Computational Engineering and Technology Innovations

www.ceti.reapress.com

Comput. Eng. Technol. Innov. Vol. 3, No. 4 (2026) 286–307.

Paper Type: Original Article

Topology Aware Adaptive Routing for Drone Assisted Emergency Communication

Seyyed Ahmad Edalatpanah[1,*], Markus Hess[2]

[1] Department of Applied Mathematics, Ayandegan Institute of Higher Education, Tonekabon, Iran; s.a.edalatpanah@aihe.ac.ir.
[2] Department of System Dynamics and Friction Physics, Faculty V-Mechanical Engineering and Transport Systems, Technische Universität Berlin, Berlin, Germany; markus-hess@tu-berlin.de.

Citation:

Received: 21 October 2025 *Revised: 26 November 2025* *Accepted: 03 May 2026*	Edalatpanah, S. A., & Hess, M. (2026). Topology aware adaptive routing for drone assisted emergency communication. *Computational Engineering and Technology Innovations,* 3(4), 286-307.

Abstract

Unmanned Aerial Vehicle (UAV) relays can provide temporary communication paths when terrestrial infrastructure is damaged, but a connected topology does not ensure that a message can traverse a route before its links expire. This paper develops the Topology Aware Adaptive Routing Algorithm (TAARA) as a network-layer design for deadline-bearing emergency traffic. The formulation separates route admissibility from route preference. Admissibility accounts for cumulative forwarding time, contact expiry, information freshness and a reserved radio-energy allowance at every participating node. A dimensionless additive cost then combines transaction success, service time, radio energy and short contact horizons. Source-pinned forwarding, explicit expiry rules and hysteresis define how selected routes are used and replaced. We derive conditional feasibility and reliability bounds, a route-ranking sensitivity result and a switching-error bound. A fully specified six-node example shows why the lowest-cost path can be inadmissible, why checking each contact in isolation is insufficient and how reliability weights alter the selected route. A deterministic switching sequence illustrates the effects of hysteresis without implying packet-delivery gains. The numerical material verifies the design logic; it is not an evaluation of an implemented aerial network. A matched, seed-paired evaluation protocol specifies executable comparators, failure-inclusive metrics, uncertainty analysis and the records needed for reproducible assessment. The contribution is an explicit routing specification with analytical and computational illustrations, establishing a defensible basis for subsequent network simulation and flight experiments.

Keywords: Flying ad hoc networks, Emergency communications, Adaptive routing, Temporal feasibility, Energy reserves, Reproducibility.

1 | Introduction

Emergency communication depends on whether information reaches an actionable destination while it remains useful. Damage to a terrestrial access point, loss of backhaul or obstruction of a local radio link can each interrupt that process. Aerial relays offer a means of connecting ground users to surviving gateways, but

Corresponding Author: s.a.edalatpanah@aihe.ac.ir
https://doi.org/10.48314/ceti.v3i4.97
Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

they introduce moving endpoints, intermittent contacts and limited vehicle energy. The survey by Matracia et al. [1] places routing within this wider recovery problem, alongside deployment, propagation and infrastructure restoration. These functions interact, yet they answer different questions: placing an aircraft where a ground terminal can reach it does not determine how the terminal's message will traverse the remaining relay network.

Flying Ad Hoc Networks (FANETs) make this distinction particularly consequential. Their three-dimensional mobility can change both adjacency and the duration for which an apparently usable link remains available [2]. Communication requirements therefore extend beyond conventional connectivity to include the age of topology information, link intermittency and the cost of maintaining a network view [3]. Unmanned Aerial Vehicle (UAV)-based wireless systems also couple communication to flight decisions [4], while the broader design space includes aerial base stations, relaying and cellular-connected vehicles [5]. The present paper considers a narrower task: forwarding messages among already deployed relays to a currently reachable gateway. Trajectories, radio hardware and gateway availability are external inputs to the routing decision.

A short path can be a poor emergency route for several independent reasons. Its final contact may disappear before the packet arrives; one relay may lack sufficient energy above its operational reserve; or a fresh-looking route may depend on a stale measurement. Conversely, repeatedly switching to an apparently better path can consume control airtime and interrupt forwarding. A scalar score alone cannot resolve these cases. If a battery reserve or a contact deadline is a requirement, a sufficiently favorable score on another attribute must not purchase its violation. This observation motivates a design in which hard admission conditions precede preference-based ranking.

Topology Aware Adaptive Routing Algorithm (TAARA) operationalizes that separation. It treats a link observation as evidence with an age and a declared forwarding-transaction definition, rather than as a timeless reliability value. It checks the time at which a packet would finish each successive hop, rather than checking contact duration only at the source. It accounts for both reception and transmission energy at an intermediate relay. Among admitted paths it minimizes a normalized cost, and it replaces an incumbent path only after a meaningful improvement or a feasibility failure. The route carried by a packet remains fixed during one forwarding attempt, making the loop-control obligation explicit.

The paper makes three connected contributions:

I. It specifies the information contract, admission rules and forwarding behavior needed to turn topology awareness into an inspectable routing policy.

II. It establishes elementary but useful properties of that policy, including the assumptions under which its reliability term has a probability interpretation.

III. It supplies exact numerical illustrations and a failure-inclusive evaluation protocol.

The claims concern the mathematical construction and its behavior on declared inputs. Comparative packet delivery, delay, goodput and endurance remain questions for an implemented, calibrated experiment.

This distinction shapes the manuscript throughout. No observed performance advantage is inferred from a route score, a diagram or an arithmetic example. Section 2 positions the design against established routing and aerial-network approaches. Sections 3–5 define the model, policy and analytical properties. Section 6 presents reproducible examples. Section 7 specifies the experiment needed to assess operational value. Section 8 discusses limitations, and Section 9 concludes.

2 | Related Work and Design Position

2.1 | From Mobile Routing to Temporally Constrained Aerial Forwarding

The development of FANET routing is a development of information and control assumptions, rather than a replacement of one universally inferior protocol by another. Conventional mobile ad hoc protocols provide

concrete discovery, dissemination and route-maintenance procedures. Ad hoc On-Demand Distance Vector (AODV) discovers routes on demand and uses destination sequence information [6]. Optimized Link State Routing (OLSR) maintains routing information proactively through optimized link-state dissemination [7]. Optimized Link State Routing version 2 (OLSRv2) specifies a later link-state design with explicitly defined information bases and metrics [8], while the Neighborhood Discovery Protocol (NHDP) formalizes neighborhood information and its validity [9]. Babel offers a different distance-vector architecture with a feasibility mechanism intended to control routing loops [10]. These specifications provide executable comparison targets and illuminate obligations that a new routing proposal must meet.

Topology-based FANET routing nevertheless depends strongly on mobility and information age. Wheeb et al. [11] review this relationship across routing families and mobility models. Rosati et al. [12] incorporate relative-motion information into predictive OLSR and assess the design using emulation and a small flight testbed. Their work directly precedes any claim that movement prediction alone distinguishes a new aerial routing method. Position-Monitoring-Based Hybrid Routing Protocol (PMHRP) combines position monitoring with geographic and topology information [13]. Its information architecture is relevant to TAARA, but reproducing it requires implementing its coordinator and position-update behavior, not substituting a shortest-path routine under its name.

The distinction between a static path and a time-respecting journey is established in the theory of time-varying graphs [14]. A set of edges can connect source and gateway when inspected without regard to ordering, yet fail to support sequential forwarding. TAARA uses a deliberately restricted version of this temporal perspective: all contacts must already be available, and forwarding must finish before their declared expiry. It does not search for future contact openings or optimize a store-carry-forward journey. For disruptions requiring persistent storage and later contact opportunities, the Bundle Protocol provides a relevant architectural alternative [15]. Such operation changes the service contract and should be evaluated as a separate mode.

2.2 | Link Quality and Information Requirements

The Expected Transmission Count (ETX) metric established that hop count can conceal large differences in link quality [16]. Expected transmission time and interference-sensitive mesh metrics further distinguish the jobs performed by loss, bandwidth and shared-channel effects [17]. TAARA builds on the same need to expose metric semantics. Its success variable refers to one bounded forwarding transaction under a declared retry policy; it is not automatically ETX, raw signal strength or end-to-end packet delivery. The negative logarithm converts a product of independent transaction-success probabilities into an additive path term. It does not eliminate dependence caused by interference, weather or shared queueing.

Learning-based approaches ask how routing decisions can improve from experience. Minhas et al. [18] develop reinforcement-learning and UAV-aided multipath routing for public-safety networks. A learned policy may use richer state and adapt beyond a fixed cost rule, but its training traffic, exploration cost and behavior under distribution shift become part of the comparison. TAARA is a transparent reference policy: every rejection and ranking decision is traceable to specified measurements, thresholds and weights. Transparency does not establish that it will outperform a learned policy.

More recent hybrid control also combines global coordination with local forwarding. Sun et al. [19] couple segmented Software-Defined Networking (SDN)-FANET routing with adaptive extended Kalman prediction. Their design illustrates that motion estimation, controller placement and forwarding organization are separate architectural choices. The 2024 review by Almansor et al. [20] provides a further taxonomy of routing strategies and open issues. Taken together, these studies support an evaluation against contemporary predictive or hybrid methods, provided their input requirements and control costs can be reproduced in the same environment.

2.3 | Architecture and Physical-Link Studies

Drone-assisted emergency architectures such as that of Ganesh et al. [21] address the organization of a rescue communication system. Drone-assisted device-to-device cooperation addresses a different set of transmission and cooperation choices [22]. RESTORE uses non-line-of-sight free-space optical communication and altitude selection in an emergency setting [23]. These studies are relevant to deployment and link provision, but their reported system outcomes cannot be treated as matched routing baselines in a single radio network without implementing the corresponding physical and control assumptions. A routing experiment must preserve the radio and mobility environment while varying the routing intervention.

Energy has an equally important scope distinction. Rotary-wing propulsion modeling and joint trajectory-communication optimization [24] concern the vehicle's mission-energy budget. A route that reduces radio joules may leave flight energy almost unchanged, or may indirectly increase flight time if communication causes loitering. TAARA therefore reports radio energy and total vehicle energy separately. Operational reserves come from the mission-energy model; they are not inferred from a favorable routing score.

Table 1 organizes selected milestones by the problem each made explicit. The periods overlap because the approaches coexist. The grouping is interpretive and does not assert a complete historical chronology.

Table 1. Selected developments that inform the routing specification.

Period	Development	Representative Sources	Consequence for This Design
1999–2004	Performance criteria, route maintenance and link-quality metrics	Mobile Ad Hoc Network (MANET) evaluation [25]; AODV [6]; OLSR [7]; ETX [16]; Expected Transmission Time (ETT) [17]	Specify control behavior and metric meaning alongside path selection.
2011–2016	Neighborhood validity, temporal connectivity and aerial mobility	NHDP [9]; time-varying graphs [14]; FANET surveys [2], [3]; predictive OLSR [12]	Treat information age and sequential contact use as model variables.
2019–2022	Mission energy, public-safety learning and hybrid position monitoring	Rotary-wing energy [24]; Reinforcement Learning (RL) routing [18]; PMHRP [13]	Separate flight constraints, observation requirements and policy costs.
2024–2025	Routing taxonomies and predictive hybrid control	Review [20]; adaptive SDN-FANET routing [19]	Compare architectural assumptions before comparing numerical outcomes.

Table 2 compares methods by the information they consume and the question they answer. The final column defines the comparison boundary for this paper; it is not a claim that a study failed to solve a problem outside its scope.

Table 2. Method and architecture comparison relevant to a matched evaluation.

Approach	Main Information and Mechanism	Evidence or Setting	Boundary for Comparison
AODV [6]	On-demand discovery; sequence-based route freshness	Protocol specification for mobile ad hoc networks	Implement its discovery, errors and timeouts; a hop-count path is insufficient.
OLSR and OLSRv2 [7], [8]	Proactive link-state dissemination and route computation	Defined control procedures and information bases	Version, timers and link metrics must be declared.
P-OLSR [12]	Motion information supplements link-quality routing	MAC-layer emulation and small UAV testbed	Match GPS availability, beacon timing and metric realization.
PMHRP [13]	Position monitoring with hybrid geographic and topology information	Three-dimensional UAV network study	Coordinator and position-update costs belong in the implementation.
RL public-safety routing [18]	Learning-assisted multipath decisions and cooperation	UAV-aided public-safety network study	Training, exploration and traffic assumptions differ from fixed-rule routing.
Adaptive SDN-FANET [19]	Segmented hybrid control and motion prediction	Optimized Network Engineering Tool (OPNET) simulation study	Controller reachability and prediction overhead affect transferability.
RESTORE [23]	Optical link architecture and altitude choice	Emergency communication system analysis	A physical-system comparator, requiring a separate cross-layer study.
TAARA	Fresh state, path admission, additive ranking and pinned forwarding	Analysis and deterministic illustrations in this paper	Network-level comparative performance remains to be measured.

2.4|The Question Addressed Here

The specific design question is how to make route selection auditable when deadlines, expiring contacts and mission reserves coexist. This is narrower than claiming that prior routing ignores mobility or energy. *Table 3* identifies the obligations addressed by the present construction and locates the corresponding contribution. The manuscript does not claim completeness over all FANET routing literature; the comparison is a targeted synthesis of the assumptions needed for this design.

Table 3. Design obligations and their treatment in the present paper.

Obligation	Relevant Prior Context	Treatment Here
Interpret reliability without changing the measurement unit	ETX and ETT [16], [17] model link-quality costs	Define a capped transaction and distinguish logarithmic ranking from route-success guarantees in Sections 3 and 5
Test a contact at the time it will be used	Temporal graphs [14]; predictive routing [12]	Impose cumulative completion-time inequalities and provide a counterexample in Section 6
Protect reserves independently of scalar preference	Vehicle-energy optimization [24]; hybrid routing [13], [19]	Impose per-node admission and reservation checks before scoring in Sections 3 and 4
Explain every performance denominator and failure	MANET evaluation guidance [25]; credibility studies [26], [27]	Specify a paired experiment and complete event records in Section 7

3 | Network Model and Observation Contract

3.1 | Scope and State

At source decision time t, let $G(t) = (V(t), E(t))$ be a directed graph containing an ingress relay s, aerial relays and a nonempty set of candidate gateways $\mathcal{G}(t)$. A directed edge represents an available forwarding transaction under the chosen radio stack. Reverse reachability is checked separately whenever acknowledgements or setup messages require it. Ground access, gateway backhaul status and flight trajectories are supplied by other subsystems. A gateway is eligible only if its application service and required backhaul are available; hearing its radio beacon alone does not establish end-to-end service.

Each application message has a source-scoped unique identifier, creation time, payload length and absolute application deadline. Its residual budget at route admission is $D > 0$. Setup and queueing before admission consume the original deadline; they are never removed from measured end-to-end delay. The reference service excludes intentional waiting for future contacts. Once an admitted attempt begins, its route is a simple ordered node list ending at one eligible gateway.

For each edge $i \to j$, the snapshot contains success estimate $q_{ij} \in (0,1]$, service envelope $\overline{d}_{ij} > 0$ in seconds, transmitter and receiver radio-energy envelopes $e_{ij}^{tx}, e_{ij}^{rx} \geq 0$ in joules, residual contact horizon $\tau_{ij} > 0$, observation time and expiry. A common snapshot timestamp or explicit clock-error bound is needed to compare these quantities. The service envelope includes the permitted medium-access waits, capped retries and acknowledgement processing. If channel access cannot be bounded, $\overline{d}$ is a predictive quantile and the associated feasibility statement becomes conditional on that quantile, rather than deterministic.

The transaction must be defined before estimating its success probability. If a MAC permits up to R attempts, q_{ij} is the probability that this entire transaction succeeds, including its acknowledgement rule. It cannot be substituted by the success of one radio attempt without changing the cost's meaning. Under independent identical attempts with success probability p, the familiar expression $1 - (1 - p)^R$ is available; correlated attempts need a different model. Energy and service envelopes must cover the same transaction, including a complete failed transaction, so that unsuccessful forwarding remains inside the accounting boundary.

3.2 | Estimation, Age and Uncertainty

The reference estimator records x successful transactions out of n completed transactions in a declared recent window. A smoothed estimate $(x + 1)/(n + 2)$ can prevent logarithms of zero, but smoothing is a modeling choice and does not turn absent evidence into a reliable link. Admission also requires $n \geq n_{min}$ and observation age at most T_{fresh}. An unseen edge is unmeasured and excluded in the conservative mode. A separate exploration mode may probe it using a declared control budget; exploration is then part of the policy's cost.

A Wilson [28] interval can describe binomial uncertainty if the window is approximately stationary and transactions are suitably independent. For example, even ten successes in ten trials do not establish perfect reliability. Under burst interference, overlapping windows and mobility-driven nonstationarity, a binomial interpretation is inadequate. Block-based calibration or independent held-out trajectories should then assess predictive coverage. The reference policy uses q as a ranking estimate; a reliability requirement requires an independently justified lower bound, denoted q^L, and an explicit conditioning population.

Contact horizons may come from position and velocity under a calibrated radio-range model or from a learned link-survival predictor. A distance-only construction must not be described as a measured channel law. Position uncertainty, turns, obstacles and clock error must either shorten the conservative horizon or appear as a documented uncertainty assumption. The positive guard g absorbs only the disturbances for which it has been sized; it cannot stand for all unmodeled radio behavior. Freshness, contact lifetime and transaction

success are related observations but serve different functions: Age governs whether evidence may be used, lifetime limits when the edge may be traversed, and success estimates the outcome of a declared transaction.

3.3 | Energy Allowances and Path Admissibility

Let B_i be a node's current battery energy. Before routing, a mission controller protects return or landing energy, avionics and other committed consumption in a quantity M_i. Let L_i be already reserved radio energy for other admitted traffic. The uncommitted allowance is

$$U_i = \max\{0, B_i - M_i - L_i\}. \tag{1}$$

The calculation is in joules and is renewed as reservations and flight plans change. It does not equate a battery percentage with an energy model. Ground relays and gateways can use their own finite allowances or a documented externally powered model. Uncertainty in the flight-energy budget should reduce the allowance conservatively; routing cannot certify the adequacy of an incorrect mission reserve.

For a simple path $P = (v_0 = s, v_1, \dots, v_k)$, write

$$S_r(P) = \sum_{\ell=1}^{r} \overline{d}_{v_{\ell-1}v_\ell}, \qquad r = 1, \dots, k. \tag{2}$$

The radio charge at node i includes its outgoing transmission and incoming reception on this path:

$$f_i(P) = \sum_{(i,j)\in P} e_{ij}^{tx} + \sum_{(j,i)\in P} e_{ji}^{rx}. \tag{3}$$

TAARA admits P only if all edge observations remain valid, its gateway is eligible, its hop count is within K_{max} and

$$S_k(P) \le D, \qquad S_r(P) + g \le \tau_{v_{r-1}v_r}\ (r = 1, \dots, k), \quad f_i(P) \le U_i\ (i \in P). \tag{4}$$

All horizons in *Eq. (4)* are measured from the attempt's admission time, after setup. A source must therefore age the original snapshot and recalculate residual horizons before committing an attempt. The cumulative term S_r is essential: A downstream link can expire while the packet is crossing earlier hops. The inequalities are an admission test on the declared envelopes, not a guarantee against disturbances outside them.

For multiple messages, a cached route is a candidate template rather than a reservation for unlimited traffic. Each packet or explicitly bounded batch requires fresh admission and atomic reservation at participating nodes. A batch uses the sum of its worst-case charges and an envelope that includes queue interaction. Without these reservations, several individually admissible paths could oversubscribe the same allowance. This restriction is especially relevant when multiple ingress relays act concurrently.

4 | Routing and Forwarding Specification

4.1 | Ranking Admitted Routes

After admission, preference is expressed by a dimensionless nonnegative cost. With reference scales $D_0, E_0, H > 0$ and prespecified weights $a, b, c, h \ge 0$ summing to one, define

$$w_{ij} = a\left(-\log q_{ij}\right) + b\frac{\overline{d}_{ij}}{D_0} + c\frac{e_{ij}^{tx} + e_{ij}^{rx}}{E_0} + h\max\left\{0, 1 - \frac{\tau_{ij}}{H}\right\}. \tag{5}$$

The selected route is

$$P^* \in \arg\min_{P\in\mathcal{F}(s,\mathcal{G},t,D)} C(P), \quad C(P) = \sum_{(i,j)\in P} w_{ij}, \tag{6}$$

where $\mathcal{F}$ is the family satisfying *Eq. (4)*. Ties are resolved by the lexicographic node sequence after gateway identifiers have been ordered. This rule is deterministic and does not silently introduce hop count as another objective. Positive reference scales keep units compatible, but their values still influence trade-offs. A change in E_0 has the same practical effect as changing the effective energy weight unless weights are adjusted accordingly.

The contact penalty assigns a soft preference for longer remaining horizons inside the feasible set. The hard timing test remains in *Eq. (4)*. Likewise, energy in *Eq. (5)* ranks affordable routes; it cannot compensate for violating a node's reserve. No positive residual-battery term is minimized directly, since doing so would reward depletion. The cost is a preference functional, not an estimate of goodput or total disaster-response utility. Its weights must be selected before the final comparison and reported as design parameters.

4.2 | Reference Search and Computational Boundary

For a finite validation graph, the reference search enumerates simple source-to-gateway paths up to K_{max}, prunes any prefix violating a cumulative contact or deadline condition, accumulates node charges and retains the least-cost admissible path. When a new outgoing edge is appended, the previous node's reception-plus-transmission charge becomes known and is checked. A candidate gateway's reception charge is also checked. This procedure is exact for the declared finite candidate family and provides an oracle for implementation verification.

Nonnegative edge costs alone do not make the constrained problem an ordinary shortest-path problem. Dijkstra's [29] algorithm applies to the unconstrained additive problem, or to a correctly constructed expanded state space. Here, route feasibility depends on cumulative time and node charges. An unconstrained shortest path followed by rejection can miss an admissible alternative. An operational implementation can use resource-constrained label search or a bounded candidate-path pool, but it must disclose whether it is exact. If only the first L candidates are checked, no admissible candidate means none was found within that pool; it does not establish that the network contains none.

With n nodes, unrestricted simple-path enumeration can grow factorially; the hop bound limits depth but does not remove combinatorial growth. The design therefore targets small emergency relay domains in its exact reference form. A deployed implementation must report domain size, candidates examined, peak memory and decision-time distribution. Computation that consumes a substantial fraction of a contact horizon must be included in the admission clock. This boundary makes scalability a measurable engineering question instead of an unsupported property of a short equation.

4.3 | Control State and Message Handling

Each relay maintains timestamped neighbor observations, a topology database, a reservation ledger and two duplicate caches with different keys. Link-state advertisements carry originator ID, monotone origin sequence, sample time, validity duration, directed neighbor records and current uncommitted allowance. A receiver accepts only a newer advertisement from that originator, retains each edge's sample age and relays it at most once within the configured dissemination domain. Receipt of a new advertisement does not refresh the age of its underlying measurement. Advertisement size, rate, fragmentation and flooding radius are implementation parameters that must be fixed and charged as control traffic.

The source maintains candidate, pending, active and invalid route states. A candidate becomes pending during reservation. Participating nodes atomically hold the charge in *Eq. (3)*, acknowledge the reservation and attach a lease expiry. Setup control traffic is charged against a separately budgeted control allowance. If any node refuses, times out or withdraws its gateway status, the source releases the successful holds; leases remove stranded holds when a release is lost. A pending route becomes active only after all holds are confirmed and *Eq. (4)* has been rechecked with the remaining deadline and contact horizons. An unresolved reservation is never treated as successful admission.

An active packet carries source ID, message ID, attempt ID, gateway ID, route sequence, hop cursor, absolute deadline and route expiry. Every relay verifies that its own identifier matches the current cursor and that the next hop matches the pinned sequence. It forwards only to that next hop and advances the cursor once. A malformed route, expired lease, repeated node in the path or expired deadline causes an explicit drop. The relay duplicate key includes attempt ID; the gateway application deduplicates on source and message ID. Consequently, a new source attempt can traverse a relay used by an earlier failed attempt without being suppressed as the same forwarding transaction.

The radio layer performs at most the declared retry count and transaction duration. It may acknowledge a duplicate radio frame without forwarding its payload again. On success, reserved worst-case energy is reconciled with measured consumption and unused energy is released. On failure, consumed energy remains charged, remaining reservations are released and an error is sent toward the source if a reverse control route is available. Because the error itself can be lost, the source also has a deadline-bounded transaction timeout. These operations require authenticated node identities and bounded cache retention in a real deployment; the routing specification does not itself provide a security protocol.

4.4 | Switching, Repair and Disconnection

For a new admission, compare the active route with the best admitted alternative under one fresh snapshot. A voluntary switch occurs only if

$$\hat{C}(P_{old}) - \hat{C}(P_{new}) > \eta, \qquad t - t_{last} \geq T_{dwell}, \tag{7}$$

where $\eta > 0$ is an absolute cost threshold and $T_{dwell} \geq 0$ is a minimum dwell time. A failed next hop, withdrawn gateway, exhausted reservation or timing violation invalidates the route immediately and bypasses *Eq. (7)*. A voluntary switch affects subsequent attempts; it does not rewrite an in-flight packet's route. Repair is initiated at the source with a new attempt identifier and the original application deadline.

If no route is admitted, the policy records the reason and either drops an expired message or keeps it in a bounded local queue until the next decision opportunity. Queueing consumes deadline and idle energy. The reference mode does not hand packets to mobile carriers for delivery at a future contact. A separate disruption-tolerant mode can use persistent storage, but it must declare its different deadlines, custody semantics and comparators. Partition time is part of the operational record even when no packets are generated during it.

Algorithm 1 states the source procedure at the level required for an independent implementation. The reservation and forwarding rules above are part of the algorithm, not optional implementation details.

Algorithm 1. TAARA source admission and route update.

I. Read the application deadline and a timestamped topology snapshot; reject expired messages and stale or unmeasured edges.
II. Enumerate the declared candidate family. Accumulate prefix service time and node charges; prune deadline, contact, reserve and hop violations.
III. Rank admissible paths by *Eqs. (5)* and *(6)*. If the active route is feasible, apply the improvement and dwell tests in *Eq. (7)*; otherwise select a feasible repair immediately.
IV. Acquire atomic node reservations with expiring leases. On refusal or timeout, release holds and try the next candidate while time remains.
V. Age all horizons by elapsed setup time and recheck *Eq. (4)*. Commit only a confirmed, feasible route, then launch a packet with a fixed node sequence and fresh attempt ID.
VI. Record success, failure or timeout; reconcile energy, release holds and preserve all failed attempts in the event log.

5 | Analytical Properties

5.1 | Conditional Timing and Energy Protection

Proposition 1. Suppose an admitted attempt starts at time zero, every hop completes or declares failure within its service envelope, the actual contacts persist at least until their declared horizons, clocks satisfy the guard assumptions and node charges do not exceed their reserved envelopes. If every hop succeeds, the packet finishes by D and each transaction completes at least g before its declared contact expiry. Its incremental radio consumption at each node does not exceed the reservation in *Eq. (3)*.

Proof: The first transaction completes by $\bar{d}_{v_0 v_1}$. Inductively, if transaction $r-1$ has completed by S_{r-1}, transaction r completes by $S_{r-1} + \bar{d}_{v_{r-1} v_r} = S_r$. *Eq. (4)* gives $S_r + g \le \tau_{v_{r-1} v_r}$ at each hop and $S_k \le D$. At node i, summing the bounded reception and transmission charges yields consumption at most $f_i(P) \le U_i$. Atomic reservations prevent other admitted attempts from spending the same allowance. This proves the stated conditional protections.

The proposition does not assert that radio transactions will succeed, or that an aircraft's mission-energy estimate is correct. If the service and contact inputs are quantiles, its conclusion holds only when those envelopes hold jointly. Marginal coverage values cannot simply be multiplied without a dependence model. A practical system should measure envelope exceedances directly and record any emergency reserve override as a separate operating mode, excluded from claims of reserve-preserving behavior.

5.2 | Reliability Interpretation and Its Limit

Proposition 2. For a fixed admitted path and a defined sequence of forwarding transactions, suppose q_r equals the success probability of hop r conditional on the snapshot and on execution of that transaction. If the corresponding success events can be represented as independent under that conditioning, then the following identities hold:

$$Q(P) = \prod_{r=1}^{k} q_r, \quad -\log Q(P) = \sum_{r=1}^{k} -\log q_r.$$

More generally, if q_r^L lower-bounds the success probability of hop r conditional on every preceding hop having succeeded, then

$$\Pr\{\text{route succeeds} \mid \text{snapshot}\} \ge \prod_{r=1}^{k} q_r^L. \tag{8}$$

Proof: Under independence, the probability of the intersection is the product and the logarithm identity follows directly. Without independence, the chain rule expresses route success as the product of conditional hop-success probabilities given all preceding successes. Replacing each conditional probability by its valid lower bound gives *Eq. (8)*.

If instead the bounds concern unconditional potential success events A_r for all path transactions under a common declared environment, the union bound gives

$$\Pr\left(\bigcap_{r=1}^{k} A_r\right) \ge \max\left\{0, 1 - \sum_{r=1}^{k} \left(1 - q_r^L\right)\right\}. \tag{9}$$

These are different access assumptions. Transactions downstream of a failed hop may never be observed, so their unconditional potential outcomes are not available from ordinary packet logs. *Eq. (9)* is useful only when its probability population is justified by a model or experiment. A set of marginal, selected-route estimates

does not establish *Eq. (8)*. The logarithmic score in *Eq. (5)* remains well defined as a ranking feature, but calling its exponential a certified delivery probability would require additional evidence.

5.3 | Weight Sensitivity and Hard Constraints

Proposition 3. Hold the snapshot, candidate paths and feasible family fixed. For weights of the form $a = \alpha$, $b = 0.5(1-\alpha)$, $c = 0.3(1-\alpha)$ and $h = 0.2(1-\alpha)$, each path cost is affine in $\alpha \in [0,1]$. Any two paths either tie throughout the interval, never tie or cross at a single point. The preferred path over a finite candidate family is the lower envelope of these affine functions. Varying α cannot admit a path that violates *Eq. (4)*.

Proof: For fixed P, collect its logarithmic risk into $R(P)$ and its other normalized terms into $V(P)$. Then $C_\alpha(P) = \alpha R(P) + (1-\alpha)V(P)$. The difference of two such functions is affine; its zeros have the stated alternatives. Minimizing over a finite family gives their lower envelope. Feasibility was fixed independently of these weights.

This result supports a direct sensitivity plot and identifies the exact trade-off at a route change. It does not establish a globally optimal weight. Nor does scalarization necessarily recover every Pareto-efficient point of a discrete path set. Hard thresholds can also make the feasible family change abruptly: Raising one node's allowance past a required charge may admit a path with a substantially lower cost. That discontinuity is an interpretable effect of admission, not numerical instability.

5.4 | Switching Error and Loop Control

Proposition 4. Suppose the estimated costs of an incumbent and an alternative differ from their true snapshot costs by at most ε each. A voluntary switch satisfying *Eq. (7)* implies a true cost improvement greater than $\eta - 2\varepsilon$. Therefore $\eta \geq 2\varepsilon$ excludes switches to a truly worse route under these error bounds. Over an interval of length T, voluntary switches separated by $T_{dwell} > 0$ number at most $1 + \lfloor T/T_{dwell} \rfloor$.

Proof: Subtract the two bounded estimation errors from the observed improvement; the worst-case reduction is 2ε. For the count, m separated switch times require at least $(m-1)T_{dwell}$ between first and last, which cannot exceed T.

The statement concerns costs at a snapshot, not future realized performance. Forced repairs are excluded from the switch-count bound, and the bound on ε must be justified before it can calibrate a threshold. Increasing hysteresis also delays beneficial adaptation. These competing effects motivate the ablations in Section 7.

Finally, for a correctly formed simple source route, an honest relay advances a strictly increasing cursor. The packet cannot revisit a node within that attempt. Radio retries do not advance the cursor twice, and a new attempt is accounted for separately. This is a structural loop-exclusion argument for pinned forwarding. It does not prove loop freedom for a different hop-by-hop repair implementation, nor does it eliminate end-to-end duplicate attempts.

6 | Deterministic Illustrations

6.1 | Declared Inputs and Arithmetic

The examples in this section are constructed calculations, not simulated or measured emergency-network performance. They require no propagation model, random seeds or learned predictor. *Fig. 1* and *Table 4* define all seven directed edges of a six-node graph. The available paths are $P_A = (S, A, G)$, $P_B = (S, B, C, G)$ and $P_D = (S, D, G)$. All observations are assumed fresh, all contacts are initially open, and G is available. The admission deadline is $D = 0.50$ s, the guard is $g = 0.05$ s, $D_0 = 0.50$ s, $E_0 = 0.20$ J and $H = 1.00$ s. We use $(a, b, c, h) = (0.85, 0.075, 0.045, 0.03)$.

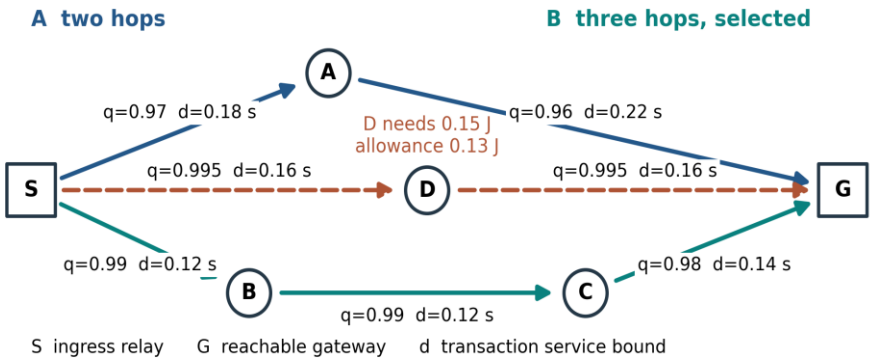


Fig. 1. Route choice in the constructed relay graph.

The three-hop path B is selected at the stated weights. Path D has the lowest unconstrained cost but requires 0.15 J at relay D, exceeding its 0.13 J allowance. Arrows show directed forwarding transactions, not a radio coverage map; labels give transaction success q and service bound d.

Table 4. Complete edge inputs for the deterministic route example.

Edge	q	Service Bounds	TX Energy J	RX Energy J	Contact Horizons
S to A	0.970	0.18	0.080	0.020	0.90
A to G	0.960	0.22	0.100	0.025	0.75
S to B	0.990	0.12	0.100	0.025	2.00
B to C	0.990	0.12	0.100	0.025	2.00
C to G	0.980	0.14	0.100	0.025	2.00
S to D	0.995	0.16	0.120	0.030	1.20
D to G	0.995	0.16	0.120	0.030	1.20

Node allowances in joules are $U_S = U_G = 1.00$, $U_A = 0.30$, $U_B = U_C = 0.25$ and $U_D = 0.13$. The gateway allowance covers reception even though it does not relay onward. *Table 5* reports quantities computed directly from *Eqs. (2)–(6)*. The product column uses independent synthetic transaction outcomes only as an interpretation of the supplied probabilities; no packet outcomes were generated.

Table 5. Computed route quantities under the declared synthetic inputs.

Path	Success Product	Service Bound s	Radio Energy J	Cost	Admission
A	0.931200	0.40	0.225	0.181714	Admitted
B	0.960498	0.38	0.375	0.175633	Admitted and selected
D	0.990025	0.32	0.300	0.124021	Rejected by relay D allowance

For example, the charge at D is $0.030 + 0.120 = 0.150$ J, while that at A is $0.020 + 0.100 = 0.120$ J. Path B spends more network-wide radio energy than A but has a smaller supplied service bound and a larger supplied success product. At the chosen reliability weight those differences yield a lower composite cost. This calculation explains a preference; it does not show that three-hop routing generally reduces delay. Hop service times were specified independently for this example.

The union lower bounds from *Eq. (9)*, under its separately stated probability interpretation, are 0.930, 0.960 and 0.990 for A, B and D. Their closeness to the products reflects these particular high input probabilities. It is not a justification for using a product under unknown dependence. Both reliability expressions also remain subordinate to admission: a high success value does not release unavailable energy.

6.2 | Weight and Reserve Sensitivity

Fig. 2 varies only the weight parameter in *Proposition 3* and the allowance at D. For A, $R_A = -\log(0.9312)$ and $V_A = 0.8075$; for B, $R_B = -\log(0.960498)$ and $V_B = 0.9425$. Solving their equality gives

$$\alpha^* = \frac{V_B - V_A}{R_A - R_B + V_B - V_A} = 0.8133616743. \quad (10)$$

When D remains inadmissible, A is preferred below this value and B above it; at equality the declared lexicographic rule selects A. When $U_D \geq 0.150$ J, D becomes affordable and has the smallest cost throughout the plotted weight interval. No weight change alone produces that admission. The resulting decision map exposes the different roles of a compensatory preference and a reserve threshold.

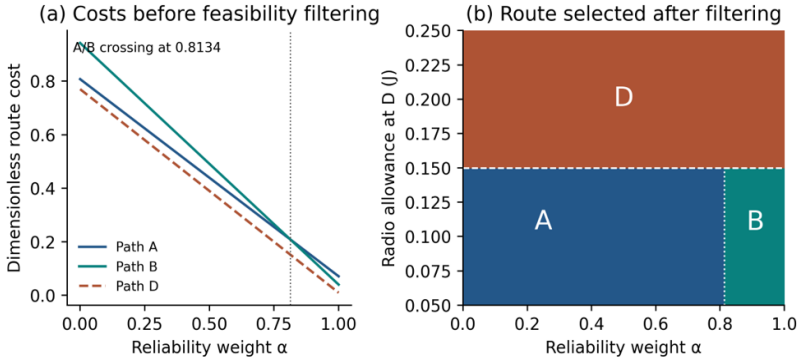


Fig. 2. Exact sensitivity of the constructed example.

Panel a shows path-cost functions before admission, with the A/B crossing at 0.8134. Panel b applies all admission conditions; below the 0.15 J allowance required at D, only the A/B trade-off remains. At or above 0.15 J, D is admitted and dominates this candidate family. The colored regions are computed decisions, not estimated performance surfaces.

6.3 | A Contact-Order Counterexample

Keep path A's service bounds but reduce only the A-to-G contact horizon to 0.35 s. Checking each edge in isolation gives $0.18 + 0.05 \leq 0.90$ and $0.22 + 0.05 \leq 0.35$, so both local tests pass. Sequential forwarding requires the second transaction and its guard to finish by $0.18 + 0.22 + 0.05 = 0.45$ s. Because $0.45 > 0.35$, the path fails *Eq. (4)*. *Fig. 3* displays the temporal mismatch. It is a minimal counterexample to local-only contact screening.

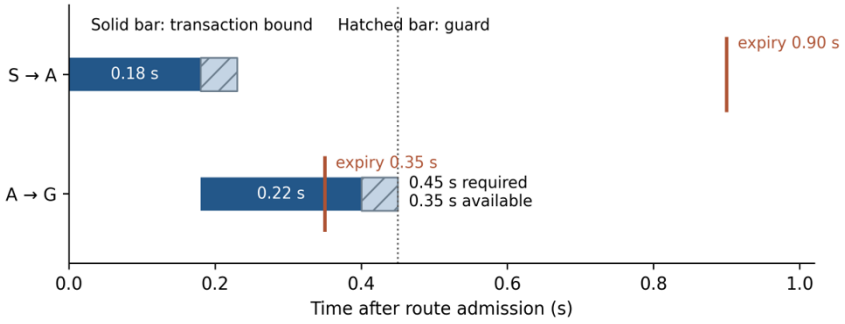


Fig. 3. Local contact duration is insufficient for path admission.

The second transaction begins after 0.18 s and needs completion plus guard by 0.45 s, but its contact expires at 0.35 s. The guard at the first hop is a safety margin against that contact's expiry; it is not an added forwarding delay before the next hop.

Two further perturbations distinguish failure modes. If the application deadline is reduced to 0.35 s while allowances remain at their base values, A and B miss the deadline and D remains unaffordable; no path is admitted. If G is withdrawn, the feasible family is empty regardless of weights. These are observable outcomes of admission. An evaluation should retain such intervals and messages rather than dropping them from a favorable delivery denominator.

6.4 | Hysteresis on a Declared Cost Sequence

Fig. 4 uses two synthetic route-cost sequences at integer epochs 0–19 s. During epochs 0–7, the pair alternates between (1.00,1.01) and (1.01,1.00). During epochs 8–13 the pairs are (1.08,1.00), (1.09,0.99), (1.10,0.98), (1.07,0.99), (1.06,1.00) and (1.05,1.00). B is withdrawn at epoch 14; subsequent cost labels do not restore its availability. Both policies begin on A.

Greedy switching changes routes at epochs 1–7 and repairs at 14, giving eight switches. With $\eta = 0.04$ and a 3 s dwell, the policy retains A during the small alternations, switches to B at 8 and repairs to A at 14, giving two switches. This verifies the action of the declared switching rule on these inputs. It establishes neither a delivery improvement nor a reduction in total energy, because no packet process, control-plane implementation or energy trace is present. A larger threshold could equally miss a useful improvement smaller than 0.04.

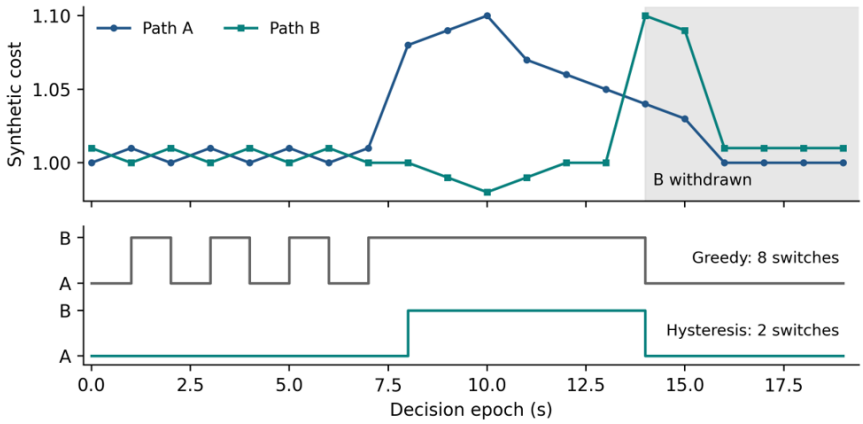


Fig. 4. Route switching on a fully declared synthetic sequence.

The upper panel supplies costs; the shaded interval marks B's withdrawal. The lower panel separates greedy and hysteretic route states vertically for readability. Eight switches become two under the stated rule, including the same forced repair; this is a policy illustration without a packet-performance claim.

7 | Reproducible Evaluation Protocol

7.1 | Experimental Unit and Scenario Design

The primary experiment should estimate the change in on-time delivery caused by the routing policy under a fixed mobility, radio and traffic environment. MANET evaluation guidance [25] provides the general basis for reporting conditions and criteria. Mobility-model choice is itself consequential [30], and credibility studies caution against drawing conclusions from undocumented configurations or inadequately analyzed simulation output [26], [27]. The experimental unit is therefore an independent scenario-seed run, not an individual packet from a long correlated trajectory.

An initial planned scenario family is shown in *Table 6*. Its values are proposed study settings, not measured environmental constants. Before final comparison, a calibration stage must justify the area, relay speeds, contact predictor, radio model and traffic mix against the intended emergency mission. A mission-constrained search-and-rescue trajectory should complement a generic random-mobility model, so that an apparent benefit does not depend entirely on one movement generator. Air-to-air and air-to-ground propagation must be distinguished, including obstacles and antenna orientation when relevant.

Table 6. Proposed evaluation factors to be frozen after independent calibration.

Factor	Planned Levels or Procedure	Purpose
Relay density	6, 12 and 24 relays over the same declared area	Sparse connectivity, redundancy and control scaling
Mobility	Hover/slow survey; mission patrol; coordinated turns	Different contact persistence and prediction error
Offered load	Low, moderate and near-saturation relative to calibrated channel capacity	Separate routing value from congestion effects
Obstruction	Open reference map; fixed obstacle map; changing blockage events	Test expiry and reliability estimates under channel changes
Gateway availability	Continuous; one scheduled loss; repeated outage schedule	Measure failover and disconnected intervals
State quality	Fresh state; added age; position bias; missing observations	Test dependence on information quality
Energy allowance	Ample; heterogeneous; near-reserve relays	Quantify admission, load concentration and unmet demand

Every policy receives the same exogenous trajectories, offered messages and gateway schedule within a paired run. Shared random channel inputs should be indexed by physical link, time and event identity where possible. Merely using the same random-number seed is insufficient if policies consume draws in different orders. Interference and queueing remain endogenous to each policy: they are simulated again because routing and control traffic change them. Forcing identical realized packet successes despite different transmissions would remove a mechanism that the experiment is intended to measure.

Freeze the simulator version, commit, build options, operating environment, radio and propagation model, channel bandwidth, rate adaptation, transmit power, receiver sensitivity, antenna assumptions, MAC retry policy, queue discipline, packet size, generation process and run duration. Store the topology and estimator timestamps that each policy actually saw. If external trajectory data are used, retain their version, license, units, schema and conversion script. A position-and-battery dataset may define scenario covariates but cannot identify the delivery outcome of an unchosen routing policy without an additional generative network model.

7.2|Comparators and Ablations

The minimum comparator set is a frozen shortest-hop route with explicit break behavior, a documented AODV implementation [6] and a documented OLSR implementation [7]. State whether OLSR means RFC 3626 or OLSRv2; they are not interchangeable labels. Add ETX-style routing [16] to isolate the contribution beyond link-quality ranking. P-OLSR [12] is a particularly relevant predictive comparator if its implementation and motion inputs can be reproduced. PMHRP [13] or the adaptive hybrid method [19] should be included only after reproducing their control architecture in the same simulator. A learned baseline [18] needs separate training and test scenarios and must count training or exploration assumptions appropriately.

Table 8 defines ablations that alter one substantive design choice at a time. All receive an equal tuning budget. Hard-constraint ablations belong in simulation because they deliberately remove protections; their violation rates must be reported beside delivery outcomes. A baseline's own standard behavior should be preserved

rather than degrading its timers or withholding observations that its specification legitimately requires. If one policy receives additional position or battery information, report an information-matched comparison as well as the operational comparison.

7.3 | Outcomes and Denominators

The primary outcome is the proportion of distinct generated application messages received by an eligible gateway before their own deadline. Let N be all generated messages in the measurement cohort and I_m indicate on-time receipt. Then

$$Y_{on} = \frac{1}{N}\sum_{m=1}^{N} I_m. \tag{11}$$

Sources must stop generation early enough to follow every message through its deadline, or append an equivalent follow-up interval. Otherwise end-of-run messages are administratively censored and cannot be called failed or successful without a declared rule. The cohort includes messages generated during disconnection, rejected by admission or lost at a queue. A run with $N = 0$ is undefined for this ratio and is reported separately. Gateway application deduplication prevents retries and multiple-gateway copies from enlarging the numerator.

Table 7 defines secondary outcomes. One-way delay follows a declared creation-to-application-receipt clock convention, consistent with the measurement concerns of RFC 7679 [31]. Loss uses an explicit waiting threshold or deadline, as required to make a finite observation operationally meaningful [32]. These definitions apply to application messages; the RFCs do not by themselves supply TAARA's application-level denominator.

Table 7. Failure-inclusive measures and minimum event records.

Measure	Definition or Reporting Rule	Required Records
On-time delivery	Unique messages received by their deadline divided by all generated messages	Generation, deadline, application receipt and unique ID
Delivery by declared horizon	Unique messages received by a separately declared follow-up horizon divided by all generated messages	Cohort membership and complete follow-up
Delay	Distribution for received messages, alongside the full cohort failure fraction	Synchronized creation and application receipt timestamps
Goodput	Unique delivered payload bits divided by measurement time	Payload length, duplicate status and observation interval
Radio energy	TX, RX, idle, listening, setup, discovery and repair joules, with failed attempts included	Per-node power-state trace or calibrated event model
Vehicle energy	Radio plus propulsion, avionics and other mission consumption, reported separately	Flight trajectory, vehicle model and energy integration
Control and stability	Control bytes, channel airtime, route changes, forced repairs and recovery delay	Message type, actual transmission, route state and failure events
Feasibility and computation	Rejection reasons, reserve violations, envelope exceedances and decision-time distribution	Admission inputs, reservations, measured consumption and runtime

Conditional delay alone is susceptible to selection: a policy can appear faster by failing on difficult messages. Report its distribution together with on-time delivery and the undelivered fraction. A restricted time-to-delivery summary may assign the application deadline to undelivered messages, but it must be named and interpreted as a composite outcome, not as their observed delay. Similarly, energy per delivered message can

rise when delivery falls, and is undefined when no message is delivered. Report total joules and delivered counts alongside every such ratio.

Radio savings must not be renamed endurance gains. For an unchanged flight plan, quantify what fraction of mission consumption the radio represents; if flight behavior changes, rerun the vehicle-energy calculation. Control bytes are counted at actual transmissions, so forwarding one advertisement across several relays is not counted as one byte contribution. Report startup and repair overhead as well as steady-state averages. A protocol that succeeds only after an expensive convergence period may be unsuitable for a short emergency mission even if its later delivery ratio is high.

7.4 | Statistical Analysis and Tuning Separation

For scenario c, seed r and comparator b, compute the paired absolute difference $\Delta_{cr} = Y_{cr}^{TAARA} - Y_{cr}^{b}$. Summarize the mean and distribution of these differences within each scenario. A confidence interval may use paired resampling of independent runs; bootstrap resampling follows the general principle of Efron [33], but the resampling unit must reflect this experiment's dependence. If several seeds share a sampled trajectory or map, resample at that higher cluster level or use a justified hierarchical model. Treating thousands of packets from one flight as independent replications would understate uncertainty.

Use an independent pilot to choose replication counts based on the desired precision for the primary outcome. As an initial planning value, 30 paired independent runs per scenario can support a pilot, but it is not a universal adequacy threshold. Freeze the final run count or a prespecified precision-stopping rule before inspecting test differences. Report the completed count, failed runs and reasons for exclusion. A simulator crash is a missing run requiring diagnosis; it is not a delivered-message failure that can be inserted into the denominator arbitrarily.

Report on-time delivery differences in percentage points. A relative change, if useful, is $100(Y^{TAARA} - Y^{b})/Y^{b}$ and must name comparator b; it is undefined when $Y^{b} = 0$. Delay and energy ratios require the same explicit denominators. For several scenarios, present scenario-specific effects and a prespecified aggregation with stated weights. A pooled packet ratio and an equally weighted mean of run ratios answer different questions. Neither should silently replace the other.

Tune estimator windows, a, b, c, h, H, guard, freshness threshold, hysteresis, dwell and candidate limits on training or calibration scenarios only. Hold out entire mobility or obstacle conditions when transfer is the question. A large sensitivity sweep can explain mechanism, but selecting its best cell after seeing the test set is additional tuning. If formal hypothesis tests are used across many outcomes, define the primary comparison in advance and state how multiplicity is handled. Practical effect size and uncertainty remain necessary even when a significance test is reported.

Table 8. Mechanism ablations and the outcomes needed to interpret them.

Variant	Single Change From the Full Design	Mechanistic Outcome
Local contact screen	Replace cumulative timing with per-edge service checks	Late downstream failures and false admissions
No reserve admission	Remove the node allowance test while retaining the score	Reserve violations and apparent delivery gained by spending protected energy
No contact preference	Set h to zero and renormalize the remaining weights	Route churn among otherwise admissible paths
No hysteresis	Set eta and dwell to zero	Voluntary switches, control airtime and delivered traffic
No freshness exclusion	Retain observations beyond the normal age limit	Decisions based on stale state and subsequent failures
Limited candidate search	Cap the candidate pool at a declared size	Decision time and missed feasible paths versus the exact small-graph reference

7.5 | Reproducibility and Progression to Flight Tests

The experimental archive must contain configuration files, immutable code identifiers, raw application and radio events, seed mappings, trajectories, gateway schedules, parameter-selection records and event-to-figure scripts. Every plotted point should be traceable to its scenario, seed, policy, cohort size and aggregation. Preserve failed transmissions and admission rejections. A result table is reproducible only if the event transformations and denominator rules are available, not merely because its final values can be copied.

Verification should precede comparative claims. First, reproduce the deterministic examples in Section 6 and compare optimized search with exhaustive enumeration on small graphs. Then test delayed advertisements, lost reservation releases, duplicate acknowledgements, packet-ID reuse after reboot, gateway withdrawal during setup and reserve contention between sources. Next, calibrate radio and service envelopes in emulation or a controlled bench network. Finally, conduct flight trials with fixed mission trajectories and synchronized energy and application logs. Each stage answers a different question: implementation correctness, model adequacy and operational effect should not be collapsed into one performance curve.

8 | Discussion and Limitations

TAARA's principal value is that its decisions can be inspected at the point where preference meets feasibility. A route rejection identifies a failed deadline, contact or allowance condition; an accepted route exposes its cost components. This makes disagreements interpretable. A mission operator can change a declared deadline or reserve policy, while a protocol engineer can investigate stale estimates or expensive dissemination. The distinction also prevents a numerical weight from silently overriding a physical restriction.

The specification has several limits. Its exact search can be expensive, and its replicated topology view can become stale during fast motion. Reservation setup consumes time precisely when contacts are short. A very conservative service or contact envelope may reject traffic that would often succeed, while an optimistic one may admit paths that fail. Hysteresis can suppress oscillation but can also retain an inferior route. These are expected trade-offs to measure, not defects removed by calling the method adaptive.

The model also separates per-message admission from network-wide congestion control. Several sources can select the same attractive relay even when each has sufficient battery allowance. Atomic energy reservations prevent one form of oversubscription; they do not reserve airtime or prove queue stability. If shared-channel contention increases service times, envelopes must be recalibrated or paired with an explicit scheduling and rate-control mechanism. The additive objective likewise does not fully represent interference coupling across paths. ETT-style considerations [17] are relevant, but simply adding another metric term would not establish a correct shared-channel model.

Control security remains external to the construction. False position, energy or gateway claims can invalidate its inputs; compromised relays can ignore a pinned route. Authentication, replay protection, clock synchronization and secure boot require an explicit system design and their own resource accounting. They should be present in a deployed emergency network, but no security guarantee is derived here. The stated loop argument assumes honest processing and well-formed packets.

Finally, the numerical illustrations demonstrate logical consequences of specified values. They do not establish calibrated success probabilities, radio ranges, network lifetime or superiority over AODV, OLSR, predictive routing or learning. The next scientific milestone is a complete implementation with a matched experiment and raw records. A positive outcome would support a conditional claim for its tested operating range. A negative result could be equally informative if it identifies the density, load or information-age regime in which admission and control overhead outweigh adaptive selection.

9 | Conclusion

This paper specifies TAARA as a constrained routing policy for a moving emergency relay network. Its defining separation is between admitting a route that satisfies declared timing and energy conditions and ranking admitted routes by a normalized preference. Cumulative contact checks, per-node reception and transmission charges, expiring reservations and fixed-route forwarding make the operational meaning of that separation explicit.

The analysis establishes conditional feasibility, distinguishes independent reliability products from justified conditional bounds, characterizes weight-induced route changes and bounds the effect of cost-estimation error on voluntary switching. Fully declared calculations show an unaffordable lowest-cost path, a contact-order failure and the behavior of hysteresis. These results support the internal logic of the construction. Establishing an emergency-network performance advantage requires the implemented, failure-inclusive and uncertainty-aware study specified here.

Data Availability

All numerical inputs used for the route calculations are given in *Table 4* and Section 6; the route set and decision rule are explicit. Appendix A provides a standalone Python calculation for *Table 5* and the sensitivity threshold. The switching sequence and policy parameters are fully declared in Section 6.4. No external dataset, field measurements or packet-level simulation results are used in the illustrations. A network implementation and its experimental archive are future work; no external repository or dataset availability is implied.

Acknowledgments

The authors acknowledge the researchers whose previous contributions to Flying Ad Hoc Networks, UAV-assisted communication, emergency networking, and adaptive routing provided the theoretical foundation for this study.

Author Contributions

S. A. E.: Conceptualization, methodology, formal analysis, algorithm development, and writing—original draft.

M. H.: Technical analysis, validation, interpretation of results, and writing—review and editing.

Both authors reviewed and approved the final manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of Interest

The authors declare that they have no conflicts of interest related to the publication of this paper.

References

[1] Matracia, M., Saeed, N., Kishk, M. A., & Alouini, M. S. (2022). Post-disaster communications: Enabling technologies, architectures, and open challenges. *IEEE Open Journal of the Communications Society*, *3*, 1177–1205. https://doi.org/10.1109/OJCOMS.2022.3192040

[2] Bekmezci, I., Sahingoz, O. K., & Temel, cSamil. (2013). Flying ad-hoc networks (FANETs): A survey. *Ad Hoc Networks*, *11*(3), 1254–1270. https://doi.org/10.1016/j.adhoc.2012.12.004

[3] Gupta, L., Jain, R., & Vaszkun, G. (2015). Survey of important issues in UAV communication networks. *IEEE Communications Surveys & Tutorials*, *18*(2), 1123–1152. https://doi.org/10.1109/COMST.2015.2495297

[4] Zeng, Y., Zhang, R., & Lim, T. J. (2016). Wireless communications with unmanned aerial vehicles: Opportunities and challenges. *IEEE Communications Magazine*, *54*(5), 36–42. https://doi.org/10.1109/MCOM.2016.7470933

[5] Mozaffari, M., Saad, W., Bennis, M., Nam, Y. H., & Debbah, M. (2019). A tutorial on UAVs for wireless networks: Applications, challenges, and open problems. *IEEE Communications Surveys & Tutorials*, *21*(3), 2334–2360. https://doi.org/10.1109/COMST.2019.2902862

[6] Perkins, C., Belding Royer, E., & Das, S. (2003). *Ad hoc on-demand distance vector (AODV) routing*. https://doi.org/10.17487/RFC3561

[7] Clausen, T., & Jacquet, P. (2003). *Optimized link state routing protocol (OLSR)*. https://doi.org/10.17487/RFC7181

[8] Clausen, T., Dearlove, C., Jacquet, P., & Herberg, U. (2014). *The optimized link state routing protocol version 2*. https://doi.org/10.17487/RFC7181

[9] Clausen, T., Dearlove, C., & Dean, J. (2011). *Mobile ad hoc network (MANET) neighborhood discovery protocol (NHDP)*. https://doi.org/10.17487/RFC6130.

[10] Chroboczek, J., & Schinazi, D. (2021). *The Babel routing protocol (RFC 8966)*. https://doi.org/10.17487/RFC8966

[11] Wheeb, A. H., Nordin, R., Samah, A., Alsharif, M. H., & Khan, M. A. (2021). Topology-based routing protocols and mobility models for flying ad hoc networks: A contemporary review and future research directions. *Drones*, *6*(1), 9. https://doi.org/10.3390/drones6010009

[12] Rosati, S., Krużelecki, K., Heitz, G., Floreano, D., & Rimoldi, B. (2015). Dynamic routing for flying ad hoc networks. *IEEE Transactions on Vehicular Technology*, *65*(3), 1690–1700. https://doi.org/10.1109/TVT.2015.2414819

[13] Ullah, S., Mohammadani, K. H., Khan, M. A., Ren, Z., Alkanhel, R., Muthanna, A., & Tariq, U. (2022). Position-monitoring-based hybrid routing protocol for 3D UAV-based networks. *Drones*, *6*(11), 327. https://doi.org/10.3390/drones6110327

[14] Casteigts, A., Flocchini, P., Quattrociocchi, W., & Santoro, N. (2012). Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, *27*(5), 387–408. https://doi.org/10.1080/17445760.2012.668546

[15] Burleigh, S., Fall, K., & Birrane, E. J. (2022). *Bundle protocol version 7*. RFC 9171. https://doi.org/10.17487/RFC9171.

[16] De Couto, D. S. J., Aguayo, D., Bicket, J., & Morris, R. (2003). A high-throughput path metric for multi-hop wireless routing. *Proceedings of the 9th annual international conference on mobile computing and networking* (pp. 134–146). ACM. https://doi.org/10.1145/938985.939000

[17] Draves, R., Padhye, J., & Zill, B. (2004). Routing in multi-radio, multi-hop wireless mesh networks. *Proceedings of the 10th annual international conference on mobile computing and networking* (pp. 114–128). ACM. https://doi.org/10.1145/1023720.1023732

[18] Minhas, H. I., Ahmad, R., Ahmed, W., Waheed, M., Alam, M. M., & Gul, S. T. (2021). A reinforcement learning routing protocol for UAV aided public safety networks. *Sensors*, *21*(12), 4121. https://doi.org/10.3390/s21124121

[19] Sun, K., Liu, M., Yin, C., & Wang, Q. (2025). Adaptive extended Kalman prediction-based SDN-FANET segmented hybrid routing scheme. *Sensors*, *25*(5), 1417. https://doi.org/10.3390/s25051417

[20] Almansor, M. J., Din, N. M., Baharuddin, M. Z., Ma, M., Alsayednoor, H. M., Al Shareeda, M. A., & Al Asadi, A. J. (2024). Routing protocols strategies for flying Ad-Hoc network (FANET): Review, taxonomy, and open research issues. *Alexandria Engineering Journal*, *109*, 553–577. https://doi.org/10.1016/j.aej.2024.09.032

[21] Ganesh, S., Gopalasamy, V., & Shibu, N. B. S. (2021). Architecture for drone assisted emergency ad-hoc network for disaster rescue operations. *2021 International conference on communication systems & networks (COMSNETS)* (pp. 44–49). IEEE. https://doi.org/10.1109/COMSNETS51098.2021.9352814

[22] Kumar, P., Darshi, S., & Shailendra, S. (2021). Drone assisted device to device cooperative communication for critical environments. *IETCommunications*, *15*(7), 957–972. https://doi.org/10.1049/cmu2.12134

[23] Esubonteng, P. K., & Rojas Cessa, R. (2022). RESTORE: Low-energy drone-assisted NLoS-FSO emergency communications. *IEEE Access*, *10*, 115282–115294. https://doi.org/10.1109/ACCESS.2022.3218014

[24] Zeng, Y., Xu, J., & Zhang, R. (2019). Energy minimization for wireless communication with rotary-wing UAV. *IEEE Transactions on Wireless Communications*, *18*(4), 2329–2345. https://doi.org/10.1109/TWC.2019.2902559

[25] Corson, S., & Macker, J. (1999). *Mobile ad hoc networking (MANET): Routing protocol performance issues and evaluation considerations*. https://doi.org/10.17487/RFC2501

[26] Kurkowski, S., Camp, T., & Colagrosso, M. (2005). MANET simulation studies: the incredibles. *ACM Sigmobile Mobile Computing and Communications Review*, *9*(4), 50–61. https://doi.org/10.1145/1096166.1096174

[27] Pawlikowski, K., Jeong, H. D., & Lee, J. S. (2002). On credibility of simulation studies of telecommunication networks. *IEEE Communications Magazine*, *40*(1), 132–139. https://doi.org/10.1109/35.978060

[28] Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, *22*(158), 209–212. https://doi.org/10.1080/01621459.1927.10502953

[29] Dijkstra, E. W. (2022). A note on two problems in connexion with graphs. *Edsger wybe dijkstra: His life, work, and legacy* (pp. 287–290). ACM. https://doi.org/10.1007/BF01386390

[30] Camp, T., Boleng, J., & Davies, V. (2002). A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing*, *2*(5), 483–502. https://doi.org/10.1002/wcm.72

[31] Almes, G., Kalidindi, S., Zekauskas, M., & Morton, A. (2016). *A one-way delay metric for IP performance metrics (IPPM)*. https://doi.org/10.17487/RFC7679

[32] Almes, G., Kalidindi, S., Zekauskas, M., & Morton, A. (2016). *A one-way loss metric for ip performance metrics (IPPM)*. https://doi.org/10.17487/RFC7680

[33] Efron, B. (1992). Bootstrap methods: another look at the jackknife. *Breakthroughs in statistics: Methodology and distribution* (pp. 569–593). Springer. https://doi.org/10.1214/aos/1176344552

Appendix

Reproducing the route calculations

The following script uses only the Python standard library. Its tuple fields are transaction success, service seconds, transmitter joules, receiver joules and residual contact seconds. It prints the three path costs, success products, service bounds and energy sums, verifies admission outcomes and computes the A/B crossing. Its scope is deterministic arithmetic, not radio simulation.

```
from math import log, prod

edges = {
    ('S','A'): (.97, .18, .08, .02, .90),
    ('A','G'): (.96, .22, .10, .025, .75),
    ('S','B'): (.99, .12, .10, .025, 2.),
    ('B','C'): (.99, .12, .10, .025, 2.),
    ('C','G'): (.98, .14, .10, .025, 2.),
    ('S','D'): (.995, .16, .12, .03, 1.2),
    ('D','G'): (.995, .16, .12, .03, 1.2),
}
paths = {'A': 'SAG', 'B': 'SBCG', 'D': 'SDG'}
allow = dict(S=1., A=.30, B=.25, C=.25, D=.13, G=1.)

def calc(path, alpha=.85):
    elapsed = energy = risk = contact = 0.
    charges, probabilities = {}, []
    temporal = True
    for i, j in zip(path, path[1:]):
        q, d, tx, rx, tau = edges[i, j]
        elapsed += d
        temporal = temporal and elapsed + .05 <= tau
        charges[i] = charges.get(i, 0.) + tx
        charges[j] = charges.get(j, 0.) + rx
        energy += tx + rx
        risk -= log(q)
        contact += max(0., 1. - tau)
        probabilities.append(q)
    other = .5*elapsed/.5 + .3*energy/.2 + .2*contact
    cost = alpha*risk + (1.-alpha)*other
    feasible = (temporal and elapsed <= .5 and
                all(v <= allow[i] + 1e-12
                    for i, v in charges.items()))
    return cost, prod(probabilities), elapsed, energy, feasible

for name, path in paths.items():
    print(name, calc(path))
assert calc(paths['A'])[4] and calc(paths['B'])[4]
assert not calc(paths['D'])[4]
assert calc(paths['B'])[0] < calc(paths['A'])[0]
ra, rb = -log(.9312), -log(.960498)
va, vb = .8075, .9425
print('alpha crossing', (vb-va)/(ra-rb+vb-va))
```