (13) at the per-node epoch scope, giving $E^{\text{sig}} = 1.19 \times 10^{-4}$ J per epoch network-wide at $N=50$, requested $K=50$, and $K_{\text{adm}}=49$. Consistent with the energy-accounting boundary of Section 3.3, it enters the battery constraint (17c) and is excluded from the per-flow SREE ratio (16).

5.2. Multi-Flow Performance

Table 4 reports the principal grid, namely the flow-count sweep $K \in \{2, \dots, 50\}$ at $N=50$ and the fleet-size sweep

Table 3

Simulation parameters across network and algorithm configurations.

Parameter	Value	Parameter	Value
<i>Experiment</i>		<i>Channel</i>	
N (UAVs)	10, 25, 50, 75, 100	f_c (GHz)	2
K (flows)	2–50	B (MHz)	10
Trial	20 / 20	P_{tx} (dBm)	23
($N=50/100$)		Path loss	A2A surrogate
		Noise figure (dB)	9
<i>Network</i>		<i>Semantic Fidelity (synthetic)</i>	
Airspace (m)	1000×1000×300	S_{max}, S_{min}	0.95, 0.05
Altitude (m)	[50, 300]	$\gamma_{0.5}$ (dB)	10
Mobility	G–M [32], $\alpha_m=0.5, \bar{v}=10$	k_s (dB ⁻¹)	0.3
d_{max} (m)	500	H_{max}	20
$AoSI_{max}$ (s)	5, 0.1	C_{max}	N
λ_{AoSI}			
<i>Energy</i>		<i>QoS & Control</i>	
n_b (bits)	10 ⁶	ϵ_{comp} (J/FLOP)	10 ⁻¹¹
P_{proc} (W)	5	T_{proc} (ms)	1
B_{bat}	300 Wh (1.08×10 ⁶ J)	F (FLOPs)	3×10 ¹⁰
$S_k^{req}, E_{max}^{(k)}$	0, ∞	T_{epoch}, T_{mis} (s)	1.0, 1800
b_{bc}, b_{fe}	256, 32 bits	P_{prop} (W)	0–500
<i>S-ACS MF Baseline Schedule</i>			
Ants M	10 / 8 / 8	α_p, β_p	1.0, 2.0
ρ_{base}, ζ	0.1, 0.3	Q	1.0
I_{max}	30 / 20 / 15	R_{LS}	10

$N \in \{10, 25, 50, 75, 100\}$ at $K=20$. Every contrast in it is obtained under three stated conditions: the load-count contention surrogate of Section 3.1, the semantic-fidelity curve calibrated to the measured deep semantic codec of Section 3.3, and both per-flow QoS budgets inactive. Each entry is a mean paired relative difference over 20 paired topology seeds. Multiplicity is controlled within this grid's own family of 40 tests, of which 26 reject at $\alpha=0.05$. The margins are reported per rival, because they differ by more than an order of magnitude. Against Sequential Dijkstra, the mean paired margin is positive at all twelve design points but small, spanning +0.024% to +0.648%. That margin survives Holm–Bonferroni correction at 5 of them. The remaining seven are positive in sign yet not separable from zero once corrected. Against Traditional ACO MF the margin spans +1.04% to +4.40% and survives correction at 9 of the 12 points. Against the parameter-free Greedy Walk MF it spans +8.84% to +46.43% and survives correction at every one of them. Accordingly, S-ACS MF attains the highest MF-SREE in 11 of 12 settings. The single exception is $K=5$ at $N=50$, where LP-MCF attains the higher mean, though the two are statistically indistinguishable ($p_{adj}=0.9838$). LP-MCF is an LP relaxation with randomized rounding, evaluated as a routing baseline. It solves the fixed-load linear relaxation, rounds that solution to integral paths, and repairs those paths against the relay-capacity constraint. Its MF-SREE is therefore measured exactly as every other baseline's is, and

it is not a bound on the objective. Its gap is nonetheless reported separately in the Supplementary Material census rather than beside the routing margins, because the LP is solved at only four of the twelve design points. Those four are $K \in \{2, 5, 10\}$ at $N=50$ and $K=20$ at $N=10$. Across them the paired gap spans 0.003% to 0.294%. The gap is nominally positive for S-ACS MF at each of the four and separates from zero at none after correction. Meanwhile, the strict win count over the 20 paired topologies grows with load, from 2 at $K=2$ to 20 at $K \geq 30$, although not monotonically. Supplementary Material Table S-XIII reports the complete paired-comparison census over all twelve design points.

Under the architectures, state–action representations, and training budgets listed in Supplementary Material Section S-C, neither retrained policy is statistically distinguishable from the parameter-free Greedy Walk MF floor. The pre-trained graph policy falls below it. That outcome reflects those specific untuned configurations rather than the corresponding method classes. The measurements, the attained statistical power, and the diagnostics not performed are reported there.

Table 5 reports the objective- and budget-matched wall-clock comparison at $N=50$, decomposed into route construction, objective evaluation, local search, and total wall clock over 20 paired topology seeds. Under this matched arm the ratio of Traditional ACO MF to S-ACS MF total wall-clock time stays below unity at $K=10, 30$, and 50 under both ant budgets of Table 3. Those six ratios are 0.942, 0.825, and 0.776 at the literal budget, and 0.696, 0.617, and 0.601 at the reduced one. Once implementation and work budget are matched, S-ACS MF is 1.44×–1.66× slower than Traditional ACO MF at the reduced budget and 1.06×–1.29× slower at the literal budget. It is therefore not faster than the baseline (Supplementary Material Fig. S-10). One qualification belongs with that statement. Under the paired estimator over the same 20 topology seeds, five of these six cells separate from parity. The literal budget at $K=10$ does not. Its paired ratio is 0.941 with a 95% confidence half-width of 0.133. The two methods are thus statistically indistinguishable there, and no ordering is claimed for that cell. When compared as originally configured, the large wall-clock margin recomputes to 14.0×–16.4× over these three flow counts. This margin decomposes into two confounds, neither of which is algorithmic. The first is an unequal ant budget, worth 5.7×–8.2×, because Traditional ACO MF ran at the literal schedule while S-ACS MF ran at the reduced one. The second is the compiled evaluator applied only to the S-ACS MF hot loop, worth a further 3.3×–3.4× against a compilation control of approximately 1× for the baseline. Their product alone reproduces a ratio of 19.2×–27.1×. The earlier figure therefore describes a configuration pairing rather than a property of either algorithm. Supplementary Material Table S-IX gives the full decomposition. The four-way split locates that time in the coupled multi-flow objective evaluation and the coordinated refinement stage, neither of which the per-flow baseline performs. The matched total grows 3.6× as the flow count grows 5×. This containment

Table 4

Paired margins of S-ACS MF over rival routing baselines under the fitted semantic-fidelity model (mean $\pm$ 95% CI; * marks statistical significance at $\alpha = 0.05$ after Holm–Bonferroni adjustment; bold marks best method).

K	N	Best method	vs. Seq. Dijkstra (%)	vs. Trad. ACO MF (%)	vs. Greedy Walk MF (%)
<i>K-sweep, N=50</i>					
2	50	S-ACS MF	+0.024 $\pm$ 0.041	+2.034 $\pm$ 2.619	+19.213 $\pm$ 12.814*
5	50	LP-MCF	+0.031 $\pm$ 0.027	+2.521 $\pm$ 2.086	+20.204 $\pm$ 15.972*
10	50	S-ACS MF	+0.165 $\pm$ 0.179	+3.601 $\pm$ 1.497*	+26.036 $\pm$ 12.614*
15	50	S-ACS MF	+0.219 $\pm$ 0.120*	+3.542 $\pm$ 1.138*	+30.032 $\pm$ 12.113*
20	50	S-ACS MF	+0.446 $\pm$ 0.373	+3.586 $\pm$ 1.024*	+31.593 $\pm$ 9.356*
30	50	S-ACS MF	+0.626 $\pm$ 0.310*	+3.567 $\pm$ 0.709*	+37.525 $\pm$ 8.666*
40	50	S-ACS MF	+0.644 $\pm$ 0.270*	+4.405 $\pm$ 0.828*	+42.792 $\pm$ 8.445*
50	50	S-ACS MF	+0.648 $\pm$ 0.220*	+4.397 $\pm$ 0.931*	+46.429 $\pm$ 9.961*
<i>N-sweep, K=20</i>					
20	10	S-ACS MF	+0.333 $\pm$ 0.370	+1.041 $\pm$ 1.013	+8.838 $\pm$ 4.918*
20	25	S-ACS MF	+0.447 $\pm$ 0.219*	+3.303 $\pm$ 0.966*	+29.787 $\pm$ 7.764*
20	75	S-ACS MF	+0.162 $\pm$ 0.163	+4.331 $\pm$ 1.163*	+32.353 $\pm$ 7.048*
20	100	S-ACS MF	+0.102 $\pm$ 0.089	+3.960 $\pm$ 1.119*	+45.993 $\pm$ 10.082*

Table 5

Computational runtime breakdown and comparison at $N=50$ (mean $\pm$ 95% CI; RC: route construction, OE: objective evaluation, LS: local search; bold marks faster runtime).

K	S-ACS MF (s)				Trad. ACO MF (s)				Ratio		Seq. Dijkstra	Greedy Walk MF	LP-MCF	Random	Fixed-Load LP Bound
	RC	OE	LS	Total	RC	OE	LS	Total	Literal	Reduced					
10	0.095	0.155	0.088	0.368	0.251	0.002	0.000	0.256	0.942	0.696	0.0018	0.0021	0.236	0.0014	0.313
	± 0.006	± 0.009	± 0.006	± 0.021	± 0.034	± 0.000	± 0.000	± 0.034			± 0.0005	± 0.0003	± 0.038	± 0.0002	± 0.061
30	0.208	0.346	0.238	0.845	0.512	0.005	0.000	0.522	0.825	0.617	0.0040	0.0058	0.669	0.0037	0.745
	± 0.009	± 0.020	± 0.014	± 0.041	± 0.051	± 0.001	± 0.000	± 0.052			± 0.0002	± 0.0005	± 0.089	± 0.0005	± 0.110
50	0.314	0.523	0.396	1.309	0.772	0.008	0.000	0.787	0.776	0.601	0.0071	0.0092	1.14	0.0058	1.27
	± 0.010	± 0.012	± 0.012	± 0.031	± 0.056	± 0.000	± 0.000	± 0.056			± 0.0003	± 0.0008	± 0.08	± 0.0004	± 0.08

occurs because the tiered ant schedule offsets the super-linear $\mathcal{O}(K^2 M H_{\max})$ local-search term. The five rightmost columns of Table 5 retain the earlier single-run timings of the baselines not re-measured under the matched configuration. They are therefore not comparable with the two matched column groups. Among them, Sequential Dijkstra and Greedy Walk MF remain millisecond-scale heuristics under 10 ms at every K . As an absolute execution latency rather than a cross-method speedup, the measured re-solve latency of the deployed engine at $N=50$ under the reduced schedule rises from 0.196 s at $K=5$ to 0.888 s at $K=50$. That latency stays below the decision deadline $\eta \Delta T_{\text{topo}} = 1.0$ s. That margin belongs to the deployed configuration alone. Specifically, it reflects the tiered ant schedule and the runtime environment of that measurement. Those latencies moreover predate the compiled evaluator kernel. They are retained deliberately, because they exceed the re-measured compiled solve times and therefore bound the staleness analysis of Section 5.6 conservatively. The matched arm, which disables the compiled evaluator for both methods, instead totals 1.309 s at $K=50$ and exceeds the same deadline. That comparison is an absolute one, so it inherits the background load disclosed below. The evaluated configuration meets the operational deadline. This result is not an unconditional real-time guarantee.

To eliminate cross-machine variability, execution times were evaluated on a dedicated server under identical background conditions (measured load average of 14.8% across

128 cores) with interleaved trials. That server was not idle for any part of the campaign. The absolute wall-clock times of this arm therefore carry the measured contention and are not idle-machine times. Interleaving distributes that contention evenly over the two methods across the campaign, so it removes the temporal confounding that would arise if one method had been timed during a busier interval than the other. It does not establish that the two methods respond identically to contention. Any residual differential sensitivity to resource contention is therefore carried by the matched contrast and by the four-way split, and neither is claimed to be free of it. In the matched arm, the compiled fast path is disabled for both methods to isolate pure algorithmic scaling. For deployed operation, the compiled evaluator is verified via automated parity tests to yield identical routing paths and MF-SREE values, accelerating execution without altering solution quality.

The fleet-size sweep of the same table, at $K=20$, shows the margin over Sequential Dijkstra peaking at intermediate density. That margin is +0.33%, +0.45%, +0.45%, +0.16%, and +0.10% for $N=10$ through $N=100$. Of those five contrasts, only the one at $N=25$ survives Holm–Bonferroni correction. Beyond $N=25$, densified relay options narrow coordination gains. Meanwhile, the margin over Traditional ACO MF rises from +1.04% at $N=10$ and then plateaus between +3.30% and +4.33%, surviving correction at four of the five fleet sizes. The paired advantage over that baseline

is thus retained across the tested range rather than amplified by it.

Repeating this fleet-size sweep under the synthetic fidelity profile of Table 3 enlarges these two contrasts. Specifically, the factor is 6.8–17.3 against Sequential Dijkstra and 5.5–8.1 against Traditional ACO MF. The enlargement is rival-dependent, however. It reaches only 1.3–1.5 $\times$ against Greedy Walk MF. At $K=40$ and $K=50$ it inverts. There, the calibrated margin exceeds the synthetic one. Supplementary Material Section S-D.9 gives that stress arm's factor at every design point and rival. The shift is a property of the profile rather than of the routing. The synthetic profile places the tested link margins on the steep part of its fidelity curve. In contrast, the calibrated curve places them in its saturated region.


5.3. Component Attribution

Table 6 reports the framework against competing methods and against removals of its own mechanisms, on identical topology seeds. All entries are measured under the load-count contention surrogate and the synthetic stress profile. No ablation variant was repeated under the measured codec curve. At the default operating point ($C_{\max}=N$, refinement cap $R_{LS}=10$), Traditional ACO MF forfeits 12.8%–20.1% of the MF-SREE that the full method attains ($d_z = 1.35$ – 4.12). Under the same setting, Sequential Dijkstra forfeits 2.7%–2.8% of it. Replacing the ACO core with a PSO core instead reduces performance by 1.2% at $K = 5$. At $K \geq 20$, the reduction is 2.2%. Under a strict relay-capacity constraint ($C_{\max}=1$), the shortfall of Sequential Dijkstra widens to 30.8% of the full method's MF-SREE at $K = 20$. At $K = 30$, that shortfall reaches 39.0%.


Single-mechanism removals isolate individual contributions. Coordinated reroute refinement is the primary contributor. Its impact scales with network contention. Its removal costs -0.5% at $C_{\max}=N$. That cost reaches -11.2% at $C_{\max}=1$ and $K=30$ ($d_z = 3.28$, $p < 0.001$, winning across all 20 seeds). Capacity-blind candidate generation degrades performance comparably, by -11.0% under $C_{\max}=1$. Conversely, neither mechanism triggers at $C_{\max}=N$ or $K=5$ due to minimal relay saturation. Every remaining removal costs under one percent there. Two structural characteristics explain those sub-percent variations. The first is the limited headroom left by the $+2.4\%$ – $+4.6\%$ margin over Sequential Dijkstra at $C_{\max}=N$ under this stress profile. The second is the substitutability of the construction and refinement stages. Supplementary Material Section S-D.8 and Tables S-V–S-VII report the per-mechanism figures, the extended operating regimes, the refinement-budget isolation that establishes that substitutability, and the congestion-dependent evaporation term. That term shows no statistically significant contribution at these operating points and is omitted from Table 6.

5.4. Robustness and Stress Tests

We stress-test across MAC contention laws, mission propulsion levels, and node capacity limits, all three within the load-count surrogate of Section 3.1. As shown in Fig. 2a,



(a) MAC model sensitivity



(b) Endurance accounting under mission propulsion

Figure 2: Evaluation of MAC-contention laws and mission-propulsion power ($N=50$).

the paired advantage of S-ACS MF over Sequential Dijkstra at $K = 20$ widens from $+1.78\%$ under logarithmic contention accumulation to $+3.13\%$ under the linear law and $+3.74\%$ under the quadratic law. That advantage grows with the modeled contention penalty under the synthetic stress profile. The sweep was not repeated under the measured codec curve. Fig. 2b then sweeps the mission propulsion level $P_{\text{prop}} \in \{0, 100, 200, 300, 500\}$ W over a mission horizon $T_{\text{mis}} = 1800$ s at $K=10$. Consistent with the energy-accounting boundary of Section 3.3, P_{prop} is route-independent and cannot alter route selection. The sweep is therefore an endurance accounting check rather than a routing result. The communication energy is bit-identical at 3.6096 J across the five levels. Total MF-SREE is exactly invariant. S-ACS MF holds a constant 0.2394-unit margin over Sequential Dijkstra (14.5107 versus 14.2713) at every level. The sweep instead evaluates fleet endurance. The post-mission battery fraction falls from 100.0% at 0 W to 16.7% at 500 W. This confirms mission feasibility across the tested range. Detailed sensitivity analyses with respect to semantic fidelity parameters k_s and $S_{\min}$, AoSI constraints, channel fading, link outages, and per-flow fairness are reported in Supplementary Material Sections S-D.2–S-D.6, Figs. S-2–S-5, and Tables S-II–S-III.

Table 6


Baseline comparison and component ablation at $N=50$: MF-SREE (mean $\pm$ std over 20 trials), relative change Δ , paired Cohen's d_z (d) and p -value (p) against the reference row of the same block, on identical topology seeds. Competing methods are listed first in each block and single-mechanism removals second. The lower block tightens the service cap to $C_{\max}=1$, forbidding relay reuse. Bold marks the best value per column in the upper block; three further regimes ($R_{LS} \in \{0, 1\}$, $C_{\max}=2$) are discussed in the text.

Variant	$K=5$			$K=20$			$K=30$		
	MF-SREE ($\Delta\%$)	d	p	MF-SREE ($\Delta\%$)	d	p	MF-SREE ($\Delta\%$)	d	p
S-ACS MF (full)	7.57 $\pm$ 1.80 (ref)	—	—	27.94 $\pm$ 4.45 (ref)	—	—	41.46 $\pm$ 4.54 (ref)	—	—
Traditional ACO MF	6.61 $\pm$ 2.12 (−12.8%)	1.35	<0.001	22.32 $\pm$ 5.02 (−20.1%)	4.12	<0.001	33.81 $\pm$ 5.90 (−18.5%)	3.46	<0.001
Sequential Dijkstra	7.36 $\pm$ 1.86 (−2.8%)	0.89	<0.001	27.19 $\pm$ 4.57 (−2.7%)	1.61	<0.001	40.34 $\pm$ 4.66 (−2.7%)	1.62	<0.001
PSO engine (vs. ACO)	7.49 $\pm$ 1.80 (−1.2%)	0.42	0.078	27.32 $\pm$ 4.51 (−2.2%)	1.31	<0.001	40.54 $\pm$ 4.52 (−2.2%)	1.34	<0.001
w/o Coordinated reroute LS	7.57 $\pm$ 1.80 (+0.0%)	0.03	0.893	27.80 $\pm$ 4.56 (−0.5%)	0.67	0.007	41.22 $\pm$ 4.61 (−0.6%)	1.08	<0.001
w/o Load-aware Dijkstra cand.	7.53 $\pm$ 1.79 (−0.6%)	0.36	0.120	27.75 $\pm$ 4.43 (−0.7%)	0.54	0.025	41.22 $\pm$ 4.59 (−0.6%)	0.53	0.029
w/o Load-aware MF heuristic	7.55 $\pm$ 1.82 (−0.3%)	0.21	0.363	27.89 $\pm$ 4.49 (−0.2%)	0.23	0.320	41.34 $\pm$ 4.60 (−0.3%)	0.34	0.148
<i>Hard-binding capacity</i> ($C_{\max}=1$, no relay reuse)									
S-ACS MF ($C_{\max}=1$)	7.57 $\pm$ 1.83 (ref)	—	—	26.24 $\pm$ 4.94 (ref)	—	—	38.07 $\pm$ 5.07 (ref)	—	—
Sequential Dijkstra	6.87 $\pm$ 2.20 (−9.3%)	0.55	0.025	18.15 $\pm$ 4.82 (−30.8%)	2.13	<0.001	23.21 $\pm$ 6.19 (−39.0%)	2.63	<0.001
w/o Coordinated reroute LS	7.54 $\pm$ 1.88 (−0.4%)	0.22	0.330	24.51 $\pm$ 5.18 (−6.6%)	1.96	<0.001	33.81 $\pm$ 5.22 (−11.2%)	3.28	<0.001
w/o Capacity-aware candidates	7.54 $\pm$ 1.88 (−0.4%)	0.22	0.330	24.55 $\pm$ 5.18 (−6.5%)	1.89	<0.001	33.86 $\pm$ 5.25 (−11.0%)	3.17	<0.001


Fig. 3 evaluates performance under constrained service availability $C_{\max}$, still within the load-count surrogate and under the synthetic stress profile. S-ACS MF holds positive paired gains at every tested capacity level, namely +41.22% ($C_{\max}=1$), +11.60% ($C_{\max}=2$), +5.43% ($C_{\max}=3$), +2.83% ($C_{\max}=5$), and +2.66% ($C_{\max}=50$) at $K = 20$. At $N = 20, K = 10, C_{\max} = 1$, S-ACS MF achieves a +115.86% mean (+47.48% median) advantage over Sequential Dijkstra by elevating flow delivery from 62.5% to 91.0%. Supplementary Material Section S-D.11 plots that delivery ratio across the sweep.

5.5. Scheduled-Objective Evaluation and Binding QoS Constraints

To evaluate routing performance when contention is directly scheduled rather than priced by relay load alone, we introduce the S-ACS MF-SINR variant. Its routing objective is evaluated under a conflict-graph TDMA schedule with spatial SINR, whereas flow admission and candidate filtering retain the load-count surrogate. The SINR admission threshold γ_{th} in (27) is set to 5.9738 dB. This value is derived from the nominal SNR at maximum separation $d_{\max}=500$ m under horizontal alignment ($\Delta_{alt}=0$ m), distinct from the 6.6 dB margin at $\Delta_{alt}=50$ m (Section 3.1). On this basis, Table 7 compares three routing regimes across eight cells on identical topology seeds: $N=20$ with $K=10$ and $N=50$ with $K \in \{10, 25, 50\}$ requested flows (49 admitted at the largest cell), crossed with binding ($C_{\max}=2$) and non-binding ($C_{\max}=N$) relay capacity. Per-flow QoS budgets remain inactive ($S_k^{req} = 0, E_{\max}^{(k)} = \infty$). In the re-scored regime, methods route under native objectives and are re-evaluated under TDMA-SINR. In the matched regime, routing directly optimizes the scheduled objective, whereas the surrogate regime uses load-count contention for both. Under re-scoring, four of eight contrasts remain significant after Holm–Bonferroni correction. Under the matched regime, the proposed variant is the top-ranked method in three of eight cells across the 60-test family (52 rejections at $\alpha = 0.05$).



(a) Paired advantage versus relay capacity



(b) Contribution of capacity-aware local search

Figure 3: Performance metrics under binding service-availability constraints.

At the densest non-binding cell ($N=50$ with 50 flows), its matched mean rank is 3.75, trailing Traditional ACO MF by −12.75% (CI [−19.72, −5.79], $p_{adj} = 0.006$).

Crossing the scheduled objective with the semantic-fidelity model and with the constraint regime yields the twelve-cell grid of Table 8. The grid takes three operating

Table 7

Ranking and paired margin evaluation of S-ACS MF under TDMA and spatial-SINR scheduling across rescored, surrogate, and matched regimes (mean rank with rank-1 frequency in parentheses; Holm–Bonferroni correction at $\alpha=0.05$).

N	K	Rank, rescored	Rank, surrogate	Rank, matched	Weakest sig. margin	Sig./tested	Baselines helped
<i>Binding relay capacity</i>							
20	10	2.08 (35%)	2.42 (10%)	1.93 (30%)	+18.62% [+9.79, +27.46] vs. Trad. ACO MF $p<0.001$	7/8	3/8
50	10	2.42 (10%)	1.85 (5%)	1.55 (35%)	+24.29% [+18.95, +29.62] vs. Trad. ACO MF $p<0.001$	7/8	2/8
50	25	1.25 (75%)	1.50 (80%)	1.15 (95%)	+26.46% [+19.00, +33.93] vs. Trad. ACO MF $p<0.001$	7/7	2/7
50	50	1.00 (100%)	2.70 (30%)	1.80 (50%)	+13.78% [+7.66, +19.89] vs. QMIX MF $p=0.002$	6/7	2/7
<i>Non-binding relay capacity</i>							
20	10	2.50 (25%)	2.33 (10%)	1.62 (50%)	+23.50% [+13.75, +33.25] vs. QMIX MF $p<0.001$	7/8	4/8
50	10	2.42 (15%)	1.75 (10%)	1.35 (50%)	+28.49% [+20.52, +36.46] vs. Trad. ACO MF $p<0.001$	7/8	2/8
50	25	1.57 (40%)	2.40 (50%)	1.10 (95%)	+20.83% [+14.18, +27.49] vs. Trad. ACO MF $p<0.001$	7/7	1/7
50	50	1.35 (65%)	5.50 (0%)	3.75 (10%)	−12.75% [−19.72, −5.79] vs. Trad. ACO MF $p=0.006$	4/7	3/7

points, namely $N=20$ with $K=10$ and $N=50$ with $K \in \{25, 49\}$. It crosses each with two fidelity curves and both constraint regimes at 20 paired topology seeds per cell. The first curve is the measured DeepSC codec model of Section 3.3, which carries the primary evidence. The second is the synthetic sigmoid profile of Table 3, retained as a stress companion rather than as a codec claim.

In contrast to Table 7, the binding regime in Table 8 jointly tightens relay capacity to $C_{\max}=2$ while activating fidelity floor S_k^{req} and route-energy cap $E_{\max}^{(k)}$, whereas the non-binding regime sets $C_{\max}=N$ with inactive budgets. Budget levels are calibrated per cell to the median delivered fidelity and route energy. Under the measured DeepSC curve, the resulting pairs $(S_k^{\text{req}}, E_{\max}^{(k)})$ are (0.9487, 0.6191 J) at $N=20, K=10$, (0.9251, 0.6279 J) at $N=50, K=25$, and (0.8559, 1.2322 J) at $N=50, K=49$. Under the synthetic profile, the corresponding pairs are (0.4230, 0.6199 J), (0.4542, 1.2293 J), and (0.3103, 1.2359 J) at the same three points. Because each curve yields distinct distributions, the levels differ by roughly a factor of two, and performance cannot be compared across curves.

At the dense cell ($N=50, K=49$), S-ACS MF-SINR ranks third of four schedule-evaluated baselines across both constraint regimes and fidelity curves. Under the measured DeepSC curve with binding constraints, it trails Traditional ACO MF by −11.57% (CI [−13.73, −9.41], $p_{\text{adj}} = 2.13 \times 10^{-7}$) and Greedy Walk MF by −11.73% (CI [−14.28, −9.18], $p_{\text{adj}} = 7.20 \times 10^{-7}$), while leading Sequential Dijkstra by +365.14% (CI [+224.87, +505.41], $p_{\text{adj}} = 8.09 \times 10^{-14}$). Because rankings evaluate summed MF-SREE while the router lexicographically maximizes admitted flows first, this standing reflects a service trade: the proposed variant admits 36.1 flows versus 27.6 for the leading rival under the measured curve, and 32.5 versus 20.1 under the synthetic profile, reporting both keys in Table 8. Under non-binding capacity where both

methods admit all 49 flows, the margin is −14.48% against Traditional ACO MF (CI [−21.72, −7.24], $p_{\text{adj}} = 0.002$), −5.58% against Greedy Walk MF (CI [−13.23, +2.06], $p_{\text{adj}} = 0.111$), and +109.21% over Sequential Dijkstra (CI [+85.15, +133.28], $p_{\text{adj}} = 5.03 \times 10^{-9}$).

Table 9 and Supplementary Material Table S-XII report paired margins across all cells. Under measured DeepSC, margins against Traditional ACO MF and Greedy Walk MF are positive across all $N=20$ cells and at $N=50, K=25$, reversing to negative only at $N=50, K=49$. Under the synthetic profile, the non-binding margin against Traditional ACO MF is negative at $N=50, K=25$. S-ACS MF-SINR secures rank one in six of twelve cells ($N=20, K=10$ or $N=50, K=25$). In an ablation maximizing metric first, rank-one standing increases from six to ten of twelve cells, reversing the non-binding dense margin against Traditional ACO MF from −14.48% to +20.91% (CI [+11.27, +30.56], $p_{\text{adj}} = 2.84 \times 10^{-4}$) at the expense of shedding traffic (49.0 down to 35.8 delivered flows). Under binding capacity, dense margins remain negative at −10.28% against Traditional ACO MF across 20 seeds and −10.45% against Greedy Walk MF. Hence, admission-first routing is retained.

Table 10 closes the remaining gap by sweeping the two per-flow QoS budgets held inactive in Section 5.1 from non-binding through to infeasible. This sweep evaluates $N=20$ with $K=10$ and $N=50$ with $K=25$ under the measured DeepSC codec recursion. Each swept level is an empirical percentile of the per-flow fidelity and route-energy values pooled across all delivered flows and seeds at the unconstrained operating point. Of the thirty-six swept levels, twelve are measured binding on the fidelity-floor axis and twelve on the per-flow energy-cap axis. Across this grid, statistical advantage is established by the Holm-corrected family of 98 paired tests, of which 42 reject at $\alpha=0.05$. The proposed method achieves the highest mean MF-SREE in twenty of the twenty-four measured binding levels. In the remaining four boundary levels, all evaluated methods tie

Table 8

Ranking and feasibility of S-ACS MF-SINR under scheduled TDMA and spatial-SINR evaluation across semantic codec curves and capacity regimes (bold marks highest admissions and rank-1; S: S-ACS MF-SINR, D: Sequential Dijkstra, T: Traditional ACO MF, G: Greedy Walk MF, L: LP-MCF; admission counts follow the method order in the Ranking column; feasibility refers to S-ACS MF-SINR).

N, K	Curve	Regime	Feasibility	Admitted	Ranking	Rank/Rank-1
20, 10	Fitted	Binding ($C_{\max}=2$)	60.0% [51.4, 68.6]	6.0/5.4/4.9/4.7/2.8	S>L>T>G>D	1/5 (55.0%)
20, 10	Fitted	Non-bind. ($C_{\max}=20$)	100.0% [100.0, 100.0]	10.0/10.0/10.0/9.9/10.0	L>S>T>G>D	2/5 (40.0%)
20, 10	Synth.	Binding ($C_{\max}=2$)	59.0% [51.7, 66.3]	5.9/4.8/3.5/3.4/2.9	S>L>T>G>D	1/5 (55.0%)
20, 10	Synth.	Non-bind. ($C_{\max}=20$)	100.0% [100.0, 100.0]	10.0/10.0/10.0/9.9/10.0	S>L>T>G>D	1/5 (55.0%)
50, 25	Fitted	Binding ($C_{\max}=2$)	77.6% [74.3, 80.9]	19.4/12.6/11.8/4.4	S>T>G>D	1/4 (85.0%)
50, 25	Fitted	Non-bind. ($C_{\max}=50$)	100.0% [100.0, 100.0]	25.0/25.0/23.6/24.9	S>T>G>D	1/4 (85.0%)
50, 25	Synth.	Binding ($C_{\max}=2$)	68.2% [63.9, 72.5]	17.1/7.4/7.0/6.0	S>G>T>D	1/4 (40.0%)
50, 25	Synth.	Non-bind. ($C_{\max}=50$)	100.0% [100.0, 100.0]	25.0/23.6/25.0/25.0	T>G>S>D	3/4 (20.0%)
50, 49	Fitted	Binding ($C_{\max}=2$)	73.8% [70.4, 77.2]	27.6/28.7/36.1/7.5	G>T>S>D	3/4 (0.0%)
50, 49	Fitted	Non-bind. ($C_{\max}=50$)	100.0% [100.0, 100.0]	49.0/46.1/49.0/48.5	T>G>S>D	3/4 (10.0%)
50, 49	Synth.	Binding ($C_{\max}=2$)	66.2% [62.8, 69.6]	20.1/15.1/32.5/8.3	G>T>S>D	3/4 (5.0%)
50, 49	Synth.	Non-bind. ($C_{\max}=50$)	100.0% [100.0, 100.0]	49.0/46.1/49.0/49.0	T>G>S>D	3/4 (0.0%)

Table 9

Paired percentage MF-SREE margins of S-ACS MF-SINR against schedule-evaluated rivals across the scheduled crossed grid of Table 8 (mean $\pm$ 95% CI over 20 paired topology seeds; * indicates Holm–Bonferroni significance at $\alpha=0.05$; —: not evaluated; † undefined relative margin (all 20 seeds won, Holm-significant); adjusted p -values and win/tie/loss records in Supplementary Material Table S-XII).

N, K	Regime	vs. Seq. Dijkstra (%)	vs. Trad. ACO MF (%)	vs. Greedy Walk MF (%)	vs. LP-MCF (%)
<i>Measured DeepSC codec profile</i>					
20, 10	Binding ($C_{\max}=2$)	—†	+19.93 $\pm$ 13.39*	+19.86 $\pm$ 10.66*	+11.21 $\pm$ 7.79*
20, 10	Non-bind. ($C_{\max}=20$)	+89.94 $\pm$ 26.36*	+16.69 $\pm$ 9.99*	+42.19 $\pm$ 19.15*	−1.69 $\pm$ 1.74
50, 25	Binding ($C_{\max}=2$)	+386.68 $\pm$ 182.26*	+17.98 $\pm$ 10.15*	+15.25 $\pm$ 5.47*	—
50, 25	Non-bind. ($C_{\max}=50$)	+164.33 $\pm$ 30.29*	+19.41 $\pm$ 6.99*	+36.68 $\pm$ 11.71*	—
50, 49	Binding ($C_{\max}=2$)	+365.14 $\pm$ 140.27*	−11.57 $\pm$ 2.16*	−11.73 $\pm$ 2.55*	—
50, 49	Non-bind. ($C_{\max}=50$)	+109.21 $\pm$ 24.06*	−14.48 $\pm$ 7.24*	−5.58 $\pm$ 7.64	—
<i>Synthetic fidelity profile</i>					
20, 10	Binding ($C_{\max}=2$)	—†	+41.63 $\pm$ 30.79*	+47.17 $\pm$ 31.93*	+22.55 $\pm$ 14.50*
20, 10	Non-bind. ($C_{\max}=20$)	+151.71 $\pm$ 52.78*	+25.23 $\pm$ 12.13*	+67.30 $\pm$ 24.67*	+5.53 $\pm$ 7.43
50, 25	Binding ($C_{\max}=2$)	+233.68 $\pm$ 102.11*	+7.24 $\pm$ 8.57	+2.79 $\pm$ 9.16	—
50, 25	Non-bind. ($C_{\max}=50$)	+108.82 $\pm$ 35.53*	−15.73 $\pm$ 10.18*	+1.30 $\pm$ 14.75	—
50, 49	Binding ($C_{\max}=2$)	+320.81 $\pm$ 79.97*	−11.87 $\pm$ 7.70*	−17.38 $\pm$ 5.40*	—
50, 49	Non-bind. ($C_{\max}=50$)	+77.96 $\pm$ 29.93*	−27.38 $\pm$ 5.80*	−15.66 $\pm$ 7.71*	—

identically across all twenty seeds with zero mean difference and zero difference variance (Table 10; Supplementary Material Table S-XI). Under those severe budgets, every method sheds all flows on most seeds and delivers the identical single flow on surviving seeds.

Under the fidelity-floor sweep at $N=20$ with $K=10$, the feasibility rate of the proposed method falls from $99.5\% \pm 1.0\%$ at the P10 level to $1.0\% \pm 1.4\%$ at P99. Rejections by the proposed method stem from the absence of a feasible route, whereas baselines reject on fidelity and energy caps (Supplementary Material Table S-X). The fidelity of the worst delivered flow rises over the sweep from 0.9155 to 0.9744, reflecting the exclusion of difficult flows from the admitted set. Notably, this admission filtering separates the two fairness scopes: while Jain's index over delivered flows remains between 0.884 and 0.968 from P10 to P75, the index over all requested flows drops monotonically from 0.880 to 0.265, reaching 0.010 at P99. A delivered-only fairness metric thus overstates equity once admissions are refused.

Both fairness columns in Table 10 report means over 20 trials. Per-trial Jain's index evaluates the per-flow SREE

vector, where a trial with zero delivered flows contributes 0. Hence, the delivered-flow column reflects an unconditional mean: at $N=20, K=10$ under the P99 floor, 2 of 20 trials deliver a single flow (scoring 1) while 18 deliver none (scoring 0), yielding the tabulated mean 0.100. Under the all-requested scope, each delivering trial contributes $1/K = 0.1$, averaging to 0.010, whereas conditional fairness among served flows is 1 across the delivering trials. Similarly, Agg. eff. normalizes total admitted MF-SREE by requested flow count K , decreasing as flows are shed. The same trend holds at $N=50$ with $K=25$. Furthermore, six of the twelve energy-binding levels replicate fidelity-floor levels identically, as both budgets reject identical flows. Table 10 lists each unique pair, while Supplementary Material Table S-X provides the full census across all thirty-six levels.

Two properties characterize the fidelity-floor placement. First, upper levels collapse onto the measured codec ceiling of 0.9744. At $N=20, K=10$, applied floors differ relatively by 4.3×10^{-5} between P75 and P90, and by 6.8×10^{-8} between P90 and P99; at $N=50, K=25$, the corresponding differences are 8.3×10^{-5} and 2.0×10^{-8} . The sharp drop in

Table 10

Fairness, feasibility, and efficiency metrics for S-ACS MF under binding per-flow QoS constraints (S_k^{req} and $E_{\text{max}}^{(k)}$) under measured DeepSC codec recursion (mean $\pm$ 95% CI).

Level	Feasibility rate	Rejected/trial	Worst fidelity	Worst energy (J)	Jain (deliv.)	Jain (all req.)	Agg. eff. (sem./J)
<i>N=20, K=10 — Semantic fidelity floor S_k^{req} sweep (dimensionless)</i>							
P10 (0.8764022)	99.5% $\pm$ 1.0%	0.05 $\pm$ 0.10	0.9155 $\pm$ 0.0090	0.872 $\pm$ 0.089	0.884 $\pm$ 0.014	0.880 $\pm$ 0.018	2.247 $\pm$ 0.153
P25 (0.9343057)	88.0% $\pm$ 3.9%	1.20 $\pm$ 0.39	0.9430 $\pm$ 0.0022	0.624 $\pm$ 0.001	0.913 $\pm$ 0.011	0.804 $\pm$ 0.038	2.127 $\pm$ 0.175
P50 (0.9494202)	54.5% $\pm$ 8.2%	4.55 $\pm$ 0.82	0.9607 $\pm$ 0.0054	0.464 $\pm$ 0.073	0.968 $\pm$ 0.018	0.531 $\pm$ 0.084	1.619 $\pm$ 0.272
P75 (0.9743538)	26.5% $\pm$ 7.9%	7.35 $\pm$ 0.79	0.9744 $\pm$ <0.0001	0.308 $\pm$ <0.001	0.950 $\pm$ 0.105	0.265 $\pm$ 0.079	0.839 $\pm$ 0.251
P90 (0.9743954)	12.0% $\pm$ 6.5%	8.80 $\pm$ 0.65	0.9744 $\pm$ <0.0001	0.307 $\pm$ <0.001	0.650 $\pm$ 0.229	0.120 $\pm$ 0.065	0.381 $\pm$ 0.208
P99 (0.9743955)	1.0% $\pm$ 1.4%	9.90 $\pm$ 0.14	0.9744 $\pm$ <0.0001	0.307 $\pm$ 0.001	0.100 $\pm$ 0.144	0.010 $\pm$ 0.014	0.032 $\pm$ 0.046
<i>N=20, K=10 — Per-flow energy cap $E_{\text{max}}^{(k)}$ sweep (J/task)</i>							
P50 (0.616)	51.5% $\pm$ 8.2%	4.85 $\pm$ 0.82	0.9664 $\pm$ 0.0047	0.388 $\pm$ 0.063	0.982 $\pm$ 0.017	0.508 $\pm$ 0.084	1.572 $\pm$ 0.273
P75 (0.625)	87.0% $\pm$ 3.8%	1.30 $\pm$ 0.38	0.9427 $\pm$ 0.0031	0.623 $\pm$ 0.001	0.914 $\pm$ 0.012	0.795 $\pm$ 0.037	2.111 $\pm$ 0.178
P90 (0.934)	97.5% $\pm$ 2.6%	0.25 $\pm$ 0.26	0.9199 $\pm$ 0.0118	0.778 $\pm$ 0.073	0.891 $\pm$ 0.014	0.869 $\pm$ 0.025	2.228 $\pm$ 0.158
<i>N=50, K=25 — Semantic fidelity floor S_k^{req} sweep (dimensionless)</i>							
P10 (0.6761514)	100.0% $\pm$ 0.0%	0.00 $\pm$ 0.00	0.8965 $\pm$ 0.0250	0.887 $\pm$ 0.097	0.875 $\pm$ 0.010	0.875 $\pm$ 0.010	2.220 $\pm$ 0.094
P25 (0.9251283)	90.8% $\pm$ 3.4%	2.30 $\pm$ 0.85	0.9350 $\pm$ 0.0031	0.655 $\pm$ 0.043	0.895 $\pm$ 0.008	0.813 $\pm$ 0.031	2.124 $\pm$ 0.114
P50 (0.9493988)	60.4% $\pm$ 5.1%	9.90 $\pm$ 1.27	0.9505 $\pm$ 0.0022	0.602 $\pm$ 0.032	0.944 $\pm$ 0.013	0.570 $\pm$ 0.048	1.666 $\pm$ 0.146
P75 (0.9743141)	27.2% $\pm$ 3.5%	18.20 $\pm$ 0.88	0.9744 $\pm$ <0.0001	0.309 $\pm$ <0.001	1.000 $\pm$ <0.001	0.272 $\pm$ 0.035	0.861 $\pm$ 0.111
P90 (0.9743955)	10.8% $\pm$ 1.8%	22.30 $\pm$ 0.46	0.9744 $\pm$ <0.0001	0.307 $\pm$ <0.001	1.000 $\pm$ <0.001	0.108 $\pm$ 0.018	0.343 $\pm$ 0.058
P99 (0.9743955)	1.0% $\pm$ 0.8%	24.75 $\pm$ 0.21	0.9744 $\pm$ <0.0001	0.307 $\pm$ <0.001	0.250 $\pm$ 0.208	0.010 $\pm$ 0.008	0.032 $\pm$ 0.026
<i>N=50, K=25 — Per-flow energy cap $E_{\text{max}}^{(k)}$ sweep (J/task)</i>							
P50 (0.617)	57.8% $\pm$ 5.0%	10.55 $\pm$ 1.25	0.9527 $\pm$ 0.0038	0.571 $\pm$ 0.052	0.951 $\pm$ 0.013	0.549 $\pm$ 0.048	1.626 $\pm$ 0.149
P75 (0.626)	88.2% $\pm$ 3.6%	2.95 $\pm$ 0.90	0.9356 $\pm$ 0.0064	0.624 $\pm$ 0.001	0.899 $\pm$ 0.008	0.793 $\pm$ 0.033	2.087 $\pm$ 0.113
P90 (0.940)	100.0% $\pm$ 0.0%	0.00 $\pm$ 0.00	0.8902 $\pm$ 0.0331	0.856 $\pm$ 0.063	0.875 $\pm$ 0.009	0.875 $\pm$ 0.009	2.221 $\pm$ 0.093

feasibility thus reflects distribution density near the ceiling rather than a substantially tighter constraint. Second, at $N=50, K=25$, the boundary-regime floor applies a bound of 0.5% of the cell's P100 anchor, serving as an unconstrained boundary baseline.

5.6. Discussion

The flow-density and nominal-SNR sweeps at $N=50$ were both obtained under the load-count contention surrogate and the synthetic stress profile. Supplementary Material Section S-D.10 reports them. The decay of the paired advantage with nominal SNR is consistent with the fidelity saturation of the sigmoid (6) which defines the stress profile. The calibrated curve is a measured fidelity recursion. However, that recursion reduces exactly to the same per-hop sigmoid family under the same composition law, as Section 3.3 establishes. The two profiles therefore differ in their curve parameters rather than in functional form, and most consequentially in their midpoints, at -0.8934 dB against 10 dB. Consequently, the same saturation mechanism acts on both arms. It simply acts earlier on the calibrated arm, which already sits in its saturated region at the tested link margins. Notably, the calibrated curve governs the primary evaluations in Tables 4, 8, and 10. By contrast, component attribution, capacity sweeps, alternative contention laws, and zero-load gap analyses were evaluated under the synthetic stress profile.


Primary scope qualifications apply to physical channel and interference dynamics. In particular, comparisons on the surrogate track are evaluated under the load-count contention surrogate (3) defined in Section 3.1. Under the worst-case full-reuse spatial evaluator in Supplementary Material Section S-D.1, the surrogate preserves link-quality ordering

only partially ($\rho = 0.668$ – 0.748 across active edges), and method rankings do not carry over. Physically, uncoordinated spatial interference penalizes multi-hop paths heavily because concurrent transmissions aggregate interference, favoring shorter, simpler routes. Because the load-count surrogate prices contention without transmitter geometry, it cannot reflect this spatial penalty. The detailed eight-cell grid and its paired contrasts are reported in Supplementary Material Section S-D.1. Correspondingly, Section 5.5 details routing performance when the optimization objective directly incorporates an explicit schedule.


Separately, centralized optimization introduces computation-induced route staleness under mobility. Across re-solve latencies of 0.196–0.888 s and UAV speeds up to 20 m/s (displacement up to 17.8 m), S-ACS MF retains 98.89%–100.00% of static MF-SREE and maintains a +1.97% to +3.28% gain over Sequential Dijkstra (Fig. 4). These measured latencies from Section 5.2 satisfy the receding-horizon deadline $\eta\Delta T_{\text{topo}} = 1.0$ s ($\Delta T_{\text{topo}} = 2.0$ s, $\eta = 0.5$). In addition, three physical dynamics are evaluated analytically in Supplementary Material Section S-D.5: Doppler frequency shift ($f_d \approx 133$ Hz at 20 m/s and $f_c = 2$ GHz), 3D altitude coupling over $z_i \in [50, 300]$ m, and fast topology updates ($\Delta T_{\text{topo}} < 0.1$ s). Continuous 3D flight trajectory co-design remains an explicit scope boundary.

5.7. Gap Analysis against the Certified Zero-Load Ceiling

We next measure how far S-ACS MF sits below the certified per-flow zero-load ceiling of Proposition 1. As shown in Fig. 5a, across 560 trials over $N \in \{20, 30, 40, 50\}$ and $K \in$



(a) Paired advantage under route staleness





(b) MF-SREE retention under fleet movement

Figure 4: Computation-induced route staleness under centralized control.

$\{2, 4, \dots, 14\}$, 95.5% of instances (535/560) exhibit zero-load-ceiling gaps below 30%. Per-cell median gaps range between 1.5% and 12.9%. The ceiling is capacity- and load-free by construction. Consequently, these are gaps against a relaxation and not optimality gaps. The relaxation loosens as relay contention grows. The gap therefore overstates the distance to the capacity-constrained optimum. The bound is self-calibrating (Fig. 5b). Specifically, it exhibits a strong negative Spearman correlation between gap and per-flow SREE ($\rho = -0.867$, $p = 6.1 \times 10^{-171}$).

Table 11 benchmarks S-ACS MF against exact global optima Z^* obtained via full state-space enumeration on small-scale networks ($N \leq 10$, $K \leq 3$). Over the enumerable instances of that table the measured gap to the exact optimum spans 0.00%–0.15%. Every optimality certification reported in this paper is limited to those small-scale enumerable instances. Each was obtained under the synthetic stress profile. At larger N and K , the only available reference is the zero-load ceiling above. That ceiling is a loose upper bound rather than an optimum. Therefore, no optimality claim is made at those scales. Empirical evaluation of Proposition 5 gives a fixed-load LP-IP gap $|Z_{LP} - Z_{IP}| < 10^{-10}$ on the topologies tested there. This verification is confined to the single-flow case $K=1$.

(a) Dinkelbach gap versus K 

(b) Self-calibrating property

Figure 5: Evaluation of certified per-flow zero-load ceiling and surrogate gap.**Table 11**

Calibration of S-ACS MF against exact global optima Z^* (obtained via full state-space enumeration) of the coupled MF-SREE objective ($N \leq 10$, $K \leq 3$, 20 topology seeds per cell, non-binding relay capacity). Relative gaps $(Z^* - \text{SREE})/Z^* \times 100\%$ are reported; best values per column are bolded.

Quantity (N, K)	(8,2)	(8,3)	(10,2)	(10,3)
Exact global optimum Z^*	2.9486	4.1821	3.1361	4.6292
Gap to Z^* (%), lower is better				
S-ACS MF	0.0000	0.0000	0.0000	0.1502
LP-MCF surrogate	1.5696	2.5150	1.9322	1.2561
S-ACS MF, 95 % CI	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[-0.16, 0.46]
Exact optimal (trials)	20/20	20/20	20/20	19/20
Bound diagnostic				
Ceiling gap above Z^* (%)	14.726	13.895	25.192	39.079

6. Conclusion

In this paper, we investigated energy-efficient multi-flow semantic routing for multi-UAV relay networks under load-count contention coupling. To trade semantic throughput against energy consumption, we formulated the Multi-Flow Semantic Routing Energy Efficiency (MF-SREE) problem and established its NP-hardness at unit service availability. Accordingly, we proposed S-ACS MF, a metaheuristic framework built on a load-aware energy-per-fidelity construction rule, coupled-objective marginal path evaluation,

and capacity-aware local search. To align the routing objective with the channel-access model that scores it, we further developed S-ACS MF-SINR, a scheduling-aware variant that constructs routes under an explicit transmission schedule.

Simulation results demonstrate that S-ACS MF achieves superior energy efficiency across most evaluated settings under the load-count contention surrogate, the codec-calibrated fidelity curve, and non-binding capacity ($C_{\max}=N$). That advantage is pronounced against metaheuristic and greedy benchmarks, and narrow against the shortest-path benchmark. It closely approaches the certified zero-load ceiling and enumerable global optima on small-scale instances. While the calibrated curve establishes the margin magnitudes, baseline orderings remain consistent across profiles. Under the same surrogate, the framework maintains its lead across most binding service-fidelity and per-flow energy levels, while demonstrating robustness against tight relay capacity, route staleness, and alternative contention laws. Furthermore, runtime evaluations reveal that matching the compilation setting and ant budget reverses the earlier wall-clock margin over traditional ant colony optimization, showing that execution-time differences stem from solver configuration rather than algorithmic design.

Under the scheduled routing objective, schedule-evaluated comparisons indicate that S-ACS MF-SINR ranks favorably at moderate flow densities, while simpler constructive heuristics attain higher summed efficiency at the densest, most contended operating points. Notably, the proposed variant admits the most flows at the densest binding points, so the comparison records a trade between service coverage and summed efficiency rather than a uniform deficit. The scheduling-aware formulation thus provides structural consistency between optimization and evaluation models rather than an unconditional performance lead.

Several scope boundaries characterize the current framework. Specifically, the formulation assumes a load-count contention surrogate, communication-layer energy accounting (excluding route-independent propulsion), centralized topology control, and frozen-snapshot optimization per epoch ΔT_{topo} . Under a worst-case full-reuse spatial evaluator, route simplicity alters method rankings. In future work, we will extend scheduling awareness to flow admission and candidate generation, incorporate explicit transmitter geometry and inter-cell interference, and integrate mobility speed directly into the age-of-semantic-information model. Additional directions include column generation for tighter optimality bounds, Doppler-resilient waveforms, 3D trajectory co-design, decentralized stigmergic coordination, and experimental testbed validation.

A. Proofs

A.1. Structural Non-Decomposability of SREE

Remark 3 (Non-decomposability of SREE). *The single-flow SREE metric $\text{SREE}(\mathcal{P}) = \prod_h S_h / \sum_h E_h$ is a non-linear fractional program with a multiplicative numerator and additive denominator. This structure violates Bellman's principle of optimality. Knowing that subpath $\mathcal{P}_{sv}^*$ is SREE-optimal from s to v does not imply that extending $\mathcal{P}_{sv}^*$ to d yields an optimal s - d path, because appending edges couples numerator degradation multiplicatively with denominator energy addition. This rules out exact polynomial-time dynamic programming, distinguishing SREE from standard shortest-path objectives that decompose additively. The multi-flow extension (Theorem 9) inherits this structural difficulty and introduces additional hardness through inter-flow load-count contention coupling.*

A.2. NP-Hardness of Multi-Flow SREE

Definition 8 (Multi-Flow SREE (MF-SREE)). *Given a directed graph $G = (\mathcal{V}, \mathcal{E})$ with per-edge semantic fidelity $S(\gamma_{ij}) \in (0, 1]$, per-edge energy $E_{ij} > 0$, and K source-destination pairs $\{(s_k, d_k)\}_{k=1}^K$, define the per-flow SREE under interference as*

$$\text{SREE}_k(\mathcal{P}_k, \mathbf{L}) = \frac{\prod_{(i,j) \in \mathcal{P}_k} S(\gamma_{ij} - 10 \log_{10}(1 + L_j))}{\sum_{(i,j) \in \mathcal{P}_k} E_{ij}(L_j)}, \quad (34)$$

where $L_j = |\{k' : j \in \mathcal{P}_{k'} \setminus \{s_{k'}, d_{k'}\}\}|$ is the relay load at node j , as defined in (2), induced by the path assignment $\mathbf{P} = (\mathcal{P}_1, \dots, \mathcal{P}_K)$, with $\mathbf{P} \in \prod_{k=1}^K \Pi(s_k, d_k)$ and $\mathbf{L} = (L_1, \dots, L_N) \in \{0, 1, \dots, K\}^N$. The per-hop energy $E_{ij}(L_j)$ is any function that is strictly positive and non-decreasing in the relay load, that is, $E_{ij}(L) > 0$ for every L , and $E_{ij}(L) \leq E_{ij}(L')$ whenever $L \leq L'$. This family is the weakest structure the hardness reduction from vertex-disjoint paths (VDP) below actually uses. Notably, it contains the simplified linear form $E_{ij}(L_j) = E_{ij}(1 + L_j)$, which is convenient for exposition. It equally contains the operational hop-energy model of (10), whose transmission term grows with load through the reduced effective capacity while its inference and processing terms remain additive. The reduction therefore establishes hardness for the operational problem directly. In particular, it does not require the two models to coincide at any value of $C_{\max}$, and they do not coincide. The MF-SREE optimization problem seeks

$$\begin{aligned} \max_{\mathcal{P}_1, \dots, \mathcal{P}_K} \quad & \sum_{k=1}^K \text{SREE}_k(\mathcal{P}_k, \mathbf{L}(\mathbf{P})) \\ \text{s.t.} \quad & \mathcal{P}_k \in \Pi(s_k, d_k), \quad L_j \leq C_{\max}, \quad \forall j, k. \end{aligned} \quad (35)$$

Every flow carries a path in (35). This problem is therefore the restriction of the lexicographic design problem (17) to assignments that admit the whole demand, that is, to the top value $K_{\text{adm}} = K$ of its first key. On that restriction the first key is constant and the second key (17a) reduces

to the sum in (35). The remaining budgets of (17) are left inactive in the constructed instances below. The per-flow fidelity floor is set to $S_k^{\text{req}} = 0$ in (17g), which then holds vacuously. The hop cap $H_{\max}$ in (17e) is set to the node count of the constructed graph, which no simple path can exceed. The staleness budget $\text{AoSI}_{\max}^{(k)}$ in (17d), the per-flow energy cap $E_{\max}^{(k)}$ in (17h), and the per-node residual battery in (17c) are each set above the largest value any simple path or node can realize on that finite graph. Each such bound is obtained in polynomial time from the edge data. Service availability (17f) is thus the only constraint that binds, and it is the one the reduction exploits.

Theorem 9 (NP-Hardness of MF-SREE). *Consider the decision version of the MF-SREE problem (35): given the instance data and a threshold $T > 0$, determine whether a feasible path assignment exists that admits every flow and attains $\sum_{k=1}^K \text{SREE}_k \geq T$. Equivalently, since $K_{\text{adm}} \leq K$ always holds, this asks whether the lexicographic objective (17a) attains a value of at least (K, T) . This decision problem is NP-hard, even for $K = 2$ and unit service availability $C_{\max} = 1$, and for every per-hop energy model in the family of Definition 8.*

Proof. We prove NP-hardness by reduction from the *Vertex-Disjoint Paths* (VDP) problem on directed graphs, which is NP-complete even for $K = 2$ pairs [16]. A VDP instance is: given a directed graph $G' = (V', E')$ and $K' = 2$ terminal pairs (s'_1, d'_1) and (s'_2, d'_2) , determine whether there exist two internally vertex-disjoint s'_k - d'_k paths in G' .

Construction. Without loss of generality, we assume that in G' each source s'_k has in-degree zero and each destination d'_k has out-degree zero. This terminal-isolated form is obtained by appending fresh terminal nodes connected to the originals by a unique edge and re-designating the terminals, which preserves VDP YES/NO [16]. Under this preprocessing, no s'_k - d'_k path can use any terminal node as an interior relay, since interior membership requires both incoming and outgoing edges incident to that node. Given this VDP instance, we construct an MF-SREE instance as follows. Set $G = G'$ with the same vertices and edges. For each edge $(i, j) \in E'$: $S(\gamma_{ij}) = S_0 \in (0, 1)$ (a constant) and $C_{\max} = 1$. The per-hop energy is drawn from the family of Definition 8 and is otherwise unrestricted; in particular, it may differ from edge to edge. Under $C_{\max} = 1$ the load argument takes only the two reachable values $L_j \in \{0, 1\}$, so we collect the two edge-wise maxima $\bar{\epsilon}_0 \triangleq \max_{(i,j) \in E'} E_{ij}(0)$ and $\bar{\epsilon}_1 \triangleq \max_{(i,j) \in E'} E_{ij}(1)$, which satisfy $0 < \bar{\epsilon}_0 \leq \bar{\epsilon}_1$ by positivity and monotonicity. Both maxima are computable in time linear in $|E'|$, so the construction remains polynomial. Since every edge has the same nominal fidelity S_0 , define the *single-flow loaded fidelity* $S_1 \triangleq S(\gamma_0 - 10 \log_{10}(2))$ where γ_0 satisfies $S(\gamma_0) = S_0$. At a relay serving exactly one flow ($L_j = 1$), the interference model (34) gives effective fidelity $S_1 < S_0$ and a per-hop energy of at most $\bar{\epsilon}_1$. At a shared relay ($L_j \geq 2$), the effective fidelity degrades further. Moreover, the service-availability constraint is violated. Set the terminal pairs $(s_k, d_k) = (s'_k, d'_k)$ for $k \in \{1, 2\}$.

Analysis of feasible solutions. Under the service-availability constraint $C_{\max} = 1$, each relay node can serve at most one flow. If the two paths share a relay node j , then $L_j = 2 > C_{\max}$, violating the constraint. Therefore, a feasible MF-SREE assignment requires the two paths to share no interior (relay) nodes, yielding $L_j \leq 1$ at every relay. By the terminal-isolation preprocessing, no path can use any terminal node as an interior relay. Consequently, the destination d_k of flow k has load $L_{d_k} = 0$ in every feasible assignment. For a path $\mathcal{P}_k$ with H_k hops, the $H_k - 1$ interior receivers each have $L_j = 1$ (fidelity S_1 , per-hop energy at most $\bar{\epsilon}_1$). The final receiver d_k has $L_{d_k} = 0$ (fidelity S_0 , per-hop energy at most $\bar{\epsilon}_0$). Its numerator is therefore exactly $S_1^{H_k-1} S_0$, while its denominator is at most $(H_k - 1)\bar{\epsilon}_1 + \bar{\epsilon}_0$, so

$$\text{SREE}_k(\mathcal{P}_k) \geq \bar{f}(H_k), \quad \bar{f}(H) \triangleq \frac{S_1^{H-1} S_0}{(H-1)\bar{\epsilon}_1 + \bar{\epsilon}_0}. \quad (36)$$

Edge-dependent energies leave (36) a bound rather than an identity, which is all the forward direction needs. Set the threshold $T = 2 \bar{f}(|V'|) > 0$.

Forward direction (VDP YES $\Rightarrow$ MF-SREE YES). If two internally vertex-disjoint paths $\mathcal{P}'_1, \mathcal{P}'_2$ exist in G' , then their corresponding MF-SREE paths satisfy $L_j \leq 1$ at every interior relay (each relay serves at most one flow), respecting $C_{\max} = 1$. The assignment routes both flows, so it admits the whole demand and $K_{\text{adm}} = K = 2$. Each path $\mathcal{P}_k$ with H_k hops achieves $\text{SREE}_k \geq \bar{f}(H_k) \geq \bar{f}(|V'|)$ by (36), since $H_k \leq |V'|$ for a simple path. The function $\bar{f}$ is strictly decreasing in H , since $\bar{f}(H+1)/\bar{f}(H) = S_1((H-1)\bar{\epsilon}_1 + \bar{\epsilon}_0)/(H\bar{\epsilon}_1 + \bar{\epsilon}_0) < S_1 < 1$ for $S_0, S_1 \in (0, 1)$ and $\bar{\epsilon}_1 > 0$. Only positivity and monotonicity of the energy are used here. Thus $\sum_k \text{SREE}_k \geq 2 \bar{f}(|V'|) = T$, and the assignment attains at least (K, T) lexicographically.

Converse direction (MF-SREE YES $\Rightarrow$ VDP YES). Suppose the decision instance answers YES, that is, some feasible assignment admits both flows and achieves total $\text{SREE} \geq T > 0$. The admission requirement is what carries this direction: it supplies an actual path $\mathcal{P}_k$ for each of $k \in \{1, 2\}$, whereas a threshold on the efficiency sum alone would also be met by admitting one short flow and refusing the other, since $\bar{f}$ is strictly decreasing and T is set at the longest admissible hop count. Given both paths, feasibility requires $L_j \leq C_{\max} = 1$ at every relay node j . If the two paths share any interior relay node j , then $L_j \geq 2 > C_{\max} = 1$, violating the service-availability constraint. Hence, the two paths must use disjoint interior relay sets. Since the graph is $G = G'$ with identical vertex and edge sets, the paths $\mathcal{P}_1, \mathcal{P}_2$ constitute internally vertex-disjoint paths in G' , solving the VDP instance.

Since VDP on directed graphs is NP-complete for $K = 2$ [16] and the reduction is polynomial-time, MF-SREE is NP-hard. $\square$

A.3. Surrogate Approximation Quality

Proposition 10 (Empirical Surrogate Gap). *For a single source-destination pair, let $\mathcal{P}_S = \arg \min_{\mathcal{P} \in \Pi(s,d)} \sum_{(i,j) \in \mathcal{P}} E_{ij}/S(\gamma_{ij})$ denote the surrogate-optimal path and $\mathcal{P}^* = \arg \max_{\mathcal{P} \in \Pi(s,d)} \text{SREE}(\mathcal{P})$ the true-SREE-optimal path. Let $\mathcal{P}^{\text{enum}}$ denote the SREE-maximizing member of the bounded path enumeration specified in the proof below, which is the reference actually evaluated here. Across all tested network sizes ($N \in \{10, 20, 50, 100\}$), the mean single-flow SREE gap ($\text{SREE}(\mathcal{P}^{\text{enum}}) - \text{SREE}(\mathcal{P}_S)/\text{SREE}(\mathcal{P}^{\text{enum}})$) remains bounded below 2.6% (ranging from 1.30% at $N = 10$ to 2.53% at $N = 20$, 1.75% at $N = 50$, and 1.31% at $N = 100$). At $N = 50$, $\mathcal{P}_S$ coincides with $\mathcal{P}^{\text{enum}}$ in 81.0% of randomly sampled source-destination pairs (and 90.7% at $N = 10$); across non-coincident paths, the worst-case pair gap is 30.5%.*

Empirical basis. For each $N \in \{10, 20, 50, 100\}$, source-destination pairs are drawn from independently generated topologies whose nodes carry identical compute capability, at most ten pairs per topology. Sampling continues until at least one hundred pairs accumulate, which yields 108 pairs at $N = 10$ and 100 pairs at each larger size. For every pair, $\mathcal{P}_S$ is obtained from the additive surrogate weight $E_{ij}/S(\gamma_{ij})$ by Dijkstra's algorithm. The reference $\mathcal{P}^{\text{enum}}$ is the SREE-maximizing member of a loopless k -shortest-path enumeration ($k \leq 200$), merged across three edge weightings: geometric distance, hop count, and the surrogate weight itself. Both paths are then scored under the identical per-path SREE, and the per-pair gap is averaged over the sample. Coincidence is recorded when the two SREE values agree to a relative tolerance of 10^{-6} . $\square$

Remark 4 (Worst-Case Reference). *A loose worst-case bound $\text{SREE}(\mathcal{P}^*)/\text{SREE}(\mathcal{P}_S) \leq 1/S_{\min}^{H_S+1}$ can be derived by bounding the numerator as $\prod_h S_h^* \leq 1$ and the denominator as $\sum_h E_h \geq S_{\min} \sum_h w_h$, where H_S is the hop count of $\mathcal{P}_S$. However, this bound is vacuously large for typical FANET parameters (e.g., $37\times$ at $H_S = 2$ vs. the measured mean gap of at most 2.53%), because it uses $S_{\min}$ for all hops whereas most links have $S \in [0.7, 0.9]$. The empirical gap (Proposition 10) is therefore the operative metric for algorithm design.*

A.4. S-ACS Convergence

Proof of Theorem 6 (S-ACS Per-Flow Convergence).

Proof. The proof adapts the MMAS convergence result of Stützle and Dorigo [39].

Step 1 (Positive construction probability). With pheromone bounds $\tau_{\min} > 0$ and heuristic $\eta_{ij} > 0$ for all connected edges, the transition probability (19) satisfies, for any feasible next-hop j ,

$$p_{ij} = \frac{[\tau_{ij}]^{\alpha_p} [\eta_{ij}]^{\beta_p}}{\sum_{\ell \in \mathcal{N}_i} [\tau_{i\ell}]^{\alpha_p} [\eta_{i\ell}]^{\beta_p}} \geq \frac{\tau_{\min}^{\alpha_p} \eta_{\min}^{\beta_p}}{|\mathcal{N}_i| \cdot \tau_{\max}^{\alpha_p} \eta_{\max}^{\beta_p}} > 0, \quad (37)$$

where $\eta_{\min} = \min_{(i,j)} \eta_{ij} > 0$ and $\eta_{\max} = \max_{(i,j)} \eta_{ij}$. For the SREE-optimal path $\mathcal{P}^* = (n_0, \dots, n_{H^*})$ with H^* hops, the probability that a single ant constructs $\mathcal{P}^*$ is

$$p^* = \prod_{h=0}^{H^*-1} p_{n_h, n_{h+1}} \geq \left(\frac{\tau_{\min}^{\alpha_p} \eta_{\min}^{\beta_p}}{N_{\max} \cdot \tau_{\max}^{\alpha_p} \eta_{\max}^{\beta_p}} \right)^{H^*} > 0, \quad (38)$$

where $N_{\max} = \max_i |\mathcal{N}_i|$ is the maximum node degree.

Step 2 (Discovery over $I_{\max}$ iterations). With M ants per iteration, the probability that at least one ant constructs $\mathcal{P}^*$ in a single iteration is $\hat{p} \geq 1 - (1 - p^*)^M > 0$. Over $I_{\max}$ iterations, the probability that $\mathcal{P}^*$ is never constructed is bounded by $(1 - \hat{p})^{I_{\max}} \rightarrow 0$ as $I_{\max} \rightarrow \infty$.

Step 3 (Best-so-far monotonicity). S-ACS maintains the best-so-far path $\mathcal{P}_k^{\text{bsf}}$ with $\text{SREE}(\mathcal{P}_k^{\text{bsf}}) = \max_{i \leq I_{\max}} \text{SREE}(\mathcal{P}_i^{\text{best}})$. Since the best-so-far SREE is monotonically non-decreasing and the path set $\Pi(s_k, d_k)$ is finite, once $\mathcal{P}_k^{\text{opt}}$ is discovered, it is retained permanently.

Step 4 (Conclusion). For any $\delta > 0$, choosing $I_0 = \lceil \log(\delta)/\log(1 - \hat{p}) \rceil$ guarantees $\Pr[\mathcal{P}_k^{\text{bsf}} \neq \mathcal{P}_k^{\text{opt}}] < \delta$ for all $I_{\max} \geq I_0$, establishing the convergence in probability claimed in (32). The congestion-dependent evaporation (22)–(23) preserves the pheromone bounds $[\tau_{\min}, \tau_{\max}]$ after each clipping step. Thus, the positive construction probability (38) is maintained at every iteration. This consistency aligns with the MMAS convergence theory for time-dependent evaporation rates [39]. $\square$

A.5. Per-Commodity Dinkelbach Upper Bound

Proof of Proposition 1 (Per-Commodity SREE Upper Bound).

Proof. The proof proceeds in two parts: validity of the per-commodity bound, and convergence of the Dinkelbach iteration for each per-flow subproblem.

Part 1 (Bound validity). Throughout, MF-SREE* denotes the second (efficiency) key of the lexicographic optimum (17a), namely the summed per-flow SREE of an optimal assignment. Consider such an assignment, writing $\mathcal{K}_{\text{adm}}^*$ for the flows it admits and $\mathbf{P}^* = (\mathcal{P}_k^*)_{k \in \mathcal{K}_{\text{adm}}^*}$ for their paths. The relay load at each node j under this assignment is $L_j^* = \sum_{k \in \mathcal{K}_{\text{adm}}^*} \mathbf{1}[j \in \mathcal{P}_k^* \setminus \{s_k, d_k\}] \geq 0$. For each admitted commodity $k \in \mathcal{K}_{\text{adm}}^*$, the per-hop effective SNR satisfies $\gamma_{ij}^{\text{eff}}(L_j^*) \leq \gamma_{ij}^{(0)}$ because $10 \log_{10}(1 + L_j^*) \geq 0$. Since $S(\cdot)$ is non-decreasing and the energy model yields $E_{ij}(L_j^*) \geq E_{ij}(0)$ (reduced effective capacity raises transmission time), we obtain

$$\begin{aligned} \text{SREE}(\mathcal{P}_k^*; \mathbf{L}^*) &= \frac{S_{\text{e2e}}(\mathcal{P}_k^*; \mathbf{L}^*)}{E_{\text{total}}(\mathcal{P}_k^*; \mathbf{L}^*)} \\ &\leq \frac{S_{\text{e2e}}^{(0)}(\mathcal{P}_k^*)}{E_{\text{total}}^{(0)}(\mathcal{P}_k^*)} \leq \text{SREE}_k^{(0)}, \end{aligned} \quad (39)$$

where the first inequality follows from fidelity degradation and energy increase under load, and the second from

the optimality of $\text{SREE}_k^{(0)}$ over all paths at zero load. Summing over admitted commodities yields $\text{MF-SREE}^* \leq \sum_{k \in \mathcal{K}_{\text{adm}}^*} \text{SREE}_k^{(0)}$. Since each omitted term $\text{SREE}_k^{(0)}$ is positive, extending the sum to all K commodities loosens the bound to establish (28). Refused flows contribute zero to the efficiency key while zero-load terms are retained, certifying the efficiency key alone without bounding admitted flow counts.

Part 2 (Computability). To establish the analytical foundation for (28), we first examine the classical Dinkelbach fixed-point characterization of the fractional objective. The per-flow bound $\text{SREE}_k^{(0)}$ can be derived via iterative fractional programming. In practice, we compute this quantity directly and exactly via the certified label-setting dynamic program presented in Part 3. Each $\text{SREE}_k^{(0)}$ is a single-ratio fractional program over $\Pi(s_k, d_k)$. Since $E_{\text{total}}^{(0)}(\mathcal{P}) > 0$ and $|\Pi(s_k, d_k)|$ is finite, the Dinkelbach iteration [12, 38] converges exactly to $\text{SREE}_k^{(0)}$ in at most $|\Pi(s_k, d_k)|$ steps, provided the inner subproblem $\max_{\mathcal{P}} [S_{\text{e2e}}^{(0)}(\mathcal{P}) - y \cdot E_{\text{total}}^{(0)}(\mathcal{P})]$ is solved to global optimality at each step.

Part 3 (Exact computation and certification). To compute $\text{SREE}_k^{(0)}$ exactly without candidate-set truncation or heuristic approximations, we propose a certified Pareto label-setting dynamic program. Solving the single-flow SREE optimization directly is non-trivial due to the non-decomposable ratio structure. The proposed dynamic program leverages Pareto label pruning and admissible reverse-Dijkstra bounding. It finds the globally optimal path $\mathcal{P}_k^*$ while providing an exact certificate of optimality. Specifically, the exactness and computational tightness of the proposed algorithm rely on the following structural properties:

- (i) The walk relaxation over the candidate space is tight. Deleting any cycle from a walk trajectory strictly decreases the zero-load energy expenditure $E_{\text{total}}^{(0)}$ while strictly increasing the end-to-end semantic fidelity $S_{\text{e2e}}^{(0)}$ (since per-hop fidelity satisfies $S_{ij} \leq S_{\text{max}} < 1$). The optimal trajectory over all valid walks is guaranteed to be a simple path, eliminating infinite loops.
- (ii) To track candidate path prefixes efficiently, each path prefix $\mathcal{P}_u$ originating at source s_k and terminating at node u is represented by a two-dimensional label vector $\mathbf{l} = (E_{\mathbf{l}}, \Lambda_{\mathbf{l}})$, where

$$E_{\mathbf{l}} = \sum_{(i,j) \in \mathcal{P}_u} E_{ij}^{(0)}, \quad \Lambda_{\mathbf{l}} = \sum_{(i,j) \in \mathcal{P}_u} \ln S_{ij}^{(0)}. \quad (40)$$

Let $\mathbf{l}_1$ and $\mathbf{l}_2$ denote two distinct labels at node u . Label $\mathbf{l}_1$ is said to *dominate* $\mathbf{l}_2$ if $E_{\mathbf{l}_1} \leq E_{\mathbf{l}_2}$ and $\Lambda_{\mathbf{l}_1} \geq \Lambda_{\mathbf{l}_2}$, with at least one strict inequality. Because Pareto dominance is invariant under arbitrary common suffix extensions to destination d_k , any dominated label $\mathbf{l}_2$ can be safely pruned without sacrificing global optimality. In particular, any revisiting walk is strictly dominated by its corresponding cycle-free prefix.

- (iii) To accelerate state-space exploration, a reverse-graph Dijkstra search is executed prior to search. We compute $\underline{E}(u)$ and $\bar{\Lambda}(u)$, representing the exact minimum residual energy and maximum residual log-fidelity from node u to destination d_k , respectively. For any label $\mathbf{l} = (E_{\mathbf{l}}, \Lambda_{\mathbf{l}})$ at node u , an admissible upper bound on the achievable end-to-end SREE is calculated as

$$\text{SREE}_k^{(0)}(\mathbf{l}) \leq \frac{\exp(\Lambda_{\mathbf{l}} + \bar{\Lambda}(u))}{E_{\mathbf{l}} + \underline{E}(u)}. \quad (41)$$

By expanding labels in decreasing order of their admissible upper bounds (41), the dynamic program prioritizes high-promise trajectories. The search terminates as soon as the maximum upper bound among all active labels falls below or equals the current incumbent SREE value y_{inc} . This condition yields the exact per-flow optimum $\text{SREE}_k^{(0)} = y_{\text{inc}}$ alongside a formal optimality certificate. As a result, the formal guarantee established in Part 1 directly applies to the computed quantity, ensuring that (28) holds as a theorem across all evaluated instances. The dynamic program certifies $B_{\text{exact}} = \sum_k \max_{\mathcal{P} \in \Pi_k} \text{SREE}_k^{(0)}(\mathcal{P})$ for the relaxed problem. In contrast, the relation $B_{\text{exact}} \geq \text{MF-SREE}^*$ rests on the load-monotonicity arguments in Part 1. The certificate therefore closes the computational gap of the uncoupled relaxation only. It does not certify the coupled MF-SREE optimum, whose distance from B_{exact} grows with the relay-load profile. Therefore, the distance between a heuristic solution and B_{exact} is a zero-load-ceiling gap rather than an optimality gap. This completes the proof. $\square$

A.6. Gap Decomposition

Proof of Theorem 4 (Gap Decomposition).

Proof. The decomposition (31) is an algebraic identity obtained by inserting the intermediate quantity $Z_{\text{IP}}(\hat{\mathbf{l}})$:

$$Z_{\text{LP}}(\hat{\mathbf{l}}) - Z_{\text{ACO}} = (Z_{\text{LP}}(\hat{\mathbf{l}}) - Z_{\text{IP}}(\hat{\mathbf{l}})) + (Z_{\text{IP}}(\hat{\mathbf{l}}) - Z_{\text{ACO}}). \quad (42)$$

It remains to prove that both components are non-negative.

Integrity gap ($Z_{\text{LP}}(\hat{\mathbf{l}}) - Z_{\text{IP}}(\hat{\mathbf{l}}) \geq 0$). Any feasible integer solution $\hat{\mathbf{x}} \in \{0, 1\}^{|\mathcal{E}| \times K}$ trivially satisfies all constraints of the LP relaxation (30). This holds because $\{0, 1\} \subset [0, 1]$, while the flow-conservation and service-availability constraints are identical for integer and fractional variables. Consequently, the LP feasible set contains all integer-feasible solutions, and maximizing over a superset yields $Z_{\text{LP}}(\hat{\mathbf{l}}) \geq Z_{\text{IP}}(\hat{\mathbf{l}})$.

Heuristic gap ($Z_{\text{IP}}(\hat{\mathbf{l}}) - Z_{\text{ACO}} \geq 0$). Both programs are posed on the common commodity set $\mathcal{K}_{\text{adm}}^{\text{ACO}}$ admitted by S-ACS MF per Theorem 4, ensuring identical flow-conservation dimensions. Setting $x_{ij}^{\pi_k} = 1$ if $(i, j) \in \mathcal{P}_{\pi_k}^{\text{ACO}}$ and 0 otherwise yields an integer-feasible solution satisfying flow conservation. Service availability is satisfied because $r_j^{\pi_k} = \mathbf{1}[j \in \mathcal{P}_{\pi_k}^{\text{ACO}} \setminus \{s_{\pi_k}, d_{\pi_k}\}]$ and $\sum_{k'} r_j^{\pi_{k'}} = L_j = \sum_{k \in \mathcal{K}_{\text{adm}}^{\text{ACO}}} \mathbf{1}[j \in \mathcal{P}_k \setminus \{s_k, d_k\}] \leq C_{\text{max}}$. Under fixed-load

coefficients $c_{ij}(\hat{L}_j)$ in (29), this assignment attains $Z_{ACO} = \sum_k \sum_{(i,j) \in \mathcal{P}_k^{ACO}} c_{ij}(\hat{L}_j)$. Since $Z_{IP}(\hat{\mathbf{L}})$ optimizes the integer problem under identical constraints, $Z_{IP}(\hat{\mathbf{L}}) \geq Z_{ACO}$.

Both terms are non-negative, completing the decomposition. $\square$

A.7. Proof of Proposition 5 (Single-Commodity LP Exactness)

Proof. When $K = 1$, the fixed-load LP (30) reduces to a single-commodity network flow problem. The service-availability constraint (30c) becomes $r_j^1 \leq C_{\max}$ for each $j \in \mathcal{V}$, where $r_j^1 = \sum_{i:(i,j) \in \mathcal{E}} x_{ij}^1 \cdot \mathbf{1}[j \neq s_1] \cdot \mathbf{1}[j \neq d_1]$ is the fractional relay usage at node j . This constraint is non-trivial only at $j \in \mathcal{V} \setminus \{s_1, d_1\}$, since the indicators force $r_j^1 = 0$ at the endpoints.

We apply the standard node-splitting reduction [3]. For each relay node $j \in \mathcal{V} \setminus \{s_1, d_1\}$, replace j by two nodes ($j_{\text{in}}, j_{\text{out}}$) connected by an internal arc of service availability $C_{\max}$, with all incoming edges redirected to j_{in} and all outgoing edges redirected from j_{out} . The source s_1 and destination d_1 are left unsplit. This yields an augmented directed graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ in which every service-availability constraint acts on an edge rather than on a node.

The constraint matrix of the resulting LP on $\mathcal{G}'$ is the node-edge incidence matrix $A' \in \{0, \pm 1\}^{|\mathcal{V}'| \times |\mathcal{E}'|}$ of a directed graph. Since every column of A' contains exactly one +1 entry and one -1 entry, the matrix is totally unimodular (TU) [23]. The right-hand side of the LP is integer: the flow-conservation demands are $b_i \in \{+1, 0, -1\}$, and the arc capacities are $C_{\max} \in \mathbb{Z}_+$ (internal arcs) or 1 (edge bounds). By the Hoffman-Kruskal theorem, every extreme point of the LP polyhedron defined by $A' \mathbf{x} = \mathbf{b}$, $\mathbf{0} \leq \mathbf{x} \leq \mathbf{u}$, with $\mathbf{x} \in \mathbb{R}^{|\mathcal{E}'|}$, $\mathbf{b} \in \mathbb{Z}^{|\mathcal{V}'|}$, and $\mathbf{u} \in \mathbb{Z}^{|\mathcal{E}'|}$, is integer. Since the LP optimum is attained at an extreme point, $Z_{LP}(\hat{\mathbf{L}}) = Z_{IP}(\hat{\mathbf{L}})$.

It remains to treat the case $c_{ij}(\hat{L}_j) < 0$, which is the only case that arises: by (29), every objective coefficient obeys $c_{ij}(\hat{L}_j) = \ln S(\gamma_{ij}(\hat{L}_j)) - \lambda E_{ij}(\hat{L}_j) \leq -\lambda E_{ij}(\hat{L}_j) < 0$, because $S(\cdot) \in (0, 1]$ renders the log-fidelity term non-positive while $\lambda > 0$ and $E_{ij}(\hat{L}_j) > 0$. Let $\mathbf{x}^*$ denote an optimal extreme point, which is integer by the argument above. A maximizer exists because the feasible set is closed and bounded by the unit box (30d). Under the unit demands $b_i^1 \in \{+1, 0, -1\}$, flow decomposition [3] writes $\mathbf{x}^*$ as a single s_1 - d_1 path superposed with a finite family of arc-disjoint directed cycles. Suppose a cycle C lies in its support. Zeroing the arcs of C preserves flow conservation and cannot increase any r_j^1 , so feasibility is retained. Yet the objective changes by $-\sum_{(i,j) \in C} c_{ij}(\hat{L}_j) > 0$ precisely because the coefficients are negative. This contradicts optimality; hence no cycle survives, and $\mathbf{x}^*$ is the incidence vector of a simple s_1 - d_1 path. Were the coefficients along some cycle non-negative, that cycle could be saturated at no cost. In that scenario, the integer optimum attaining $Z_{LP}(\hat{\mathbf{L}}) = Z_{IP}(\hat{\mathbf{L}})$ would encode a circulation rather than a feasible single-commodity route. $\square$

References

- [1] 3GPP, 2018. Study on Enhanced LTE Support for Aerial Vehicles. Technical Report TR 36.777 V15.0.0. 3rd Generation Partnership Project (3GPP).
- [2] Aggarwal, S., Zaman, M.A.u., Bastopcu, M., Başar, T., 2026. Semantic communication in multitask dynamic games: A mean field perspective. *IEEE Transactions on Automatic Control* 71, 49–64. doi:10.1109/TAC.2025.3583625.
- [3] Ahuja, R.K., Magnanti, T.L., Orlin, J.B., 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall.
- [4] Alam, M.M., Moh, S., 2024. Joint trajectory control, frequency allocation, and routing for UAV swarm networks: A multi-agent deep reinforcement learning approach. *IEEE Trans. Mobile Comput.* 23, 11989–12005. doi:10.1109/TMC.2024.3403890.
- [5] Almasan, P., Suárez-Varela, J., Rusek, K., Barlet-Ros, P., Cabellos-Aparicio, A., 2022. Deep reinforcement learning meets graph neural networks: Exploring a routing optimization use case. *Computer Communications* 196, 184–194. doi:10.1016/j.comcom.2022.09.029.
- [6] Bernárdez, G., Suárez-Varela, J., López, A., Wu, B., Xiao, S., Cheng, X., Barlet-Ros, P., Cabellos-Aparicio, A., 2021. Is machine learning ready for traffic engineering optimization?, in: *Proc. IEEE 29th Int. Conf. Network Protocols (ICNP)*, pp. 1–11. doi:10.1109/ICNP52444.2021.9651930.
- [7] Bian, C., Shao, Y., Gunduz, D., 2025. A deep joint source-channel coding scheme for hybrid mobile multi-hop networks. *IEEE J. Sel. Areas Commun.* 43, 2543–2559. doi:10.1109/JSAC.2025.3559134.
- [8] Boursoulatz, E., Kurka, D.B., Gündüz, D., 2019. Deep joint source-channel coding for wireless image transmission. *IEEE Trans. Cogn. Commun. Netw.* 5, 567–579. doi:10.1109/TCCN.2019.2919300.
- [9] Chen, L., Gong, J., 2024. Multi-source scheduling and resource allocation for age-of-semantic-importance optimization in status update systems, in: *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, pp. 1–6. doi:10.1109/WCNC57260.2024.10570932.
- [10] Colas, C., Sigaud, O., Oudeyer, P.Y., 2018. How many random seeds? statistical power analysis in deep reinforcement learning experiments. *arXiv preprint arXiv:1806.08295*. doi:10.48550/arXiv.1806.08295, arXiv:1806.08295.
- [11] Ding, R., Chen, J., Wu, W., Liu, J., Gao, F., Shen, S., 2022. Packet routing in dynamic multi-hop UAV relay network: A multi-agent learning approach. *IEEE Trans. Veh. Technol.* 71, 10059–10072. doi:10.1109/TVT.2022.3182335.
- [12] Dinkelbach, W., 1967. On nonlinear fractional programming. *Management Science* 13, 492–498. doi:10.1287/mnsc.13.7.492.
- [13] Dorigo, M., Gambardella, L.M., 1997. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Trans. Evol. Comput.* 1, 53–66. doi:10.1109/4235.585892.
- [14] Erfaniantaghvayi, A., Zhao, X., Chan, M., Verma, R., Swami, A., Segarra, S., 2024. Ant backpressure routing for wireless multi-hop networks with mixed traffic patterns, in: *Proc. IEEE Military Commun. Conf. (MILCOM)*, pp. 1174–1179. doi:10.1109/MILCOM61039.2024.10774011.
- [15] Even, S., Itai, A., Shamir, A., 1976. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing* 5, 691–703. doi:10.1137/0205048.
- [16] Fortune, S., Hopcroft, J., Wyllie, J., 1980. The directed subgraph homeomorphism problem. *Theoretical Computer Science* 10, 111–121. doi:10.1016/0304-3975(80)90009-2.
- [17] Gabriel, S.A., Bernstein, D., 1997. The traffic equilibrium problem with nonadditive path costs. *Transportation Science* 31, 337–348. doi:10.1287/trsc.31.4.337.
- [18] Gao, R., Xu, Y., Li, H., Zhang, Q., Yang, Z., 2024. Semantic-aware jointed coding and routing design in large-scale satellite networks: A deep learning approach. *IEEE/ACM Transactions on Networking* 32, 5415–5429. doi:10.1109/TNET.2024.3464540.
- [19] Gunduz, D., Qin, Z., Aguerri, I.E., Dhillon, H.S., Yang, Z., Yener, A., Wong, K.K., Chae, C.B., 2023. Beyond transmitting bits: Context, semantics, and task-oriented communications. *IEEE J. Sel. Areas Commun.* 41, 5–41. doi:10.1109/JSAC.2022.3223408.

- [20] Guo, B., Xiong, Z., Wang, B., Quek, T.Q.S., Han, Z., 2024. Semantic communication-aware end-to-end routing in large-scale leo satellite networks, in: 2024 IEEE International Conference on Metaverse Computing, Networking, and Applications (MetaCom), pp. 137–142. doi:10.1109/MetaCom62920.2024.00033.
- [21] Hamilton, W.L., Ying, Z., Leskovec, J., 2017. Inductive representation learning on large graphs, in: Advances in Neural Information Processing Systems (NIPS), pp. 1024–1034. arXiv:1706.02216.
- [22] He, H., Sun, K., Chen, S., Jiang, X., Zhu, R., Yang, J., 2024. Deep reinforcement learning based routing for non-cooperative multi-flow games in dynamic AANETs. IEEE Trans. Veh. Technol. 73, 19495–19510. doi:10.1109/TVT.2024.3440949.
- [23] Hoffman, A.J., Kruskal, J.B., 1956. Integral boundary points of convex polyhedra, in: Linear Inequalities and Related Systems. Princeton Univ. Press. volume 38 of *Annals of Mathematics Studies*, pp. 223–246.
- [24] Holm, S., 1979. A simple sequentially rejective multiple test procedure. Scandinavian Journal of Statistics 6, 65–70.
- [25] Hu, H., Zhu, X., Zhou, F., Wu, W., Hu, R.Q., Zhu, H., 2025. Resource allocation for multi-modal semantic communication in uav collaborative networks. IEEE Transactions on Communications 73, 7599–7616. doi:10.1109/TCOMM.2025.3552303.
- [26] Huang, X., Wang, C., Tian, X., Li, Z., Zhao, C., Xiao, M., 2025. A multi-hop semantic communication system enhanced by semantic importance. IEEE Access 13, 140685–140693. doi:10.1109/ACCESS.2025.3593341.
- [27] Javadpour, A., Nafei, A., Ja'fari, F., Pinto, P., Zhang, W., Sangaiah, A.K., 2023a. An intelligent energy-efficient approach for managing IoE tasks in cloud platforms. Journal of Ambient Intelligence and Humanized Computing 14, 3963–3979. doi:10.1007/s12652-022-04464-x.
- [28] Javadpour, A., Sangaiah, A.K., Pinto, P., Ja'fari, F., Zhang, W., Abadi, A.M.H., Ahmadi, H., 2023b. An energy-optimized embedded load balancing using DVFS computing in cloud data centers. Computer Communications 197, 255–266. doi:10.1016/j.comcom.2022.10.019.
- [29] Kim, B., Kong, J.H., Moore, T.J., Dagefu, F.T., 2025. Deep reinforcement learning for multi-flow routing in heterogeneous wireless networks. arXiv preprint arXiv:2511.02030 URL: <https://arxiv.org/abs/2511.02030>, doi:10.48550/arXiv.2511.02030, arXiv:2511.02030. preprint.
- [30] Li, T., Li, C., Yang, C., Shao, J., Zhang, Y., Pang, L., Chang, L., Yang, L., Han, Z., 2022. A mean field game-theoretic cross-layer optimization for multi-hop swarm UAV communications. J. Commun. Netw. 24, 68–82. doi:10.23919/jcn.2021.000035.
- [31] Li, X., Wang, H., Yi, S., Yao, X., Zhu, F., Zhai, L., 2018. Redundancy-guaranteed and receiving-constrained disaster backup in cloud data center network. IEEE Access 6, 47666–47681. doi:10.1109/ACCESS.2018.2859427.
- [32] Liang, B., Haas, Z.J., 2003. Predictive distance-based mobility management for multidimensional PCS networks. IEEE/ACM Trans. Netw. 11, 718–732. doi:10.1109/TNET.2003.815301.
- [33] Liu, H., Yang, Z., Zhao, N., Gu, Y., Yuen, C., 2024. Interference-aware multihop routing in UAV networks: A harmonic-function-based potential field approach. IEEE Internet Things J. 11, 19406–19420. doi:10.1109/JIOT.2024.3366580.
- [34] Martins, E.Q.V., 1984. On a multicriteria shortest path problem. European Journal of Operational Research 16, 236–245. doi:10.1016/0377-2217(84)90077-8.
- [35] Paul, R., Cohen, K., Kedar, G., 2024. Multi-flow transmission in wireless interference networks: A convergent graph learning approach. IEEE Trans. Wireless Commun. 23, 3691–3705. doi:10.1109/TWC.2023.3310353.
- [36] Rashid, T., Farquhar, G., Peng, B., Whiteson, S., 2020. Weighted QMIX: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning, in: Proc. 34th Conf. Neural Information Processing Systems (NeurIPS), pp. 10199–10210. arXiv:2006.10800.
- [37] Rashid, T., Samvelyan, M., de Witt, C.S., Farquhar, G., Foerster, J., Whiteson, S., 2018. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning, in: Proc. Int. Conf. Mach. Learn. (ICML), pp. 4295–4304.
- [38] Schaible, S., 1976. Fractional programming. II, on Dinkelbach's algorithm. Management Science 22, 868–873. doi:10.1287/mnsc.22.8.868.
- [39] Stützle, T., Dorigo, M., 2002. A short convergence proof for a class of ant colony optimization algorithms. IEEE Trans. Evol. Comput. 6, 358–365. doi:10.1109/TEVC.2002.802444.
- [40] Tang, B., Huang, L., Li, Q., Pandharipande, A., Ge, X., 2024. Cooperative semantic communication with on-demand semantic forwarding. IEEE Open Journal of the Communications Society 5, 349–363. doi:10.1109/OJCOMS.2023.3344882.
- [41] Wang, X., Fu, L., Cheng, N., Sun, R., Luan, T.H., Quan, W., Al-dubaikhy, K., 2022. Joint flying relay location and routing optimization for 6G UAV-IoT networks: A graph neural network-based approach. Remote Sensing 14, 4377. doi:10.3390/rs14174377.
- [42] Wu, H., Xu, Y., Cao, S., Liu, J., Takakura, H., Shiratori, N., 2025. Sleeping multi-armed bandit-based path selection in space-ground semantic communication networks, in: 2025 IEEE Wireless Communications and Networking Conference (WCNC), pp. 1–6. doi:10.1109/WCNC61545.2025.10978165.
- [43] Xie, H., Qin, Z., Li, G.Y., Juang, B.H., 2021. Deep learning enabled semantic communication systems. IEEE Trans. Signal Process. 69, 2663–2675. doi:10.1109/TSP.2021.3071210.
- [44] Yin, Y., Huang, L., Li, Q., Tang, B., Pandharipande, A., Ge, X., 2023. Semantic communication on multi-hop concatenated relay networks, in: Proc. IEEE/CIC Int. Conf. Commun. China (ICCC), pp. 1–6. doi:10.1109/ICCC57788.2023.10233446.
- [45] Yu, C., Velu, A., Vinitsky, E., Gao, J., Wang, Y., Bayen, A., Wu, Y., 2022. The surprising effectiveness of PPO in cooperative multi-agent games, in: Advances in Neural Information Processing Systems (NeurIPS), pp. 24611–24624.
- [46] Zhang, G., Hu, Q., Cai, Y., Yu, G., 2024. Alleviating distortion accumulation in multi-hop semantic communication. IEEE Commun. Lett. 28, 308–312. doi:10.1109/LCOMM.2023.3339776.