Cooperative Risk-Sensitive Survivor-Hub Planning for Resilient Multi-Operator RANs

Mohammad Reza Jokar, Greta Vallero, Daniela Renga, and Michela Meo
Department of Electronics and Telecommunications, Politecnico di Torino, Turin, Italy
Email: {mohammadreza.jokar, greta.vallero, daniela.renga, michela.meo}@polito.it

***Abstract*—Power outages can disrupt dense urban radio access networks (RANs), particularly when several nearby base-station (BS) sites become unavailable simultaneously. Since equipping every site with long-duration backup energy is costly, this paper investigates the selection of battery-backed *survivor hubs* in a shared multi-operator RAN. These hardened sites remain operational during outages and recover traffic from failed neighboring sites, including those owned by other operators. We formulate hub selection, battery sizing, and scenario-dependent traffic recovery as a two-stage stochastic mixed-integer linear program. The model uses a Conditional Value-at-Risk objective at the 95% confidence level ($\mathrm{CVaR}_{0.95}$), augmented by site-relevance, operator-balance, and recovery-oriented penalties, to limit priority-weighted service loss in severe outage scenarios. Radio-capacity and battery-energy constraints account for the combined native and migrated traffic carried by each hub, while outage duration determines the energy required to maintain operation. We evaluate the framework in a fixed-seed planning case study based on real multi-operator deployment records from Turin, Italy, combined with synthetic traffic and spatially correlated outage scenarios. At the 500 k€ reference budget, the proposed method achieves lower empirical tail loss than No Cooperation and the other evaluated reference methods. It is also the only evaluated method whose aggregate demand-weighted recovery exceeds the configured thresholds for all three traffic classes. These results support cooperative survivor-hub planning under the stated case-study assumptions, without constituting an operational performance forecast.**

***Index Terms*—RAN resilience, multi-operator network sharing, stochastic optimization, backup-energy planning.**

I. Introduction

Mobile networks support services that are expected to remain available during emergencies, including industrial IoT, transportation, emergency response, and digital healthcare. The growing role of 5G, and in the future 6G, makes service continuity increasingly important. However, the radio access network (RAN) remains vulnerable to failures in the power grid. A prolonged outage affecting several nearby base-station (BS) sites can substantially reduce connectivity across an urban area. Maintaining essential network functions during severe and uncertain disruptions is therefore an important objective in the design of resilient mobile networks [1], [2].

Installing backup batteries at every BS site is a direct way to improve resilience, but it can be costly in dense urban networks, where neighboring sites often have partially overlapping service areas. An alternative is to harden a selected subset of sites, which we refer to as *survivor hubs*. During a power outage, these sites remain active using backup energy and serve part of the traffic reassigned from failed neighboring sites. Selecting the hubs alone is not sufficient. A hub may have enough radio capacity to accept additional traffic but still exhaust its battery if the installed energy capacity was determined using only its native load. The planning problem must therefore decide both which sites to harden and how much backup energy to install, while accounting for native demand, reassigned emergency traffic, and uncertain outage conditions.

Multi-operator sharing can further improve recovery because an operational site may serve traffic originally associated with another operator. Several studies have examined related forms of cooperation and post-failure recovery. Backup-power sharing among co-located cellular towers can reduce emergency-power provisioning costs [3]. National roaming and infrastructure sharing can also improve network robustness under random and geographically correlated failures [4], [5]. Other approaches use UAV-assisted outage compensation [6], while energy-side flexibility has been investigated through vehicle-to-grid resources for dynamic cooperative energy and coverage management [7]. Network-sharing architectures, renewable-energy cooperation, sustainable operation, and fault recovery have also been considered in [8]–[11]. These studies address important aspects of energy sharing, roaming, and operational recovery, but they do not jointly select terrestrial sites for hardening, allocate backup batteries, and determine cross-operator traffic recovery under uncertain, spatially correlated outages. This joint planning problem is the focus of this paper.

We formulate survivor-hub selection, battery sizing, and traffic recovery as a two-stage stochastic Mixed-Integer Linear Program (MILP). In the first stage, the model selects the sites to harden and determines their battery capacities before the outage scenario is known. After a scenario is revealed, the second stage assigns traffic to available hubs subject to coverage-overlap, slice-compatibility, radio-capacity, and battery-energy constraints. The second-stage decisions constitute an offline planning representation in which the scenario parameters, including outage duration, are available when recovery is evaluated. The formulation should therefore not be interpreted as an online controller with advance knowledge of the duration of an actual outage.

The primary objective term minimizes the Conditional Value-at-Risk at the 95% confidence level, $\mathrm{CVaR}_{0.95}$, of priority-weighted service loss. This term places greater em-

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

phasis on high-loss scenarios than an expected-loss objective. The implemented objective also accounts for recovery-target shortfalls, operator-level recovery balance, URLLC loss, and the relevance of selected sites. The recovery model distinguishes among Ultra-Reliable Low-Latency Communications (URLLC), enhanced Mobile Broadband (eMBB), and massive Machine-Type Communications (mMTC) traffic [12], [13].

We evaluate the framework in a fixed-seed planning case study based on real spatial deployment records from three operators in Turin, Italy, combined with synthetic traffic and spatially correlated outage scenarios. At the reference budget of 500 k€, the proposed method achieves lower empirical tail loss than the evaluated reference methods and obtains the highest aggregate recovery rate for each traffic class. It is also the only evaluated method whose aggregate demand-weighted recovery exceeds the configured thresholds for all three classes. These results describe the complete cooperative design, including hub placement, battery allocation, and cross-operator recovery; they should not be interpreted as isolating the effect of battery sizing alone or as an operational performance forecast.

The main contributions are:

- A cooperative survivor-hub planning model for multi-operator RAN resilience, in which hardened sites can support traffic recovery across operator boundaries.
- Joint optimization of survivor-hub selection and backup-battery allocation while accounting for both native and reassigned emergency traffic.
- A two-stage stochastic MILP with a $\mathrm{CVaR}_{0.95}$-centered objective for limiting priority-weighted service loss under severe outage scenarios.
- A fixed-seed case study based on a real three-operator deployment in Turin, with comparisons against No Cooperation and heuristic reference methods.

The remainder of this paper is organized as follows. Section II formalizes the problem setting. Section III presents the two-stage stochastic MILP formulation. Section IV describes the traffic model and outage-scenario generation. Section V introduces the Shapley-based cost-allocation procedure. Sections VI and VII present the evaluation setup and numerical results, respectively. Section VIII concludes the paper.

II. Problem Overview

We consider a shared multi-operator RAN comprising a set $\mathcal{I}$ of candidate operator-specific BS sites. Each site $i \in \mathcal{I}$ represents the infrastructure of one operator and has a unique owner $o_i \in \mathcal{O}$, where $\mathcal{O}$ is the set of participating operators. Sites deployed by different operators at the same or nearby locations are represented as distinct operator-specific sites rather than as a single multi-owner site. Cross-operator sharing is modeled through recovery decisions that allow a survivor hub to serve traffic originally associated with an unavailable site of another operator, subject to coverage and slice compatibility. Traffic is partitioned into the standard 5G service classes $\mathcal{K} = \{\text{URLLC}, \text{eMBB}, \text{mMTC}\}$ [12], [13]. Each scenario $s \in \mathcal{S}$ specifies the service-unavailable sites $\mathcal{F}_s \subseteq \mathcal{I}$, the outage duration T_s, and the slice-dependent demand $L_{j,s,k}$ originally served by each $j \in \mathcal{F}_s$. A site in $\mathcal{F}_s$ cannot act as a recovery hub; loss of grid supply at other sites is handled through the battery-energy constraints.

The planning problem is structured in two stages, introduced here and formalized in Section III.

- **First stage (design):** Before the outage scenario is known, the model selects the candidate sites to harden as survivor hubs and determines their installed battery capacities. These decisions are represented by $z_i \in \{0,1\}$ and $E_i \geq 0$, respectively, and together define the deployment plan $\{(z_i, E_i) : i \in \mathcal{I}\}$.
- **Second stage (recovery):** After scenario s is realized, the deployment remains fixed and the model determines how much traffic from each unavailable site is redirected to the selected hubs. The variable $x_{i,j,s,k} \in [0,1]$ denotes the fraction of slice-k demand from site j migrated to hub i, while $u_{j,s,k} \geq 0$ denotes the remaining unserved demand. This stage is an offline recourse representation in which the scenario parameters are available when recovery is evaluated.

A hub can support an unavailable site only when four conditions hold simultaneously: 1) the hub is selected ($z_i = 1$); 2) the source site is unavailable ($j \in \mathcal{F}_s$); 3) the hub remains available ($i \notin \mathcal{F}_s$); and 4) the sites have overlapping coverage and the hub supports the relevant traffic slice. Here, $O_{ij} \in [0,1]$ denotes the fraction of the service area of source site j overlapped by receiver i, and $A_{i,j,k} \in \{0,1\}$ indicates whether site i can serve slice-k traffic from site j. Migration is therefore possible only if $O_{ij} > 0$ and $A_{i,j,k} = 1$. When these conditions hold, a hub owned by o_i may serve traffic associated with another operator, $o_j \neq o_i$.

Two resource limits govern recovery. The *radio-capacity limit* prevents a hub from being overloaded by its native traffic and migrated demand. The *battery-energy limit* ensures that it has sufficient installed energy to remain active for duration T_s while carrying that combined load. This coupling distinguishes the proposed model from sharing approaches that do not jointly dimension backup energy for migrated traffic [11]. Thus, a hub with sufficient radio capacity may still be unable to support a recovery assignment if its battery was sized only for native traffic.

The two stages are optimized jointly in a single stochastic MILP. The deployment is common to all scenarios, whereas recovery adapts to the realized scenario. The complete formulation is presented next, and the subsequent allocation of the cooperative investment cost is described in Section V.

III. Risk-Sensitive Survivor-Hub Optimization

A. Decision Variables and Feasible Migration Pairs

The first-stage variables are the hub-selection indicator $z_i \in \{0,1\}$ and the installed battery capacity $E_i \geq 0$, defined for every $i \in \mathcal{I}$. These variables are scenario-independent: the deployment is fixed before the outage scenario is revealed.

For each scenario $s \in \mathcal{S}$, three second-stage variables describe traffic recovery. The fraction of slice-k demand

TABLE I
MAIN NOTATION USED IN THE OPTIMIZATION MODEL.

Symbol	Meaning
Sets and indices	
$\mathcal{I}$	Set of candidate operator-specific BS sites.
$\mathcal{O}$	Set of participating operators.
$\mathcal{K}$	Set of traffic slices.
$\mathcal{S}$	Set of outage scenarios.
$\mathcal{F}_s$	Set of service-unavailable sites in scenario s.
$\mathcal{E}$	Set of feasible migration pairs (i, j).
$\mathcal{R}, \mathcal{I}_r$	Planning zones and candidate sites in zone r.
Input parameters	
o_i	Owner/operator of site i.
w_s	Probability weight of scenario s.
T_s	Outage duration in scenario s.
$f_{i,s}$	Equal to 1 if site i is unavailable in scenario s, and 0 otherwise.
O_{ij}	Fractional coverage overlap between receiver i and source site j.
$A_{i,j,k}$	Equal to 1 if site i can serve slice-k traffic from site j.
$L_{j,s,k}$	Slice-k demand originally served by unavailable site j in scenario s.
$L^{\mathrm{nat}}_{i,s,k}$	Native slice-k demand at site i in scenario s.
$C_{i,s}$	Available radio capacity of site i in scenario s.
a_k	Radio-resource coefficient of slice k.
P^{base}_i	Base power consumption of site i.
$p_{i,k}$	Incremental power per Mbps of slice k at site i.
V_k	Priority weight of unserved slice-k demand.
ρ_k	Minimum recovery target for slice k.
ω_{rec}	Penalty weight for recovery-target violations.
F_i	Fixed hardening cost of site i.
$c_{E,i}$	Unit cost of battery capacity at site i.
$E^{\max}_i$	Maximum installable battery capacity at site i.
B_{cap}	Total investment budget.
M^C_i, M^E_i	Big-M bounds for radio-capacity and battery-energy constraints.
β	CVaR confidence level.
Decision and auxiliary variables	
z_i	Equal to 1 if site i is selected as a survivor hub.
E_i	Installed battery capacity at site i.
$x_{i,j,s,k}$	Fraction of slice-k demand from unavailable site j migrated to hub i.
$u_{j,s,k}$	Unserved slice-k demand of site j in scenario s.
$\sigma_{j,s,k}$	Fractional slack for the slice recovery target.
$Y_{i,s,k}$	Total carried slice-k load at site i in scenario s.
ℓ_s	Priority-weighted service loss in scenario s.
η, ξ_s	Auxiliary variables for CVaR linearization.

migrated from unavailable site j to hub i is $x_{i,j,s,k} \in [0, 1]$. The corresponding unserved demand is $u_{j,s,k} \geq 0$, and $\sigma_{j,s,k} \geq 0$ is the fractional shortfall from the prescribed slice recovery target. Consistent with Section II, these scenario-indexed variables form an offline recourse representation after the scenario is revealed.

To keep the formulation compact, migration variables are defined only for *feasible migration pairs*:

$$\mathcal{E} = \{(i, j) : i, j \in \mathcal{I},\ i \neq j,\ O_{ij} > 0,\ \exists\, k \in \mathcal{K} \text{ s.t. } A_{i,j,k} = 1\}, \tag{1}$$

where $O_{ij} \in [0, 1]$ is the fraction of the source-site service area overlapped by receiver i, and $A_{i,j,k} \in \{0, 1\}$ indicates whether site i can serve slice-k traffic from site j. Thus, O_{ij} also bounds the fraction of source demand that can be migrated through pair (i, j). The set $\mathcal{E}$ may contain both intra-operator and cross-operator pairs; pairs with no overlap or no admissible slice are excluded. The variable domains are

$$z_i \in \{0, 1\}, \qquad \forall\, i \in \mathcal{I}, \tag{2}$$
$$0 \leq E_i \leq E^{\max}_i, \qquad \forall\, i \in \mathcal{I}, \tag{3}$$
$$0 \leq x_{i,j,s,k} \leq 1, \qquad \forall\, s \in \mathcal{S},\ (i, j) \in \mathcal{E},\ k \in \mathcal{K}, \tag{4}$$
$$u_{j,s,k} \geq 0, \qquad \forall\, s \in \mathcal{S},\ j \in \mathcal{F}_s,\ k \in \mathcal{K}, \tag{5}$$
$$0 \leq \sigma_{j,s,k} \leq 1, \qquad \forall\, s \in \mathcal{S},\ j \in \mathcal{F}_s,\ k \in \mathcal{K}, \tag{6}$$
$$\eta \in \mathbb{R}, \quad \xi_s \geq 0, \qquad \forall\, s \in \mathcal{S}. \tag{7}$$

B. Hub Activation, Budget, and Geographic Diversity

Battery capacity can be installed only at selected hubs:

$$E_i \leq E^{\max}_i z_i, \qquad \forall\, i \in \mathcal{I}. \tag{8}$$

The total investment is bounded by the available capital budget:

$$\sum_{i \in \mathcal{I}} (F_i z_i + c_{E,i} E_i) \leq B_{\mathrm{cap}}. \tag{9}$$

The candidate set is partitioned into planning zones $\mathcal{R}$, with $\mathcal{I}_r \subseteq \mathcal{I}$ denoting the candidates in zone r. Under limited budgets, the diversity requirement is applied to an affordable subset $\mathcal{R}_B \subseteq \mathcal{R}$, prioritizing zones with higher average candidate relevance:

$$\sum_{i \in \mathcal{I}_r} z_i \geq 1, \qquad \forall\, r \in \mathcal{R}_B. \tag{10}$$

When the budget permits, at least one hub is also selected for each participating operator. These conditions promote deployment diversity but do not guarantee continuous geographic coverage. Actual recovery remains subject to O_{ij}, $A_{i,j,k}$, radio capacity, and installed energy.

C. Traffic Balance and Slice Recovery Targets

For every unavailable site $j \in \mathcal{F}_s$, scenario $s \in \mathcal{S}$, and slice $k \in \mathcal{K}$, demand is divided between recovered and unserved portions:

$$\sum_{i:\,(i,j) \in \mathcal{E}} x_{i,j,s,k} L_{j,s,k} + u_{j,s,k} = L_{j,s,k}. \tag{11}$$

A minimum recovery fraction $\rho_k \in [0, 1]$ is prescribed for each slice:

$$\sum_{i:\,(i,j) \in \mathcal{E}} x_{i,j,s,k} + \sigma_{j,s,k} \geq \rho_k, \quad \forall\, s \in \mathcal{S},\ j \in \mathcal{F}_s,\ k \in \mathcal{K}. \tag{12}$$

Because $x_{i,j,s,k}$ represents a fraction of $L_{j,s,k}$, constraint (12) imposes a per-site, per-scenario recovery floor for source–slice

combinations with positive demand. The slack $\sigma_{j,s,k}$ preserves feasibility when the target cannot be achieved and is penalized in the objective. Hence, ρ_k is a soft target rather than an unconditional feasibility guarantee.

D. Migration Feasibility

Traffic can be migrated from source site j to receiver i only when the receiver is selected as a hub, the source is unavailable, and the receiver remains operational:

$$x_{i,j,s,k} \le O_{ij} A_{i,j,k} z_i, \tag{13}$$
$$x_{i,j,s,k} \le O_{ij} f_{j,s}, \tag{14}$$
$$x_{i,j,s,k} \le O_{ij}(1 - f_{i,s}), \tag{15}$$

for all $s \in \mathcal{S}$, $(i,j) \in \mathcal{E}$, and $k \in \mathcal{K}$. Here, $f_{i,s} = 1$ denotes a service failure that installed backup energy cannot restore, rather than loss of utility supply alone. A site with $f_{i,s} = 0$ may still be disconnected from the grid for T_s, in which case its battery-energy constraint remains active. Constraints (13)–(15) enforce hub selection and compatibility, source-site unavailability, and receiver availability, respectively.

E. Radio-Capacity Constraint

The total slice-k load carried by hub i in scenario s is

$$Y_{i,s,k} = L^{\mathrm{nat}}_{i,s,k} + \sum_{j:\,(i,j)\in\mathcal{E}} x_{i,j,s,k} L_{j,s,k}, \tag{16}$$

where $L^{\mathrm{nat}}_{i,s,k}$ is the native slice-k demand at site i, set to zero when $i \in \mathcal{F}_s$. The weighted radio-resource demand must satisfy

$$\sum_{k\in\mathcal{K}} a_k Y_{i,s,k} \le C_{i,s} + M^C_i(1 - z_i + f_{i,s}), \qquad \forall\, i \in \mathcal{I},\ s \in \mathcal{S}. \tag{17}$$

Here, a_k converts slice-k traffic into radio-resource demand. The constraint is active only for a selected and operational hub.

F. Battery-Energy Constraint

The installed battery must sustain hub i for the outage duration T_s under its actual carried load:

$$E_i \ge T_s \left(P^{\mathrm{base}}_i + \sum_{k\in\mathcal{K}} p_{i,k} Y_{i,s,k} \right) - M^E_i(1 - z_i + f_{i,s}), \qquad \forall\, i \in \mathcal{I},\ s \in \mathcal{S}. \tag{18}$$

The big-M term deactivates the energy requirement for non-selected or unavailable sites. Replacing $Y_{i,s,k}$ by $L^{\mathrm{nat}}_{i,s,k}$ defines the radio-only sharing variant. If hub locations are reoptimized, however, this variant compares complete location-and-sizing designs and does not isolate battery sizing alone.

G. Risk Objective and CVaR Linearization

The priority-weighted service loss in scenario s is

$$\ell_s = \sum_{j\in\mathcal{F}_s} \sum_{k\in\mathcal{K}} V_k u_{j,s,k}, \tag{19}$$

where V_k is the priority weight of slice k. At $\beta = 0.95$, its CVaR is represented using the standard linear reformulation [14]:

$$\xi_s \ge \ell_s - \eta, \qquad \forall\, s \in \mathcal{S}, \tag{20}$$
$$\xi_s \ge 0, \qquad \forall\, s \in \mathcal{S}. \tag{21}$$

The implemented objective also includes the site-relevance reward C_{hub}, operator-balance penalty C_{fair}, recovery-shortfall penalty C_{sla}, and URLLC-specific loss penalty C_{U}, defined as

$$C_{\mathrm{hub}} = \sum_{i\in\mathcal{I}} S^{\mathrm{hub}}_i z_i, \qquad C_{\mathrm{sla}} = \sum_{s\in\mathcal{S}} \sum_{j\in\mathcal{F}_s} \sum_{k\in\mathcal{K}} \sigma_{j,s,k},$$
$$C_{\mathrm{U}} = \sum_{s\in\mathcal{S}} w_s \sum_{j\in\mathcal{F}_s} u_{j,s,\mathrm{URLLC}}, \quad C_{\mathrm{fair}} = \sum_{s\in\mathcal{S}} w_s \sum_{o\in\mathcal{O}} d_{o,s}. \tag{22}$$

Here, S^{hub}_i is the candidate-site score defined in Section IV-B. The nonnegative variable $d_{o,s}$ bounds the absolute difference between the fraction of unserved demand of operator o and the corresponding fraction for the other operators in scenario s; the absolute value is implemented by its standard pair of linear inequalities. The complete objective is

$$\begin{aligned} \min \quad & \eta + \frac{1}{1-\beta} \sum_{s\in\mathcal{S}} w_s \xi_s - \omega_{\mathrm{hub}} C_{\mathrm{hub}} + \omega_{\mathrm{fair}} C_{\mathrm{fair}} \\ & + \omega_{\mathrm{rec}} C_{\mathrm{sla}} + \omega_{\mathrm{U}} C_{\mathrm{U}}. \end{aligned} \tag{23}$$

The first two terms form the CVaR risk criterion. The remaining terms favor relevant candidate sites and penalize operator imbalance, recovery-target shortfalls, and unserved URLLC demand. The nonnegative coefficients ω_{hub}, ω_{fair}, ω_{rec}, and ω_{U} control their relative weights. The scenario weights satisfy $\sum_{s\in\mathcal{S}} w_s = 1$ and are described in Section IV.

H. Model Size

The dominant second-stage variables are the migration variables $x_{i,j,s,k}$. Defining them only for feasible pairs $\mathcal{E}$ reduces their count from $O(|\mathcal{I}|^2|\mathcal{S}||\mathcal{K}|)$ to $O(|\mathcal{E}||\mathcal{S}||\mathcal{K}|)$, thereby reducing the extensive-form MILP used for the Turin case study.

IV. Traffic Model and Scenario Generation

This section describes how the traffic and outage inputs of the MILP are generated. For each scenario, the model requires the native load of operational sites $L^{\mathrm{nat}}_{i,s,k}$, the recoverable demand of unavailable sites $L_{j,s,k}$, the unavailable-site set $\mathcal{F}_s$, the outage duration T_s, and the scenario weight w_s. Because measured per-site traffic and outage traces are not available, these inputs are constructed from the observed multi-operator deployment in Turin and the case-study assumptions described below. The objective is to produce heterogeneous, spatially grounded inputs for resilience planning, rather than to

reproduce individual user mobility or claim a fully empirical network trial.

A. Turin Deployment Data

The deployment dataset contains BS antenna locations, operator labels, radio-technology attributes, and the number of active cells per site. Each candidate site $i \in \mathcal{I}$ is therefore associated with a geographic location, an owner $o_i \in \mathcal{O}$, an active-cell count, and recorded radio attributes. The dataset does not contain measured traffic, battery telemetry, backhaul status, outage records, or deployed slice configurations; these quantities are modeled separately using inputs common to all compared methods.

The overlap coefficient $O_{ij} \in [0,1]$ is a planning proxy derived from inter-site distance and nominal service footprints. It bounds the fraction of source-site j's demand that is geographically reachable from receiver i, rather than representing a measured handover-success probability. The binary parameter $A_{i,j,k}$ represents the assumed technology and sharing-policy compatibility for serving slice k from source j at receiver i. Because the dataset does not report deployed slice configurations or roaming agreements, $A_{i,j,k}$ is a case-study input that would be replaced by operator-specific capability and policy information in an operational deployment.

B. Candidate-Site Scoring

Allowing every site in the full deployment to be hardened would make the MILP unnecessarily large (Section III-H). We therefore construct a balanced candidate set of 90 sites, with 30 sites per operator, using fixed-seed infrastructure-biased sampling. Sites with 5G availability, low-band support, or more active cells receive higher sampling probability, while the random component avoids treating this preprocessing step as an optimized deployment. The binary variable z_i subsequently determines the final survivor-hub selection among these candidates.

For each candidate site i, a normalized score is computed as

$$S_i^{\mathrm{hub}} = \sum_{q \in \mathcal{Q}} \alpha_q S_{i,q}, \tag{24}$$

where $\mathcal{Q}$ is the set of scoring criteria, $S_{i,q} \in [0,1]$ is the normalized value of criterion q, and $\sum_{q \in \mathcal{Q}} \alpha_q = 1$. The criteria and weights in Table II are obtained through an AHP-style pairwise comparison [15]. Low-band support and active-cell count receive the largest weights because they provide infrastructure-based proxies for coverage reach and load-handling capability. Other-operator proximity and critical-location proximity provide smaller spatial adjustments. The critical locations are synthetic planning references rather than a municipal critical-infrastructure inventory. The score is used to prioritize affordable planning zones and as a secondary site-relevance reward in the objective; it does not replace the CVaR criterion or determine the final deployment.

TABLE II
CRITERIA USED IN THE CANDIDATE-SITE SCORE.

Criterion q	**Symbol**	**Weight**
Low-band support	S_i^{prop}	0.364
Active cells	S_i^{ant}	0.258
Radio technology	S_i^{rad}	0.192
Other-operator proximity	S_i^{share}	0.084
Critical-location proximity	S_i^{crit}	0.102

C. Traffic Proxy

The MILP requires a traffic value for every site, scenario, and slice. In the absence of measured traffic traces, we construct a deployment-density proxy under the assumption that areas supported by more radio infrastructure tend to carry more aggregate demand. This assumption introduces spatial heterogeneity but is not presented as a calibrated traffic model for Turin.

For each site i, two normalized demand weights are computed: b_i^{own}, representing infrastructure density associated with the owning operator, and b_i^{tot}, representing the surrounding multi-operator emergency-demand density. The latter already incorporates the factor ϕ, so that only a fraction ϕ of normal traffic is retained as recoverable emergency demand during a major outage. Traffic inputs are generated as

$$L_{i,s,k}^{\mathrm{nat}} = \bar{L}_k \, b_i^{\mathrm{own}} \, m_k(h_s) \, \varepsilon_{i,s,k}^{\mathrm{nat}}, \tag{25}$$

$$L_{j,s,k} = \bar{L}_k \, b_j^{\mathrm{tot}} \, m_k(h_s) \, \varepsilon_{j,s,k}^{\mathrm{emg}}, \tag{26}$$

where $\bar{L}_k$ is the base traffic level of slice k, h_s is the outage onset hour, and $m_k(h_s)$ is its normalized hourly multiplier. The independent factors $\varepsilon_{i,s,k}^{\mathrm{nat}}$ and $\varepsilon_{j,s,k}^{\mathrm{emg}}$ introduce scenario-level traffic variation, with dispersion decreasing as active-cell count increases. To avoid extreme proxy values, emergency demand is capped at three times native demand, while native slice loads are proportionally scaled if their weighted total exceeds 90% of site radio capacity. Identical traffic realizations are used for all methods within each scenario.

D. Outage Scenario Generation

Each outage scenario $s \in \mathcal{S}$ is defined by the unavailable-site set $\mathcal{F}_s$, outage duration T_s, and onset hour h_s. Spatially correlated disruption states are generated around a sampled set of epicenters $\mathcal{P}_s$. For each epicenter $e \in \mathcal{P}_s$, the probability of site unavailability decreases with distance:

$$p_{j,e}^{\mathrm{fail}} = p_0 \exp\left(-\frac{d_{j,e}}{\lambda_{\mathrm{fail}}}\right), \tag{27}$$

where p_0 is the peak unavailability probability, $d_{j,e}$ is the distance between site j and epicenter e, and λ_{fail} controls the spatial spread. Epicenter sites are included in $\mathcal{F}_s$, and another site becomes unavailable if it is affected by at least one epicenter. This procedure produces clustered rather than independent site failures.

Consistent with Section III-D, $f_{j,s} = 1$ denotes a service-unavailable state that installed battery energy cannot restore;

TABLE III
MAIN SETTINGS USED FOR TRAFFIC AND OUTAGE GENERATION.

Quantity	Value
$\bar{L}_{\mathrm{URLLC}}$	50 Mbps
$\bar{L}_{\mathrm{eMBB}}$	200 Mbps
$\bar{L}_{\mathrm{mMTC}}$	1 Mbps
ϕ	0.20 of normal traffic
p_0	0.65
λ_{fail}	1.5 km
T_s	1–18 h
h_s	Sampled over 24 h
Raw scenarios	500 training + 100 diagnostic
$\lvert\mathcal{S}\rvert$	100 weighted representatives
Main random seed	2024

loss of utility supply alone does not set $f_{j,s} = 1$. Sites with $f_{j,s} = 0$ may nevertheless be disconnected from the grid for T_s and must satisfy the battery-energy constraint. The parameters p_0 and λ_{fail} define a case-study stress model and are not presented as fitted Turin failure statistics.

Using the fixed main seed 2024, 600 raw scenarios are generated: 500 form the scenario-reduction pool and 100 are reserved for diagnostic evaluation. A max–min diversity procedure over failure geography, outage duration, onset time, traffic, and capacity stress retains 100 representative scenarios, following the scenario-reduction principle in [16]. Each raw training scenario is assigned to its closest representative, and w_s is the resulting cluster frequency divided by 500. Consequently, $w_s \geq 0$ and $\sum_{s\in\mathcal{S}} w_s = 1$. All methods use the same representative scenarios, weights, and traffic realizations. The reported results describe this fixed-seed case study and are not interpreted as seed-independent statistical guarantees.

V. COST ALLOCATION

The optimization model selects survivor hubs and battery capacities for the multi-operator system but does not determine how the investment should be divided among operators. We therefore apply a separate post-processing step for cost allocation. This step does not affect the deployment or recovery decisions.

Let z_i^* and E_i^* denote the cooperative deployment obtained at the reporting budget. Its actual investment cost is

$$C^{\mathrm{tot}} = \sum_{i\in\mathcal{I}} \left(F_i z_i^* + c_{E,i} E_i^*\right). \tag{28}$$

The installed cost may be lower than the available budget if additional investment does not improve the objective.

We use the Shapley value [17]–[19] to quantify the operators' marginal contributions. For a nonempty coalition $\mathcal{C} \subseteq \mathcal{O}$, the coalition problem retains the sites and demand of its members and uses the same budget, scenario realizations, scenario weights, and operational rules as the full-cooperation problem. Let $\mathrm{CVaR}_{0.95}(\mathcal{C})$ denote the empirical tail loss of the resulting coalition deployment. For the empty coalition, we set $\mathrm{CVaR}_{0.95}(\emptyset) = \mathrm{CVaR}_{0.95}^{\mathrm{base}}$, where the latter is the empirical tail loss without hardened hubs. The coalition value is

$$v(\mathcal{C}) = \mathrm{CVaR}_{0.95}^{\mathrm{base}} - \mathrm{CVaR}_{0.95}(\mathcal{C}). \tag{29}$$

Thus, $v(\emptyset) = 0$, and a greater value indicates lower tail loss relative to the common no-hardening reference. Because coalition membership determines both the included demand and available sites, the resulting values are interpreted within this specified game.

The Shapley value of operator o, denoted by φ_o, is its average marginal contribution over all possible coalitions:

$$\varphi_o = \sum_{\mathcal{C}\subseteq\mathcal{O}\setminus\{o\}} \frac{|\mathcal{C}|!\,(|\mathcal{O}| - |\mathcal{C}| - 1)!}{|\mathcal{O}|!} \left[v(\mathcal{C}\cup\{o\}) - v(\mathcal{C})\right]. \tag{30}$$

This quantity is computed after the coalition problems are solved and is not an additional optimization objective.

The investment assigned to operator o is proportional to its Shapley contribution:

$$\chi_o = C^{\mathrm{tot}} \frac{\varphi_o}{v_{\mathcal{O}}}, \qquad v_{\mathcal{O}} = \sum_{o\in\mathcal{O}} \varphi_o. \tag{31}$$

By Shapley efficiency, $v_{\mathcal{O}} = v(\mathcal{O})$. In the evaluated instance, the grand-coalition value and all operator contributions are positive, so the allocation is budget-balanced: $\sum_{o\in\mathcal{O}} \chi_o = C^{\mathrm{tot}}$. Section VII compares this rule with traffic-proportional and BS-site-proportional allocations.

TABLE IV
PER-OPERATOR INFRASTRUCTURE CHARACTERISTICS IN THE TURIN STUDY AREA.

Metric	MNO1	MNO2	MNO3	Total
Sites	306	550	404	1,260
Active cells	3,572	3,595	4,140	11,307
Mean cells per site	11.67	6.54	10.25	–
Sites carrying a 5G tag	73%	45%	50%	53.5%
Sites with low-band support	99%	59%	92%	79.5%

VI. EVALUATION SETUP

This section summarizes the deployment data, candidate reduction, model settings, and reference methods used in the evaluation.

A. Dataset, Candidate Reduction, and Baselines

We evaluate the framework using spatial deployment records for three operators in Turin, Italy (MNO1–MNO3), comprising 1 260 sites and 11 307 active cells (Table IV). The data include site locations, operator labels, active-cell counts, and radio attributes. Traffic, radio-capacity, battery, and outage quantities are synthetic case-study inputs; the evaluation is therefore deployment-grounded rather than a fully empirical network trial.

The full deployment is used to construct the traffic proxy and candidate attributes. Using the procedure in Section IV-B, we form a balanced set of 90 candidate sites, with 30 per

operator. The MILP and outage scenarios are represented over these sites, while the surrounding deployment contributes to their demand weights. Candidate reduction therefore limits both possible battery locations and the spatial resolution of the planning model.

All methods use the same scenarios, weights, traffic inputs, resource parameters, recovery targets, and investment cap, and are evaluated using the same recovery and loss calculations. The slice-priority weights are $(V_{\mathrm{URLLC}}, V_{\mathrm{eMBB}}, V_{\mathrm{mMTC}}) = (500, 50, 10)$, and the recovery targets are $(0.65, 0.60, 0.30)$. The objective coefficients are $\omega_{\mathrm{hub}} = 50$, $\omega_{\mathrm{fair}} = 2000$, $\omega_{\mathrm{rec}} = 500$, and $\omega_{\mathrm{U}} = 200$. The budget ranges from 200 k€ to 900 k€ in 100 k€ increments, with 500 k€ used for detailed comparisons. We evaluate:

- **Proposed**: the cooperative stochastic MILP, jointly optimizing hub selection, battery sizing, and scenario-dependent recovery while allowing cross-operator traffic migration.
- **Greedy**: a cooperative heuristic that first assigns an affordable high-scoring hub to each operator and then adds hubs sequentially according to estimated tail-risk reduction and site relevance per unit cost. Its deployment is evaluated using the same operational model as Proposed.
- **No Cooperation**: a no-sharing MILP in which demand can be assigned only to hubs owned by the same operator. It retains the system-wide investment cap and isolates cross-operator migration within the centralized planning setting; it does not represent three independent operator planning processes.
- **Uniform**: a non-risk-aware reference that initially divides the budget equally among operators and selects spatially distributed sites without optimizing the $\mathrm{CVaR}_{0.95}$ objective. Any unused operator allocation is returned to a common budget pool.

Equation (18) also defines a radio-only variant in which battery energy is sized using native rather than combined native and migrated load. Because this variant reoptimizes both hub locations and battery capacities, it does not isolate the causal effect of battery sizing alone. A matched fixed-siting comparison would be required for that purpose; consequently, the reported gains are attributed to the complete cooperative planning framework.

VII. Results

A. Tail-Risk Reduction at the Reference Budget

Fig. 1 compares the distribution of per-scenario priority-weighted service loss at the 500 k€ reference budget. The horizontal axis reports loss in 10^3 units, and a leftward shift indicates better recovery. Proposed lies to the left of No Cooperation and Uniform over most of the displayed range and improves upon Greedy across much of the medium- and high-loss region. The longer right tails of No Cooperation and Uniform indicate larger residual losses in some representative outage scenarios. On the common representative set, Proposed also achieves a substantially lower empirical upper-tail loss than No Cooperation.

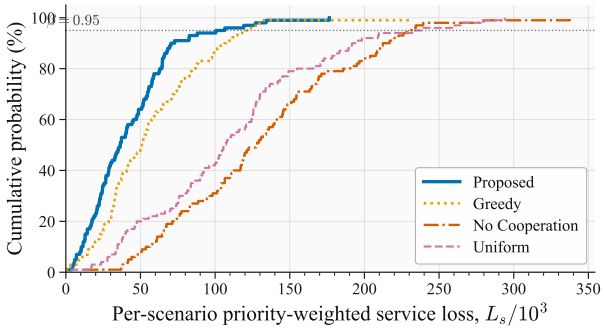


Fig. 1. Empirical CDF of priority-weighted service loss at $B_{\mathrm{cap}} = 500$ k€. The dotted line marks the 0.95 cumulative-probability level.

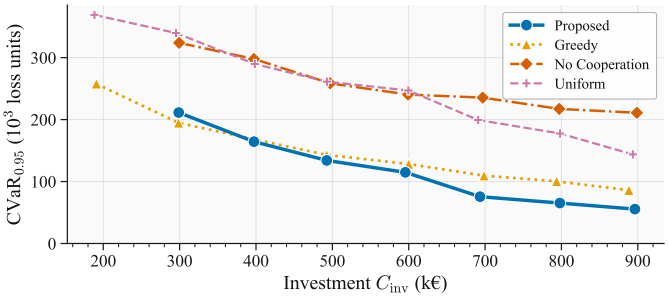


Fig. 2. Empirical $\mathrm{CVaR}_{0.95}$ across the investment sweep. Lower values indicate lower tail risk.

The CDF and reported empirical tail summaries give equal visual mass to the 100 representative scenarios, whereas the optimization itself uses the scenario weights w_s in (23). The horizontal 0.95 line is therefore a visual tail reference and does not identify the exact weighted CVaR tail. The observed ordering is consistent with the methods' construction: No Cooperation excludes cross-operator recovery, Uniform does not optimize tail risk, and Greedy selects hubs sequentially. Proposed instead coordinates hub selection, battery sizing, and scenario-dependent recovery. These findings describe the fixed-seed case study and do not imply dominance for every scenario sample or parameter setting.

B. Budget Sensitivity and Slice Recovery

Fig. 2 reports how empirical tail risk changes as the available budget increases. Each point is generated under a corresponding budget cap B_{cap}, while the horizontal axis reports the resulting investment $C_{\mathrm{inv}} \le B_{\mathrm{cap}}$. The first feasible Proposed deployment is obtained at 300 k€ under the imposed planning constraints. Greedy gives a slightly lower value at this tight budget, and the two cooperative methods are close at 400 k€. From 500 k€ onward, Proposed has the lowest reported tail risk, with the separation becoming clearer at higher budgets. No Cooperation remains above Proposed throughout their common feasible range and has the largest residual tail risk at high budgets. Thus, within this experiment, additional investment is used more effectively when placement, battery, and cross-operator recovery are coordinated, although the gain depends on the available budget.

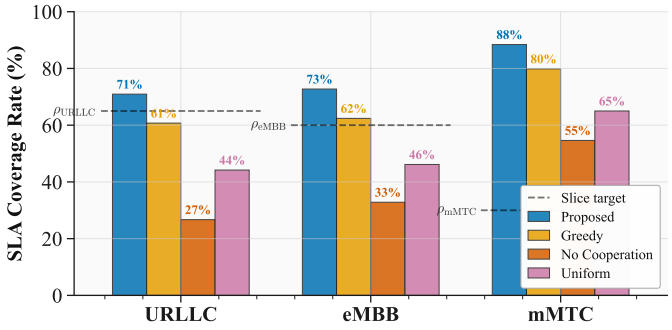


Fig. 3. Demand-weighted slice recovery at $B_{\text{cap}} = 500$ k€. Dashed lines show the configured targets.

TABLE V
COST ALLOCATION FOR THE PROPOSED DEPLOYMENT AT $B_{\text{cap}} = 300$ K€.

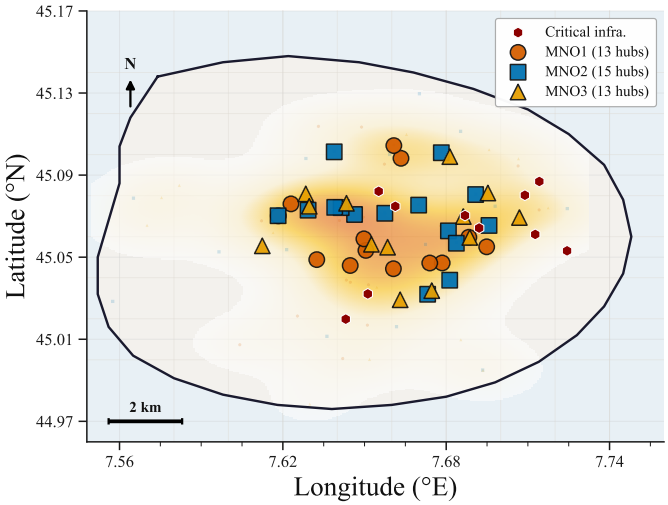
Operator	Shapley-based (k€)	City-wide traffic (k€)	City-wide sites (k€)
o_1	101.9	94.3	72.5
o_2	104.9	94.9	130.3
o_3	91.6	109.3	95.7

At $B_{\text{cap}} = 500$ k€, Fig. 3 compares aggregate demand-weighted recovery with the configured targets $\rho_{\text{URLLC}} = 65\%$, $\rho_{\text{eMBB}} = 60\%$, and $\rho_{\text{mMTC}} = 30\%$. Proposed is the only evaluated method that exceeds all three thresholds, recovering 71% of URLLC, 73% of eMBB, and 88% of mMTC demand. Greedy reaches 61%, 62%, and 80%, respectively, satisfying the eMBB and mMTC targets but remaining below the URLLC target. No Cooperation recovers 27%, 33%, and 55%, while Uniform recovers 44%, 46%, and 65%. Both satisfy the lower mMTC target but remain below the URLLC and eMBB thresholds. The largest recovery gains from cooperative planning occur for URLLC and eMBB, for which restricting traffic to same-operator hubs leaves fewer feasible recovery options. Together, Figs. 2 and 3 show that the lower tail risk is accompanied by improved recovery across the three traffic classes.

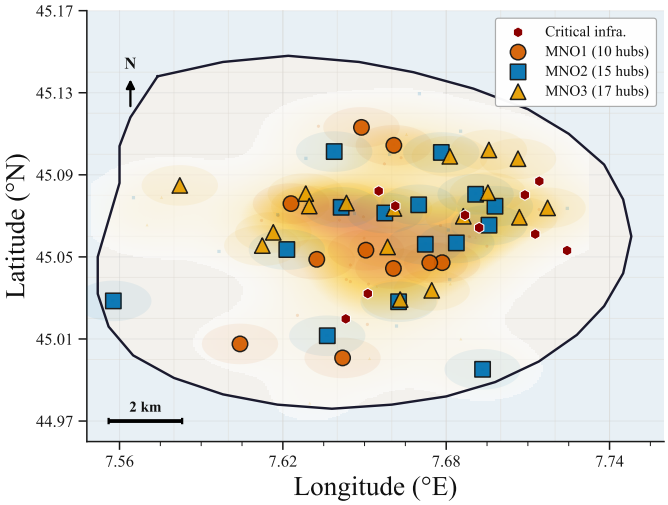
C. Spatial Deployment and Cost Allocation

Figs. 4a and 4b compare the spatial deployments at the tighter 300 k€ budget. Operator-specific markers denote selected hubs, while the small red points are synthetic critical-location references used in candidate scoring. Under No Cooperation, migrated demand can be assigned only to hubs owned by the same operator. Consequently, operator-specific recovery capacity is retained in several overlapping parts of the central high-demand region. Proposed can instead use a surviving hub to recover admissible demand from another operator, changing the spatial value of candidate sites and moving some selections beyond the central cluster.

The legends indicate 41 selected hubs for No Cooperation and 42 for Proposed. The maps therefore show a redistribution



(a) No Cooperation



(b) Proposed

Fig. 4. Survivor-hub deployments at $B_{\text{cap}} = 300$ k€: (a) No Cooperation and (b) Proposed.

of hardening locations rather than a reduction in the total number of hubs. They suggest that cross-operator recovery can reduce the need for nearby owner-specific selections in some overlapping areas while retaining hubs where the demand proxy is highest. Because no scalar spatial-overlap metric is reported, the maps illustrate different deployment patterns but do not independently establish a quantitative reduction in redundancy.

Table V and Fig. 5 apply three cost-sharing rules to the same Proposed deployment. Operators o_1, o_2, and o_3 denote the anonymized MNOs. The Shapley-based rule uses the marginal contributions defined by the coalition game in Section V, whereas the proportional references use aggregate traffic or BS-site shares. The label "City-wide Antennas" retained in Fig. 5 refers to this BS-site-count rule. The actual deployment cost is approximately 298.4 k€, slightly below the 300 k€ cap.

The Shapley-based payments are 101.9, 104.9, and 91.6 k€

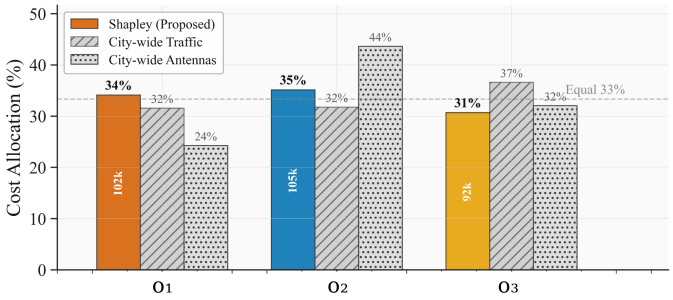


Fig. 5. Cost-allocation rules for the same Proposed deployment at $B_{\text{cap}} = 300$ k€. The dashed line indicates an equal split.

for o_1, o_2, and o_3, corresponding to approximately 34%, 35%, and 31% of the shared cost. The traffic-proportional allocation assigns approximately 32%, 32%, and 37%, whereas the BS-site-proportional rule assigns approximately 24%, 44%, and 32%, respectively. Minor differences in displayed totals arise from rounding. The larger site-based share assigned to o_2 reflects its greater number of site records in Table IV. These comparisons do not establish that one rule is universally preferable. They show instead that deployment and cost sharing are separate decisions: the physical solution remains fixed, but operator payments change with the selected accounting rule.

VIII. Conclusion

This paper presented a cooperative, risk-sensitive framework for planning battery-backed survivor hubs in multi-operator RANs during prolonged power outages. Hub selection, battery sizing, and slice-aware traffic recovery are formulated as a two-stage stochastic MILP whose primary term minimizes the scenario-weighted $\mathrm{CVaR}_{0.95}$ of priority-weighted service loss. The formulation permits cross-operator recovery when coverage and compatibility conditions are satisfied and accounts for both native and migrated emergency traffic in the battery-energy constraint. It therefore coordinates deployment and recovery decisions across severe, spatially correlated outage scenarios.

In the deployment-grounded Turin case study, which combines real multi-operator spatial records with synthetic traffic and outage inputs, the proposed method achieves substantially lower empirical upper-tail loss than the evaluated reference methods at the 500 k€ budget. It is also the only method whose aggregate demand-weighted recovery exceeds the configured targets for all three traffic classes. These improvements characterize the complete cooperative framework and should not be attributed to battery-aware sizing alone. The spatial results show that cross-operator recovery changes the placement of hardening investments and can reduce nearby owner-specific selections in some overlapping areas, without implying fewer selected hubs overall. The Shapley analysis further shows that deployment and cost sharing are separate decisions, as different allocation rules produce different payments for the same physical solution. These conclusions remain conditional on the representative scenarios and case-study assumptions. Future work will consider operational data, multi-seed out-of-sample evaluation, matched fixed-siting battery ablations, and partial-coalition settings.

References

[1] ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond," International Telecommunication Union, Tech. Rep. Rec. ITU-R M.2160-0, Nov. 2023.

[2] R.-J. Reifert, Y. Karacora, C. Chaccour, A. Sezgin, and W. Saad, "Resilience and criticality: Brothers in arms for 6G," 2024.

[3] M. N. H. Nguyen, N. H. Tran, M. A. Islam, C. Pham, S. Ren, and C. S. Hong, "Fair sharing of backup power supply in multi-operator wireless cellular towers," *IEEE Transactions on Wireless Communications*, vol. 17, no. 3, pp. 2080–2093, Mar. 2018.

[4] L. Weedage, S. Rangel, C. Stegehuis, and S. Bayhan, "On the resilience of cellular networks: How can national roaming help?" *IEEE Transactions on Network and Service Management*, vol. 21, no. 2, pp. 1702–1714, Apr. 2024.

[5] Z. Fang, G. Yang, W. Lyu, Z. Hong, S. Zhong, W. Zuo, Y. Xie, Y. Yang, G. Wang, Y. Liu, and D. Zhang, "Cellular infrastructure sharing for network robustness: A citywide empirical study," *IEEE Transactions on Mobile Computing*, vol. 24, no. 11, pp. 11 386–11 400, 2025.

[6] M. Y. Selim, A. Alsharoa, and A. E. Kamal, "Short-term and long-term cell outage compensation using UAVs in 5G networks," in *Proceedings of the IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2018, pp. 1–6.

[7] M. R. Jokar, A. Martiny, and M. Meo, "Dynamic cooperative energy and coverage management for V2G-enhanced RAN resilience," in *Proceedings of the 2026 IEEE International Conference on Pervasive Computing and Communications Workshops and Other Affiliated Events (PerCom Workshops)*. IEEE, 2026, pp. 946–951.

[8] E. Zeydan, S. S. Arslan, and Y. Turk, "Exploring blockchain architectures for network sharing: Advantages, limitations, and suitability," *IEEE Transactions on Network and Service Management*, vol. 21, no. 2, pp. 1791–1801, 2024.

[9] D. Renga, M. Ni, M. Ajmone Marsan, and M. Meo, "Network sharing to enable sustainable communications in the era of 5G and beyond," in *Proceedings of the IEEE International Conference on Communications (ICC)*. IEEE, 2024, pp. 2840–2846.

[10] M. Ajmone Marsan, A. P. Couto da Silva, M. Meo, and D. Renga, "Sharing renewable energy in a network sharing context," in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2018, pp. 1–6.

[11] M. Ni, D. Renga, M. Meo, and M. Ajmone Marsan, "Network sharing for fault resilience," in *Proceedings of the 14th International Workshop on Resilient Networks Design and Modeling (RNDM)*. IEEE, 2024, pp. 1–7.

[12] 3GPP, "Study on scenarios and requirements for next generation access technologies," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.913, 2022, release 17.

[13] P. Popovski, K. F. Trillingsgaard, O. Simeone, and G. Durisi, "5G wireless network slicing for eMBB, URLLC, and mMTC: A communication-theoretic view," *IEEE Access*, vol. 6, pp. 55 765–55 779, 2018.

[14] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[15] T. L. Saaty, *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York: McGraw-Hill, 1980.

[16] J. Dupačov'a, N. Gr"owe-Kuska, and W. R"omisch, "Scenario reduction in stochastic programming," *Mathematical Programming*, vol. 95, pp. 493–511, 2003.

[17] L. S. Shapley, "A value for n-person games," in *Contributions to the Theory of Games, Volume II*, ser. Annals of Mathematics Studies, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ, USA: Princeton University Press, 1953, vol. 28, pp. 307–317.

[18] M. Oikonomakou, A. Antonopoulos, L. Alonso, and C. Verikoukis, "Evaluating cost allocation imposed by cooperative switching off in multi-operator shared HetNets," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 12, pp. 11 352–11 365, Dec. 2017.

[19] D. Fioriti, T. Ferrucci, and D. Poli, "Fairness and reward in energy communities: Game-theory versus simplified approaches," in *Proceedings of the 2023 IEEE International Conference on Environment and Electrical Engineering and 2023 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*. IEEE, 2023, pp. 253–258.