Service Continuity in RAN Blackouts with Landed UAV Under Uncertain Restoration

Aurora Martiny
Politecnico di Torino
Turin, Italy
aurora.martiny@polito.it

Mohamed Amine Mbarek
Politecnico di Torino
Turin, Italy
mohamed.mbarek@polito.it

Michela Meo
Politecnico di Torino
Turin, Italy
michela.meo@polito.it

Daniela Renga
Politecnico di Torino
Turin, Italy
daniela.renga@polito.it

Greta Vallero
Politecnico di Torino
Turin, Italy
greta.vallero@polito.it

***Abstract*—Long-duration power outages can simultaneously disable multiple terrestrial base stations, threatening cellular network resilience until grid restoration. Landed UAV-mounted base stations can restore service from rooftop locations while avoiding the dominant propulsion cost of hovering operation. Yet, their finite battery must be allocated before the repair time is known. This paper introduces a restoration-time-aware stochastic scheduling framework for landed UAV-BS recovery. For each candidate rooftop position, a MILP jointly controls user admission and transmit power over hourly blackout slots. The restoration time is modeled through a discrete lognormal prior, and the objective maximizes expected service continuity by penalizing the weakest pre-restoration service level for each possible outage duration. Results show that proximity to the outage footprint is more important than altitude alone, that the proposed scheduler avoids late service collapse without worst-case over-conservatism, and that accounting for service-floor compliance changes the interpretation of throughput-oriented recovery policies. Optimistic repair assumptions can reduce expected service continuity by up to 24%.**

***Index Terms*—Radio access network resilience; UAV-mounted base station; landed UAV deployment; blackout recovery; restoration-time uncertainty; service continuity; stochastic MILP**

I. Introduction

As everyday life becomes increasingly dependent on cellular networks, resilience has become a central requirement in the design of modern Radio Access Networks (RANs) [1]. Among the various challenges that may affect network operation, long-duration power outages are particularly critical, as they can disable multiple base stations (BSs) simultaneously and persist beyond the autonomy of local backup systems. According to the 2024 ENISA report, system failures accounted for 60% of reported incidents, with power-supply outages being the second most impactful technical cause within this category [2]. The duration of a blackout strongly affects which recovery mechanism is effective. Conventional backup resources are typically dimensioned for limited autonomy, while severe disruptions may last longer because of infrastructure damage, accessibility constraints, or repair logistics [3]. In such cases, recovery mechanisms must account not only for how much service can be restored immediately, but also for how long that service can be sustained.

In this context, unmanned aerial vehicles (UAVs) equipped with communication capabilities offer a flexible means to assist users who are left unserved by unavailable or capacity-limited terrestrial infrastructure [4]. However, their use in long-term blackouts is constrained by finite battery energy [5]. To extend UAV-mounted base station (UAV-BS) operation, landed deployments have recently been considered. In this approach, the UAV-BS operates from a fixed rooftop position instead of continuously hovering, thereby reducing propulsion energy consumption [6], [7]. However, it makes the recovery problem more coupled: the selected rooftop determines the reachable residual demand, propagation conditions, required transmit power, and battery depletion rate. Therefore, landed UAV-BS recovery requires joint reasoning over landing position, user admission, transmit power, and time.

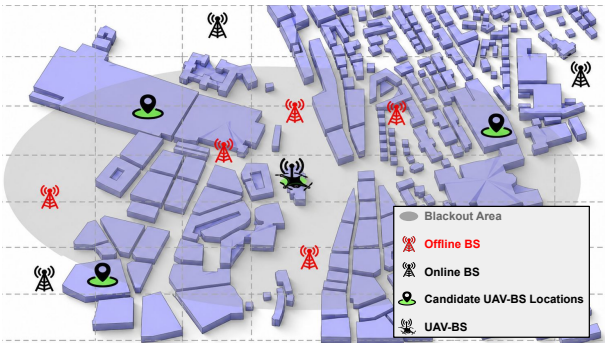


Fig. 1: Visual schema of landed UAV-BS recovery during a RAN blackout.

Although landed UAV-BS deployments can extend service duration, the restoration time remains uncertain and is generally unknown at the time of UAV-BS scheduling. Indeed, restoration time depends on the severity of the failure, infrastructure accessibility, and repair logistics. Maximizing instantaneous coverage or throughput can therefore be misleading in long outages: aggressive early service may deplete the battery before grid restoration, while overly conservative operation may waste recovery opportunities if the outage ends early. This motivates a restoration-aware scheduling model in which the outage duration is treated as a random variable and the schedule is evaluated across possible restoration times. To the best of our knowledge, existing studies do not provide a restoration-aware UAV-BS scheduling framework that explicitly balances immediate traffic recovery and sustained service continuity under uncertain blackout durations. In this paper, we address this problem through the following contributions:

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

- We introduce a restoration-time-aware service-continuity model for finite-energy landed UAV-BS recovery, where the service continuity associated with each possible restoration duration is measured by the minimum residual-demand recovery achieved before restoration.
- We formulate the resulting admission and transmit-power scheduling problem as a stochastic mixed-integer linear program (MILP) that embeds the restoration-duration distribution directly in the objective while enforcing battery, capacity, link-budget, and minimum-service constraints.
- We evaluate the framework using real BS deployments, urban morphology, and data-driven traffic demand, quantifying the impact of landing position, deterministic-horizon baselines, transmit-power allocation, and restoration-prior mismatch on service continuity.

II. Related Work

UAV-assisted cell outage compensation has been studied through UAV placement, user association, and transmit-power optimization for failed terrestrial BSs, while distinguishes between short- and long-term failures [8]. However, the UAV-to-user channel is usually modeled under a simplified line-of-sight (LoS) assumption, which abstracts blockage, non-line-of-sight (NLoS) attenuation, and the dependence of coverage on detailed urban morphology. This abstraction limits the ability to capture the strong position-dependent effects that arise in dense urban deployments. More recent learning-based approaches focus on dynamic settings, where UAV trajectories, transmit power, and user association are controlled online to support users affected by inactive or failed cells [9]. Nevertheless, the restoration time is treated as part of the scenario description rather than as an uncertain quantity that directly enters the service-continuity objective.

Another line of work addresses UAV-BS energy limitations through reinforcement learning based energy-aware control [10] and landed UAV-BS operation [7]. Rooftop-based landed deployment has been investigated specifically for RAN blackout recovery, showing that site selection affects the ability of the UAV-BS to compensate for failed terrestrial infrastructure [6]. These studies motivate landed operation and site selection, but they do not address how the remaining communication energy should be scheduled since the blackout duration is unknown.

Stochastic optimization has been applied to disaster-response UAV networks to account for uncertain demand, scenario severity, and resource requirements [11]. However, UAVs coordinate and execute service requests from ground users, rather than acting as substitutes for failed terrestrial BSs serving residual cellular traffic. Moreover, the modeled uncertainty concerns demand severity rather than restoration duration. Power-system outage data further show that restoration times are highly variable and difficult to predict, with distributions that are heavy-tailed and with a mean that is highly sensitive to the observation period [12]. This variability makes the fixed-horizon assumption difficult to justify in practice and motivates treating the outage duration as a random variable when designing UAV-BS recovery policies.

Taken together, existing works demonstrate the potential of UAV-assisted outage recovery, landed operation, and energy-aware control. However, prior formulations generally optimize deterministic or instantaneous recovery objectives and do not explicitly model restoration-duration uncertainty within the service-continuity objective itself. Consequently, the interaction between finite UAV-BS energy, urban-propagation-dependent landing feasibility, and stochastic restoration time remains insufficiently addressed.

III. System Model

We consider a simulation model in which time is discretized into hourly slots of duration Δt. A blackout starts at time t_0, and the planning horizon is defined as $\mathcal{T} = \{1, \dots, T_{\max}\}$, where slot $t = 1$ denotes the first hour after blackout onset. The urban cellular network is composed of a set of terrestrial BSs $\mathcal{B}$ and a set of spatial traffic tiles $\mathcal{U}$. Each tile $u \in \mathcal{U}$ represents an aggregate demand point, located at its centroid $q_u = (x_u, y_u)$, and generates hourly traffic demand $d_u(t)$ at slot t. The model takes as input an urban morphology map, used to identify feasible rooftop UAV-BS locations and to compute propagation conditions between traffic tiles, terrestrial BSs, and UAV-BS candidates.

The blackout duration τ is uncertain and modeled as a discrete random variable supported on the same finite horizon $\mathcal{T}$. A realization $\tau = \ell$ means that the terrestrial infrastructure is restored after ℓ blackout slots. Power-grid restoration times are well characterized by lognormal distributions in the reliability literature [12], so the probability that the grid recovers exactly at slot ℓ is

$$\pi_\ell = \Pr(\tau = \ell), \qquad \ell \in \mathcal{T}, \tag{1}$$

with $\sum_{\ell=1}^{T_{\max}} \pi_\ell = 1$. Letting $F_\tau(\cdot)$ denote the lognormal CDF, each π_ℓ captures the probability mass in the interval $(\ell-1, \ell]$, renormalized over the truncated support:

$$\pi_\ell = \frac{F_\tau(\ell) - F_\tau(\ell-1)}{\sum_{m=1}^{T_{\max}} [F_\tau(m) - F_\tau(m-1)]}, \qquad \ell \in \mathcal{T}. \tag{2}$$

During the blackout, a subset $\mathcal{B}^{\text{off}} \subseteq \mathcal{B}$ of terrestrial BSs becomes unavailable and is assigned a zero capacity. The set of still available BSs is $\mathcal{B}' = \mathcal{B} \setminus \mathcal{B}^{\text{off}}$. The failed BSs are selected through a spatially correlated outage-generation process, in which neighboring BSs are progressively deactivated. This captures the fact that failures are often geographically correlated rather than independent across BS sites. To mitigate the loss of connectivity, the deployment of a single landing UAV-BS denoted by d is considered in addition to the set $\mathcal{B}'$. The service provisioning follows a two-stage association procedure:

- **Baseline terrestrial BS association phase:** users are first associated with the available ground BSs;
- **UAV-assisted association phase:** unserved users after the first phase are considered for association with the UAV-BS, subject to its resource and energy constraints.

A. Terrestrial Association Policy

At each time $t \in \mathcal{T}$, users are sequentially associated with the terrestrial BSs in $\mathcal{B}'$ according to a minimum-path-loss criterion under capacity constraints. Let $i \in \mathcal{B}'$ denote an available terrestrial BS, and let $PL_{u,i}(t)$ be the path loss between user u and BS i at slot t. At the beginning of each slot, the residual capacity of each available BS is initialized as $C_i^R(t)$. For a user u, the set of feasible serving BSs is

$$\mathcal{B}_u(t) = \left\{ i \in \mathcal{B}' : PL_{u,i}(t) \le PL_{\max},\ C_i^R(t) \ge d_u(t) \right\}, \tag{3}$$

where $PL_{\max}$ is the maximum allowable path loss required by the link budget, and $d_u(t)$ is the traffic demand of user u at slot t. To account for competition among users for the finite BS capacities, association is performed sequentially. Users are first ordered according to their best available path loss,

$$\widetilde{PL}_u(t) = \min_{i \in \mathcal{B}'} PL_{u,i}(t), \tag{4}$$

and processed in ascending order of $\widetilde{PL}_u(t)$. For each processed user u, if $\mathcal{B}_u(t) \neq \emptyset$, the serving BS is selected as

$$i_u^{\mathrm{BS}}(t) = \arg\min_{i \in \mathcal{B}_u(t)} PL_{u,i}(t), \tag{5}$$

and the residual capacity of the selected BS is updated as

$$C^R_{i_u^{\mathrm{BS}}(t)}(t) \leftarrow C^R_{i_u^{\mathrm{BS}}(t)}(t) - d_u(t). \tag{6}$$

If $\mathcal{B}_u(t) = \emptyset$, user u is marked as unserved by the terrestrial network and is included in the UAV-assisted association phase. This sequential policy prioritizes users with favorable terrestrial channel conditions while accounting for the residual capacity available at each BS.

B. UAV-BS Candidate Locations

Candidate UAV-BS locations are restricted to feasible rooftop positions within the considered area. To characterize the spatial center of the outage, we define the blackout centroid as the geometric centroid of the unavailable terrestrial BSs:

$$q_{\mathrm{bc}} = \frac{1}{|\mathcal{B}^{\mathrm{off}}|} \sum_{i \in \mathcal{B}^{\mathrm{off}}} q_i, \tag{7}$$

where $q_i = (x_i, y_i)$ denotes the planar position of BS i. The blackout centroid provides a geometric reference for the spatial concentration of the failed infrastructure and is used to construct reference UAV-BS locations. We distinguish between two classes of candidate locations:

- *Cluster-based candidates* $\mathcal{D}^{\mathrm{clust}}$: terrestrial BSs are grouped according to geographical proximity, enforcing a minimum inter-cluster distance of 200 m. For each cluster, one candidate location $q_d = (x_d, y_d)$ is selected as the rooftop of the tallest building within the tile containing the cluster centroid. This construction provides spatially distributed candidate locations while favoring elevated rooftops with potentially larger coverage footprints.
- *Geometry-driven reference candidates* $\mathcal{D}^{\mathrm{ref}}$: additional locations are included to isolate the effect of spatial placement on service continuity. These include the rooftop closest to the blackout centroid, the rooftop with the maximum absolute altitude in the area, and a hovering position directly above q_{bc} at a fixed altitude.

The complete set of candidate UAV-BS locations is therefore defined as $\mathcal{D} = \mathcal{D}^{\mathrm{clust}} \cup \mathcal{D}^{\mathrm{ref}}$. The reference candidates are used only as benchmarks and are not intended as proposed deployment rules. For each candidate location $d \in \mathcal{D}$, the wireless channel between user u and serving node $i \in \mathcal{B} \cup \{d\}$ is modeled through LoS/NLoS path loss according to the 3GPP Urban Macro (UMa) specification [13].

C. Residual Demand and UAV-BS Power Model

Users that remain unserved after the terrestrial association phase form the set $\mathcal{U}_0(t) \subseteq \mathcal{U}$, and the corresponding residual demand is $D_{\mathrm{left}}(t) = \sum_{u \in \mathcal{U}_0(t)} d_u(t)$. The UAV-BS is activated only if $D_{\mathrm{left}}(t) > 0$ and serves the demand, either fully or partially, subject to capacity and energy constraints. Modeling its power consumption is necessary, since the amount of traffic served and the transmission effort directly determine the energy depletion over the blackout horizon. In the literature, BS power is expressed as the sum of a fixed component and variable terms that depend on load and transmit power, as introduced in the EARTH framework and subsequent analytical models [14]. Measurement-based studies on LTE and NR BSs show that linear and interpretable models based on load and transmit power accurately capture power consumption and outperform more complex nonlinear predictors [15]. Although neural-network-based models can improve prediction accuracy, their black-box structure is less suited to the MILP formulation considered here. Following this established line, we adopt the linear decomposition as established in [15]

$$P(t) = b_{\mathrm{DL}}^{\mathrm{eff}}\, \rho(t) + b_{P_{\max}}\, p(t) + b_0^{\mathrm{eff}}, \tag{8}$$

where $P(t)$ denotes the electrical power consumption of the UAV-BS communication system, $p(t)$ is the radiated transmit power controlling coverage feasibility, $b_{\mathrm{DL}}^{\mathrm{eff}}$ captures the load-dependent term, $b_{P_{\max}}$ scales the transmit-power contribution, and b_0^{eff} represents baseline consumption. The load is defined as $\rho(t) = R(t)/C_d$, where $R(t)$ denotes the aggregate admitted traffic (see Section IV) and C_d is the licensed maximum traffic volume at rooftop d.

IV. Problem Formulation

Building on the stochastic restoration model introduced in Section III, we formulate the UAV-BS scheduling problem for a fixed candidate position $d \in \mathcal{D}$. The objective is to sustain service during the blackout by maximizing the worst recovered-demand fraction experienced before restoration, in expectation over the uncertain restoration time. The UAV-BS jointly decides: i) the subset of unserved users to admit at each time slot $y_{u,d}(t)$, and ii) the transmit power profile $p(t)$ over the blackout horizon, under received-power feasibility, capacity constraints, and finite battery energy. The position d is treated as a parameter; candidate positions are compared after solving the problem independently for each d. Let

$y_{u,d}(t) \in \{0,1\}$ indicate whether user $u \in \mathcal{U}_0(t)$ is admitted by the UAV-BS located at d during slot $t \in \mathcal{T}$. The recovered traffic is

$$R(t) = \sum_{u \in \mathcal{U}_0(t)} d_u(t) y_{u,d}(t), \qquad t \in \mathcal{T}. \tag{9}$$

Hence, $R(t)/D_{\text{left}}(t)$ is the fraction of residual demand recovered at slot t. For each possible restoration duration $\ell \in \mathcal{T}$, we define an auxiliary variable η_ℓ as the minimum recovered-demand fraction over the slots that would be experienced if restoration occurred at ℓ:

$$\eta_\ell = \min_{t \in \mathcal{T}:\, t \le \ell} \frac{R(t)}{D_{\text{left}}(t)}. \tag{10}$$

The quantity η_ℓ represents the service level guaranteed throughout the outage prefix $\{1, \ldots, \ell\}$. Unlike cumulative recovered-traffic objectives, which may favor high service in early slots while allowing later service degradation, the prefix minimum evaluates each restoration outcome by its weakest pre-restoration slot. Therefore, any temporary collapse in recovered demand directly reduces the value assigned to all restoration durations extending beyond that collapse. The metric consequently captures sustained recovery in residual demand rather than aggregate throughput. Since the minimum operator is nonlinear, relation (10) is enforced through linear inequalities. For candidate position d, the problem is

$$(\mathbf{P}): \quad \max_{y,p,E,\eta} \sum_{\ell \in \mathcal{T}} \pi_\ell \eta_\ell \tag{11}$$

s.t.

$$R(t) \ge \eta_\ell D_{\text{left}}(t), \qquad \forall \ell \in \mathcal{T},\ \forall t \in \mathcal{T} : t \le \ell, \tag{12}$$

$$R(t) \ge \rho_{\min} D_{\text{left}}(t), \qquad \forall t \in \mathcal{T}, \tag{13}$$

$$R(t) \le C_d, \qquad \forall t \in \mathcal{T}, \tag{14}$$

$$p(t) - PL_{u,d}(t) \ge P_{\text{rx}}^{\min} - M(1 - y_{u,d}(t)), \qquad \forall t \in \mathcal{T},\ \forall u \in \mathcal{U}_0(t), \tag{15}$$

$$0 \le p(t) \le p_{\max}, \qquad \forall t \in \mathcal{T}, \tag{16}$$

$$E(1) = E_0, \tag{17}$$

$$E(t+1) = E(t) - (P_{\text{fix}} + P(t))\,\Delta t, \qquad \forall t \in \mathcal{T}, \tag{18}$$

$$E(t) \ge 0, \qquad \forall t \in \mathcal{T}, \tag{19}$$

$$0 \le \eta_\ell \le 1, \qquad \forall \ell \in \mathcal{T}, \tag{20}$$

$$y_{u,d}(t) \in \{0,1\}, \qquad \forall t \in \mathcal{T},\ \forall u \in \mathcal{U}_0(t). \tag{21}$$

Constraint (12) is the linear representation of the prefix minimum in (10). For each restoration duration ℓ, it forces η_ℓ to be no larger than the recovered-demand fraction in every slot $t \le \ell$. Maximizing $\sum_{\ell \in \mathcal{T}} \pi_\ell \eta_\ell$, (11), therefore maximizes the expected minimum service level over the random outage duration τ. The expectation is taken after evaluating the worst slot associated with each possible restoration duration, meaning that each realization of the restoration time is assessed according to the service continuity actually experienced before recovery. Constraint (13) imposes a deterministic per-slot service floor $\rho_{\min}$, independently of the restoration prior. Constraint (14) limits the recovered traffic by the UAV-BS capacity. Constraint (15) enforces link-budget feasibility for admitted users through a big-M implication: if $y_{u,d}(t) = 1$, the received-power condition must hold; otherwise the implication is deactivated. Constraint (16) imposes the transmit-power limit. Constraints (17)–(19) define the finite-energy evolution of the landed UAV-BS. Constraints (20) and (21) define the domains of the continuity and admission variables. For a fixed candidate position d, the formulation contains $\sum_{t \in \mathcal{T}} |\mathcal{U}_0(t)|$ binary admission variables. The stochastic restoration model introduces $T_{\max}(T_{\max}+1)/2$ continuity constraints through (12). Thus, the number of continuity constraints induced by restoration uncertainty grows quadratically with the planning horizon, while the admission complexity scales with the number of residual-demand users. The resulting problem is therefore a restoration-aware MILP (RA-MILP), solved in the numerical analysis using the Gurobi optimizer.

V. Experimental Results

The proposed framework is evaluated using realistic urban morphology, real BS deployments, and data-driven residual traffic demand. As shown in Fig. 2, we consider the three distinct spatially correlated outage realizations produced by the proximity-based blackout-generation procedure for the case in which 70% of the 10 terrestrial BSs are unavailable. For each outage case, the UAV-BS is scheduled over a maximum recovery horizon of $T_{\max} = 10$ hourly slots.

The stochastic restoration process is represented through three discretized lognormal distributions: *Short*, *Moderate*, and *Prolonged*, as depicted in Fig. 3.

For each restoration-time distribution and each outage realization, the RA-MILP and benchmark methods are solved independently for every candidate UAV-BS position. Results reported in Section V are averaged over the three outage realizations to reduce dependence on a specific failed-BS set and to evaluate candidate UAV-BS locations under different blackout footprints with the same outage severity.

The candidate set of UAV-BS locations follows the construction described in Section III-B. It includes three cluster-based rooftop locations (CL1, CL2, and CL3); the rooftop with maximum absolute altitude (MAXALT); the rooftop closest to the blackout centroid (CEN); and a hovering point above the blackout centroid (HOV). The hovering point is retained only as a geometric reference. Its hovering propulsion cost is not included in the energy model; therefore, it should be interpreted as a placement upper-bound reference rather than as a directly comparable deployment option. A full energy-budget analysis including propulsion energy for hovering or relocation phases is left for future work. The simulation parameters are summarized in Table I.

A. Scenario

The simulation is initialized with the following inputs.

Urban morphology: A 3D shapefile of Zurich provides building geometry, including ground elevation, planar coordinates, and height. It enables accurate LoS/NLoS classification.

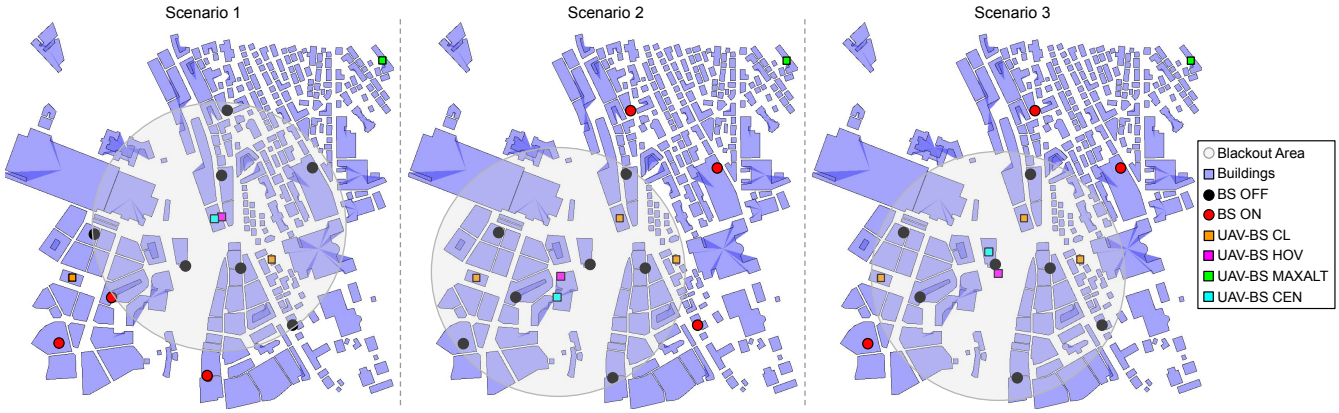


Fig. 2: Three representative outage realizations with 70% unavailable BSs and candidate UAV-BS locations.

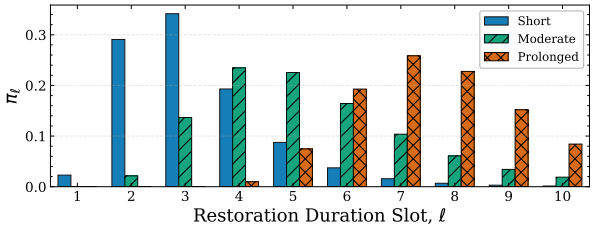


Fig. 3: Restoration-duration priors over the planning horizon.

TABLE I: Main simulation parameters.

BS Parameters	
Frequency	2100 MHz
Bandwidth	100 MHz
Transmit Antenna Gain	18 dBi
Rx Antenna Element Gain	0 dBi
Radiated Power	49 dBm
Sampling Factor	1.536
Rx Noise Figure	8 dB
Rx Antenna Height	1.5 m
UAV-BS Parameters	
Battery Capacity E_0	0.526 kWh
Fixed Power P_{fix}	0.02 kW
Max Transmit Power $p_{\max}$	1 W
Capacity C	10 Mbps
Time Slot Δt	1 hour
Service floor $\rho_{\min}$	0.5

BS deployment: Real BS coordinates[1] are used within a 1 km^2 area of Zurich. Only macro cells are retained, yielding 10 active terrestrial BSs under normal operation.

UAV-BS deployment: A single UAV-BS is considered, with battery modeled after the DJI Matrice 350 RTK dual-battery system[2] equipped with TB65 intelligent batteries[3].

Traffic demand: Demand traces are drawn from the Net-Mob23 dataset [16], which provides real-world mobile traffic data aggregated over 0.01 km^2 tiles at hourly resolution. Since the traces are normalized, we scaled them to the Zurich scenario using a population density of $5{,}084$ people/km^2. The resulting traffic demand is computed assuming a peak activity factor of 4 Mbit/s and 50% active users.

[1] https://www.bakom.admin.ch/en/location-of-radio-transmitters
[2] https://enterprise.dji.com/matrice-350-rtk/specs
[3] https://store.dji.com/it/product/tb65-intelligent-flight-battery

TABLE II: Geometric and feasibility characterization of landed UAV-BS candidates across the three outage realizations.

Candidate	Distance [m]	Altitude [m]	Feasibility [%]
CEN	49.4	425.3	100
CL2	144.7	429.2	100
CL3	256.8	438.8	100
CL1	351.5	433.0	66.7
MAXALT	797.3	508.0	33.3

B. Benchmark Methods

We compare RA-MILP with two deterministic MILP baselines and one greedy policy. All MILP methods use the same admission variables, transmit-power variables, battery dynamics, capacity, received-power, and service-floor $\rho_{\min}$ constraints; they differ only in how restoration duration is modeled in the objective.

Full-Horizon MILP (FH-MILP): This baseline represents a worst-case deterministic planning policy. It assumes that the outage lasts until the end of the planning horizon $T_{\max}$ and maximizes the worst recovered-demand fraction over all slots $t \in \mathcal{T}$. It therefore protects delayed restoration cases without using the restoration-time distribution.

Expected-Horizon MILP (EH-MILP): This baseline represents a mean-duration deterministic planning policy. It replaces the random restoration time by $T_{\text{exp}} = \text{round}(\mathbb{E}[\tau])$ and solves the same deterministic worst-slot problem as FH-MILP over the restricted horizon $\{1, \dots, T_{\text{exp}}\}$. Service after T_{exp} is not explicitly optimized, reflecting a policy planned around an average repair-time estimate.

Greedy Energy-Feasible Admission (GEA): This policy operates slot by slot over $T_{\max}$. At each slot, reachable users are sorted by demand and admitted while the UAV-BS capacity, received-power condition, and residual battery permit service. The transmit power is then set to the minimum value required by the admitted users. GEA has no lookahead over future slots and does not enforce the service floor $\rho_{\min}$.

C. Candidate Position Assessment

To better interpret the position-dependent results, Table II reports both the geometric properties and the feasibility rate of landed UAV-BS candidates. The feasibility rate indicates

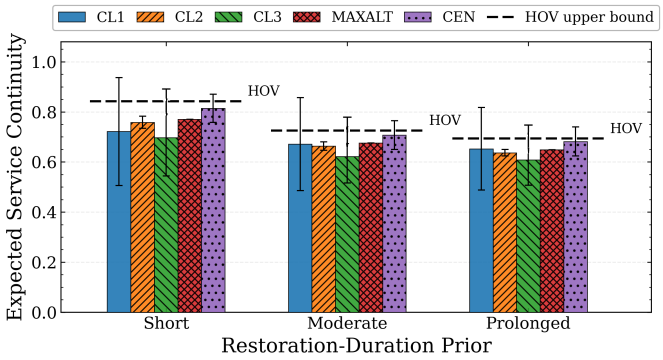


Fig. 4: Expected service continuity across UAV-BS candidate positions under the three restoration-time priors.

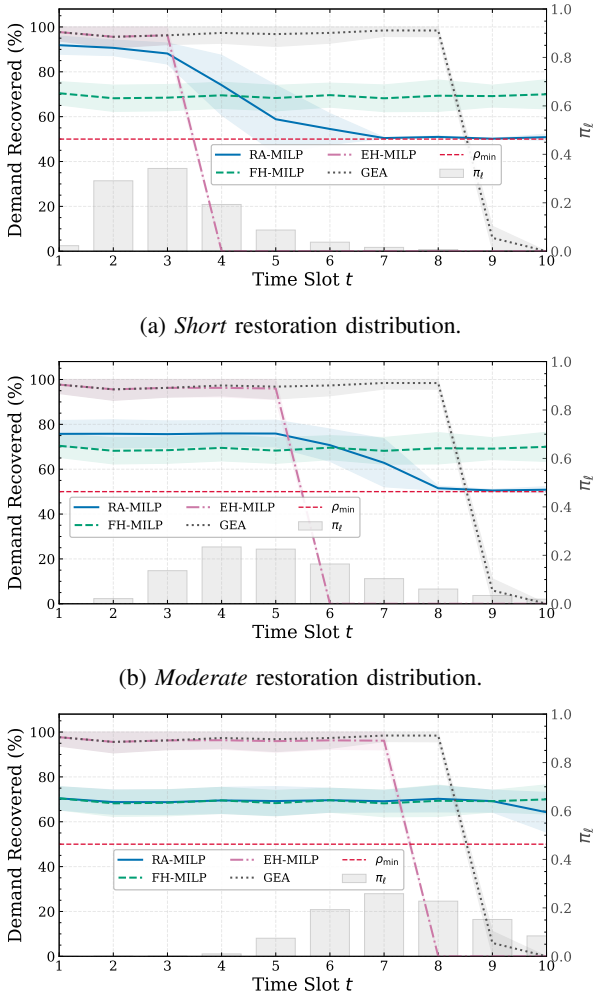
the fraction of outage realizations for which the optimization problem admits a solution satisfying the service-floor, capacity, link-budget, and battery constraints. Distance is measured from the outage centroid. The table shows that CEN is feasible in all realizations and is only 49.4 m from the outage centroid, whereas MAXALT has the highest altitude but lies 797.3 m away and is feasible in only 33.3% of the cases. This confirms that altitude alone is not sufficient: effective landed recovery requires spatial alignment with the residual-demand footprint.

Fig. 4 compares the expected service continuity of the RA-MILP across candidate UAV-BS positions and restoration-duration priors. The hovering reference HOV provides the highest values because of its greater positional flexibility, but, as discussed above, it represents an upper-bound reference rather than a complete hovering deployment model. Among landed candidates, CEN gives the most stable performance, while cluster-based locations provide outage-agnostic alternatives when the blackout footprint is not accurately known. As the prior shifts from *Short* to *Prolonged*, expected continuity decreases for all positions, reflecting the increasing cost of sustaining service over later slots under finite UAV-BS energy.

D. Service Continuity Under Uncertain Restoration

A resilient recovery policy should restore service without exhausting the resources needed if the outage lasts longer than expected. This tradeoff is visible in Fig. 5. GEA provides high early recovery because it admits users whenever the current capacity, link budget, and battery state allow it. This is effective for short outages, but it does not preserve enough energy for delayed restoration. As a result, GEA falls below $\rho_{\min}$ in late slots and provide no recovery. EH-MILP represents recovery planning around a mean repair horizon. It provides strong service before $T_{\exp}$, but delayed restoration outcomes are not explicitly protected. FH-MILP takes the opposite approach by protecting the full horizon, which is robust to delayed repair but conservative when recovery is likely to occur early.

RA-MILP adapts between these regimes. Under the *Short* prior, it allocates more service to early slots than FH-MILP because early restoration is likely. Under the *Prolonged* prior, it approaches FH-MILP because maintaining service in late slots becomes essential. The transmit-power profiles in Fig. 6



(a) *Short* restoration distribution.

(b) *Moderate* restoration distribution.

(c) *Prolonged* restoration distribution.

Fig. 5: Per-slot residual-demand recovery of the compared scheduling policies under the three restoration-time priors.

explain this adaptation: RA-MILP front-loads power under *Short* restoration and moves toward sustained power allocation under *Prolonged* restoration. Therefore, RA-MILP avoids the late collapse of myopic recovery while also avoiding unnecessary worst-case conservatism when the restoration probability favors early repair. Since policies with similar average recovered demand may nevertheless differ in their ability to maintain $\rho_{\min}$, the next subsection evaluates their sensitivity using a compliance-weighted recovery metric.

E. Sensitivity to the Service-Floor Requirement

The previous results show that policies differ not only in the amount of residual demand they recover, but also in their ability to maintain service above the required floor over time. To quantify this effect while preserving the restoration-aware nature of the analysis, we introduce a penalized expected service-continuity metric. We define the prior-weighted service continuity as

$$R_{\text{prior}} = \sum_{\ell=1}^{T_{\max}} \pi_\ell \eta_\ell.$$

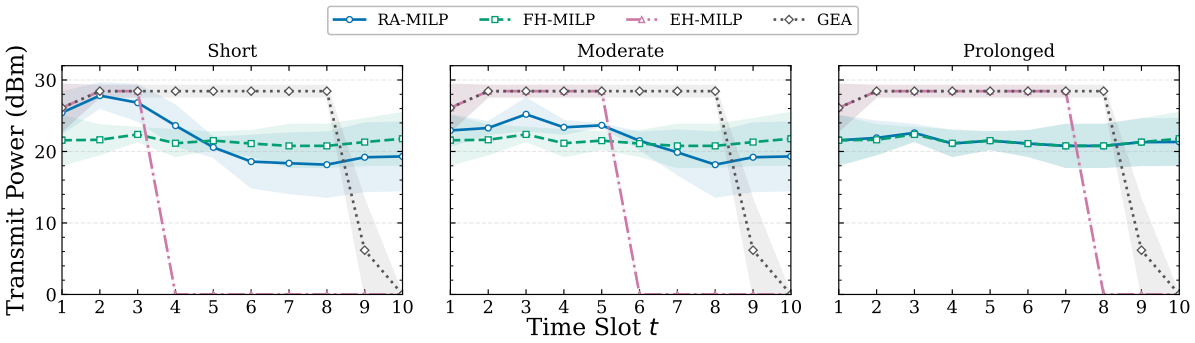


Fig. 6: Per-slot transmit-power allocation of the compared scheduling policies under the three restoration-time priors.

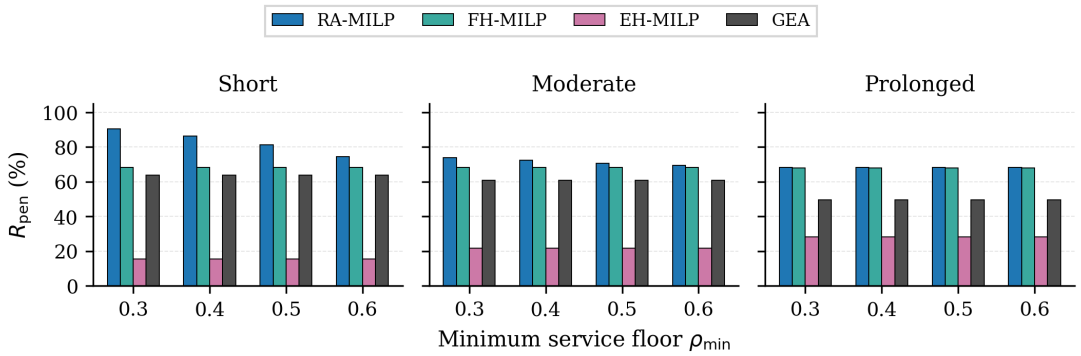


Fig. 7: Penalized expected service continuity at CEN under different minimum service-floor requirements.

This term is consistent with the restoration-aware objective, since each possible restoration duration is evaluated through the minimum service level experienced before restoration. However, R_{prior} alone does not explicitly penalize intermittent violations of the minimum service floor. We therefore define the service-floor violation ratio as

$$V_\rho = \frac{1}{T_{\max}} \sum_{t=1}^{T_{\max}} \mathbf{1}\left\{r(t) < \rho_{\min}\right\}.$$

The final penalized service-continuity metric is

$$R_{\text{pen}} = R_{\text{prior}} \exp\left(-\lambda V_\rho\right),$$

where λ controls the severity of the penalty. We use $\lambda = 2$, which penalizes below-threshold slots while avoiding a binary feasible/infeasible classification. Thus, a policy is rewarded when it maintains a high prefix-minimum recovery under the restoration prior, and is penalized when it violates the required service floor during the planning horizon.

Since the position analysis identified CEN as the most reliable landed candidate, we evaluate the sensitivity to $\rho_{\min}$ at this position. Fig. 7 reports R_{pen} for $\rho_{\min} \in \{0.3, 0.4, 0.5, 0.6\}$ under the *Short*, *Moderate*, and *Prolonged* restoration priors. RA-MILP achieves the highest score under the *Short* and *Moderate* priors because it exploits the restoration-time distribution while still preserving service above the required floor. Under the *Prolonged* prior, RA-MILP approaches FH-MILP, reflecting the need to protect later slots. FH-MILP remains stable across all priors because it protects the full horizon, but it can be conservative when early restoration is likely. EH-MILP is strongly penalized because service after the expected repair time is not explicitly protected, while GEA loses performance because its high early recovery comes at the cost of later service degradation. These results confirm that blackout resilience must account not only for average recovered demand, but also for $\rho_{\min}$ compliance over time.

F. Sensitivity to Restoration-Time Prior Mismatch

After assessing the impact of $\rho_{\min}$ requirement, we assess the cost of optimizing the schedule with an incorrect restoration-time prior. For each true prior, Fig. 8 reports the expected continuity loss relative to the schedule optimized with the correct prior. The largest loss occurs when a *Short* prior is assumed but restoration is actually *Prolonged*, causing about 24% expected continuity loss. Using the *Moderate* prior in the same true *Prolonged* case reduces the loss to about 12%, showing the value of preserving some late-horizon service when repair is delayed. The reverse mismatch is less severe but still visible: if restoration is actually *Short*, planning with the *Prolonged* prior loses about 16%, while the *Moderate* prior loses about 8%. Thus, mismatch is most damaging when the assumed prior assigns too little probability to the restoration times that actually occur. This suggests a conservative modeling rule: when repair-time information is uncertain, overly optimistic fast-restoration assumptions are riskier than priors that retain some probability mass in the tail.

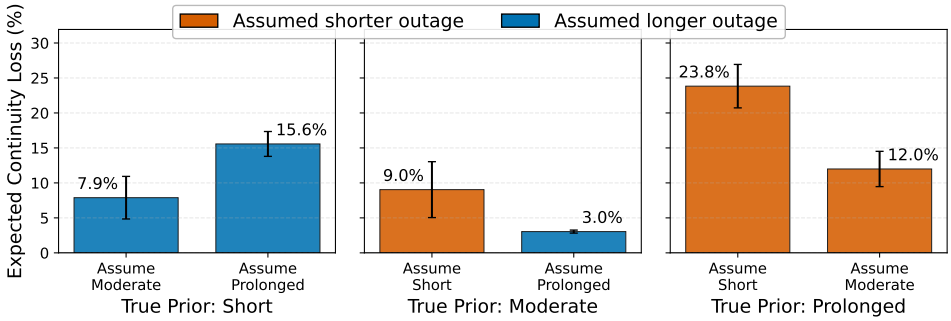


Fig. 8: Expected service-continuity loss caused by restoration-prior mismatch.

G. Deployment Perspective

The proposed optimization may be implemented as a long-horizon planning and re-optimization function in an Open Radio Access Network (O-RAN) architecture. Since landing-site selection and UAV-BS scheduling rely on slowly varying information, such as outage reports, residual-demand forecasts, battery state, and urban morphology, the RA-MILP is naturally suited to the Service Management and Orchestration (SMO) layer and the Non-Real-Time RAN Intelligent Controller (Non-RT RIC). During the blackout, updated restoration priors or demand forecasts can trigger re-optimization, while faster admission or transmit-power corrections can be handled by Near-Real-Time RAN Intelligent Controller (Near-RT RIC) applications. Thus, the proposed framework should be interpreted as a blackout-time planning tool rather than as a slot-level radio scheduler.

VI. Conclusion

This paper studied landed UAV-BS recovery for RAN service continuity during power-grid blackouts with uncertain restoration time. We formulated a restoration-aware MILP that schedules user admission and transmit power under finite battery, link-budget, capacity, and service-floor constraints. The results show that restoration-time priors significantly affect the optimal recovery profile: short-repair priors favor earlier service, while delayed-repair priors require more conservative battery use. The results also show that landing position is critical, since proximity to the outage footprint can be more important than altitude alone. Finally, the service-floor and prior-mismatch analyses show that maximizing recovered traffic is not sufficient when sustained service continuity is required. Future work will extend the framework to closed-loop recovery with updated restoration priors, coordinated multi-UAV deployment, and complete energy-budget modeling including propulsion energy.

Acknowledgment

This paper was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”, Focused Project R4R, 6GWINET and the PRIN 2022 PNRR M4.C2.1.3 “BREADCRUMBS” project, grant number P2022K7ERB, CUP J53D23014990001.

References

[1] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, “Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines,” *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.

[2] European Union Agency for Cybersecurity (ENISA), “Telecom Security Incidents 2024,” ENISA Reports, Tech. Rep., 2025.

[3] A. Cabrera-Tobar, F. Grimaccia, and S. Leva, “Energy resilience in telecommunication networks: A comprehensive review of strategies and challenges,” *Energies*, vol. 16, no. 18, 2023.

[4] B. Karaman, I. Basturk, S. Taskin, E. Zeydan, F. Kara, E. A. Beyazit, M. Camelo, E. Björnson, and H. Yanikomeroglu, “Solutions for sustainable and resilient communication infrastructure in disaster relief and management scenarios,” 2024.

[5] B. Galkin, J. Kibilda, and L. A. DaSilva, “Uavs as mobile infrastructure: Addressing battery lifetime,” *IEEE Communications Magazine*, vol. 57, no. 6, pp. 132–137, 2019.

[6] A. Martiny, M. Meo, D. Renga, and G. Vallero, “Towards resilient rans through uav integration: Landing positioning strategy during blackouts,” in *2025 8th International Conference on Advanced Communication Technologies and Networking (CommNet)*, 2025, pp. 1–7.

[7] A. I. Abubakar, M. S. Mollel, O. Onireti, M. Ozturk, I. Ahmad, S. M. Asad, Y. Sambo, A. Zoha, S. Hussain, and M. A. Imran, “Coverage and throughput analysis of an energy efficient uav base station positioning scheme,” *Computer Networks*, vol. 232, p. 109854, 2023.

[8] M. Y. Selim, A. Alsharoa, and A. E. Kamal, “Short-term and long-term cell outage compensation using uavs in 5g networks,” in *2018 IEEE GLOBECOM*, 2018, pp. 1–6.

[9] D. L. V. Dinh, A. N. T. Mai, H. Tran, G. Q. L. Vu, T. D. Ho, Z. Pan, V. N. Van, S. Chatzinotas, and D.-H. Tran, “Uav-assisted resilience in 6g and beyond network energy saving: A multi-agent drl approach,” *arXiv preprint arXiv:2511.07366*, 2025.

[10] B. Omoniwa, B. Galkin, and I. Dusparic, “Energy-aware optimization of uav base stations placement via decentralized multi-agent q-learning,” in *2022 IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*, 2022, pp. 216–222.

[11] G. Colajanni, P. Daniele, and D. Sciacca, “On the provision of services with uavs in disaster scenarios: A two-stage stochastic approach,” *Operations Research Forum*, vol. 3, no. 1, p. 18, 2022.

[12] S. Kancherla and I. Dobson, “Heavy-tailed transmission line restoration times observed in utility data,” *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 1145–1147, 2018.

[13] ETSI, “5g; study on channel model for frequencies from 0.5 to 100 ghz (3gpp tr 38.901 version 15.0.0 release 15),” *Tech. Rep., ETSI*, 2018.

[14] G. Auer, V. Giannini, C. Desset, I. Gódor, P. Skillermark, M. Olsson, M. A. Imran, D. Sabella, M. J. Gonzalez, O. Blume, and A. Fehske, “How much energy is needed to run a wireless network?” *IEEE Wireless Communications*, vol. 18, no. 5, pp. 40–49, 2011.

[15] G. Vallero, G. Perin, M. Meo, M. S. Garino, S. D’Elia, and D. Vaccarono, “A new explainable power demand model for 4g lte and 5g nr base stations,” in *ICC 2025-IEEE International Conference on Communications*. IEEE, 2025, pp. 4173–4178.

[16] O. E. Martínez-Durive, S. Mishra, C. Ziemlicki, S. Rubrichi, Z. Smoreda, and M. Fiore, “The netmob23 dataset: A high-resolution multi-region service-level mobile data traffic cartography,” 2023.