POLITECNICO DI TORINO

Master's Degree in

ICT for Smart Societies

Master's Thesis

6G Multi-HAPS Networks: Coverage, Resilience and Priority-Aware Strategies

Supervisor:

Prof. Michela Meo

Candidate:

Ten. ing. Marco Cutarelli

Co-supervisor:

Dr. Greta Vallero

Academic Year 2024/2025

Alla mia famiglia,

la mia fonte di amore e sostegno in ogni momento.

Ai miei genitori, che mi hanno insegnato il valore dell'onestà.

Abstract

When terrestrial communications partially or completely fail, rapidly deployable non-terrestrial overlays can be decisive to preserve essential services.

This thesis presents a modular simulation framework for multi-HAPS (High-Altitude Platform Station) networks targeting urban environments and large-scale outages. The framework integrates realistic MU-MIMO air-interface modeling, antenna directivity, and weather-aware propagation losses (free-space, atmospheric gases, rain, cloud/fog) within a priority-aware management layer. The architecture extends to HAPS-to-HAPS operation, enabling cooperative overlays when terrestrial infrastructure is degraded or unavailable.

Evaluation over Paris, using Orange mobile-traffic traces, compares user association, scheduling, and content placement under fair-weather and worst-day meteorology.

A Best-HAPS association balancing link quality, platform load, and backhaul headroom, combined with admission strategies inspired by weighted proportional-fair objectives, delivers robust per-area beam capacity (averaging $\approx$ 100 Mb/s).

Capacity remains higher in suburban/rural sectors and tighter in dense urban zones, reflecting path loss and load patterns.

Priority-aware allocation consistently favors critical services (e.g., hospitals, public safety), reducing blocking and tail latency. The simulator couples traffic, channel, and allocation policies, providing a reproducible testbed for multi-HAPS strategies.

Limitations include simplified small-scale dynamics, coarse granularity in mobility and demand, and limited inter-HAPS coordination. Future work will incorporate richer channel and energy models, dynamic role switching, multi-hop relaying, and learning-based control for adaptive prioritization under uncertainty. Overall, the results support

multi-HAPS overlays as a practical path to resilient communications for dual-use civil and military operations.

Contents

List of Figures

List of Tables

List of Acronyms

3GPP	3rd Generation Partnership Project
AAW	Anti-Air Warfare
ADS-B	Automatic Dependent Surveillance–Broadcast
AIS	Automatic Identification System
AIS/ADS-B	Automatic Identification System / Automatic Dependent Surveillance–Broadcast
BS	Base Station
CIRA	Centro Italiano Ricerche Aerospaziali
COMINT	Communications Intelligence
CRS	Coordinate Reference System
CSV	Comma-Separated Values
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
ECC	Electronic Communications Committee
ELINT	Electronic Intelligence
ENISA	European Union Agency for Cybersecurity
EO/IR	Electro-Optical / Infrared
EPW	EnergyPlus Weather
ESA	European Space Agency
EWMA	Exponentially Weighted Moving Average
FR1	Frequency Range 1 (sub-6 GHz)

FWA Fixed Wireless Access

GBAD Ground-Based Air Defence

GEO Geostationary Earth Orbit

gNB Next-Generation NodeB

HAPS High-Altitude Platform Station

HHAA Hybrid High-Altitude Airship

IAB Integrated Access and Backhaul

IMT International Mobile Telecommunications

IoT Internet of Things

ISR Intelligence, Surveillance, Reconnaissance

ITU International Telecommunication Union

ITU-R ITU Radiocommunication Sector

KPI Key Performance Indicator

LAN Local Area Network

LEO Low Earth Orbit

LTE Long Term Evolution

MEC Multi-access Edge Computing

MIMO Multiple-Input Multiple-Output

MTOW Maximum Take-Off Weight

MNO Mobile Network Operator

MU-MIMO Multi-User Multiple-Input Multiple-Output

NLOS Non-Line-of-Sight

NR New Radio

NTN Non-Terrestrial Network

OSINT Open-Source Intelligence

PF Proportional-Fair

PSAP	Public Safety Answering Point
QoS	Quality of Service
RAN	Radio Access Network
RF	Radio Frequency
SAM	Surface-to-Air Missile
SCADA	Supervisory Control and Data Acquisition
SE	Spectral Efficiency
SINR	Signal-to-Interference-plus-Noise Ratio
SLA	Service-Level Agreement
SNR	Signal-to-Noise Ratio
THz	Terahertz
TMY	Typical Meteorological Year
UAV	Uncrewed Aerial Vehicle
UE	User Equipment
UGV	Uncrewed Ground Vehicle
WGS84	World Geodetic System 1984
WRC	World Radiocommunication Conference

Chapter 1

Introduction

Modern societies are increasingly structured around digital infrastructures and advanced communication technologies. While this enables unprecedented efficiency and connectivity, it also exposes vulnerabilities in terms of resilience and reliability.

Dependence on Internet connectivity and network services makes societies critically exposed to disruptions in the underlying infrastructures.

On April 28, 2025, the Iberian Peninsula experienced a large-scale blackout that left millions of citizens in Spain and Portugal without electricity for several hours.

Public transport halted, airports and urban mobility systems were disrupted, and routine hospital operations were suspended; telecom and internet access were also affected.[2] [3]

This event highlighted the intrinsic fragility of terrestrial power and communication infrastructures and their limited ability to withstand widespread failures. In such scenarios, where terrestrial communication infrastructures partially or completely collapse, the rapid deployment of non-terrestrial solutions may play a decisive role in ensuring continuity of essential services.[4]

This chapter presents a modular simulation framework for multi-HAPS networks intended for complex urban scenarios. The architecture combines realistic MU-MIMO radio models, antenna directivity, and channel effects (including weather). The design is extensible and lean, enabling cooperative operation under variable load and partial unavailability of the terrestrial network.

1.1 System Architecture

The framework is composed of the following modules:

- **Communication module:** Implements the multi-user MIMO air interface, integrating small-scale fading, configurable antenna patterns, path loss, beam directivity, and weather-dependent propagation losses (e.g., rain, fog, snow).

- **Management module:** Dynamically decides whether a user request is served locally or by another HAPS and applies the selected allocation and prioritization strategies.

1.2 Backhaul and Relay Concepts

The backhaul, representing the HAPS–core network link, is modeled as a limited and costly resource.

Each HAPS is endowed with a finite backhaul budget. In the current implementation, the Best-HAPS selection is primarily driven by radio capacity, but the backhaul headroom can be summarized by a factor $\eta_{\mathrm{BH}} \in (0, 1]$ and integrated in the score to penalize attachment to platforms with saturated backhaul in future extensions.

Relay functionality (left for future work) would allow HAPS to forward data to one another over inter-platform links, so as to substitute missing backhaul or extend coverage. In this perspective, HAPS-to-HAPS relays could provide three main benefits:

- **Backhaul substitution:** a HAPS without direct core connectivity could connect through another HAPS.

- **Coverage extension:** a relay HAPS could serve users beyond the footprint of any single station.

- **Latency and miss reduction:** relays could reduce the need for cloud requests by reusing cached content within the aerial layer.

1.3 Thesis Objectives

The simulation environment is centered on Paris and the study area is partitioned into 108 zones, covering an area of approximately 107 km^2. The main objectives are:

- to develop a Python-based simulation tool for MU-MIMO communication between HAPS and ground users;
- to model the radio channel and coverage of HAPS in dense-urban, urban, and suburban environments, accounting for small-scale fading, path loss, building penetration, weather, and antenna directivity;
- to obtain a complete, simulation-ready dataset starting from raw data;
- to integrate priority metrics into the selection of the "Best HAPS" for each zone, ensuring that sensitive infrastructures (hospitals, emergency services, control centers) are prioritized;
- to implement traffic-aware mechanisms for deciding whether requests are served locally, via another HAPS, or through the backhaul;
- to simulate inter-HAPS cooperation strategies;
- to evaluate achievable capacity and the percentage of critical sites served, comparing selection strategies (nearest-based, capacity/SNR-based, cooperative priority-aware).

This thesis uses realistic simulations to assess how HAPS-based non-terrestrial networks can improve the resilience of communication infrastructures during large-scale terrestrial outages.

Chapter 2

State of the Art

This chapter summarises the main concepts and works on High Altitude Platform Stations (HAPS), their integration within Non-Terrestrial Networks (NTN), and resource management strategies for service continuity.

The objective is to provide an overview of the current state of the art and to highlight the research gap addressed in this thesis: the adoption of cooperative and priority-aware HAPS strategies to ensure resilience of critical infrastructures during large-scale blackout scenarios.

2.1 High Altitude Platform Stations (HAPS)

HAPS are aeronautical platforms deployed in the stratosphere, typically at altitudes between 15 and 25 km, equipped with telecommunication payloads capable of delivering wide-area connectivity [5, 6]. Compared to Low Earth Orbit (LEO) and Geostationary Orbit (GEO) satellites, HAPS offer several advantages:

- **Low latency**: the shorter distance significantly reduces propagation delay;
- **Regional coverage**: beams can be dynamically tailored to urban or suburban scenarios;
- **Operational flexibility**: platforms are retrievable, maintainable, and reconfigurable;

- **Cost-effectiveness**: deployment and maintenance are potentially cheaper than satellite alternatives.

Within the 3GPP framework, HAPS are explicitly modeled as part of NTN. TR 38.811 and TR 38.821 define scenarios, requirements, and channel models for stratospheric systems, which form the basis for simulation environments such as the one developed in this thesis.

2.2 Architecture and Operation

A generic HAPS system integrates:

- a **communication payload** with multi-beam antenna arrays and beam-steering capabilities;
- **direct HAPS-to-ground links**;
- **backhaul connections** via satellite or terrestrial gateways;
- **inter-HAPS links**, enabling cooperation;
- a **sensing payload** (EO/IR, hyperspectral, RF sensing, AIS/ADS-B, weather) for mission-driven observation;
- **sense-and-avoid/surveillance sensors** (ADS-B In/Out, transponder, optional optical/lidar/radar) for safe integration into controlled airspace.

Beamforming allows HAPS to simultaneously serve multiple clusters with controlled interference. Multi-user MIMO techniques, largely studied for terrestrial massive MIMO systems [7], are applicable to HAPS as well.

In this thesis, MU-MIMO performance is modeled by a small-scale fading module and a configurable MU-MIMO configuration block, implemented in Python as the classes `SmallScaleFading` and `MIMOConfig`. These components generate the MIMO channel matrices and set the antenna and link-budget parameters used in the capacity evaluation.

2.3 Real-World HAPS Projects

The feasibility of HAPS-based communication systems has been tested in several projects:

- **Google Loon** (2013–2021): deployed stratospheric balloons providing LTE connectivity in underserved regions. Despite demonstrating coverage feasibility, the project was discontinued due to economic and regulatory challenges.
- **Airbus Zephyr**: a solar-powered stratospheric UAV capable of flights over 60 days. Zephyr remains active and targets both civil and defence applications, representing one of the most mature HAPS platforms. However, there are limitations related to longitudinal coverage and the efficiency of solar PVs.
- **EuroHAPS/CIRA**: within the EU-funded *EuroHAPS* programme, CIRA leads the development of a tactical Hybrid High Altitude Airship (HHAA). The platform combines aerostatic buoyancy and aerodynamic lift in a lenticular airship configuration, enabling multi-month missions around 20 km altitude with medium payloads. The Italian demonstrator is designed for intelligence, surveillance and reconnaissance (ISR) and broadband communications, with flight tests planned over the Mediterranean/Canary Islands, and provides a relevant reference for hybrid airship-based HAPS concepts.[1]

Other initiatives include HAPSMobile (SoftBank) and Thales *Stratobus*, focusing on 5G, IoT, and hybrid civil–military services. These examples confirm the strong industrial and institutional interest in HAPS as a complement to terrestrial networks.

2.3.1 Platform assumptions (CIRA HHAA)

Stratospheric operations at $\sim$ 18–20 km are assumed, sustained by a hybrid solar–battery architecture that powers the payload 24/7, with battery sizing for long night durations ($\sim$ 15 h). The platform can hold position against headwinds in the $\sim$ 7–25 m/s range and targets multi-month endurance (on the order of $\sim$ 4 months). These assumptions bound feasible link budgets and duty cycles in the simulator and are consistent with the CIRA HHAA mission profile [1].

Figure 2.1: Artist's rendering of a CIRA hybrid high-altitude airship with a nadir-looking payload (courtesy: CIRA) [1].

2.3.2 Mission profile and operational constraints (CIRA HHAA)

The CIRA high-altitude hybrid airship (HHAA) combines aerostatic buoyancy and aerodynamic lift along a three-phase mission profile [1]:

- **Phase 1: Buoyant segment (vertical launch and ascent).** The vehicle launches vertically in a *partially inflated* configuration and climbs under aerostatic buoyancy only. During the ascent, the helium lifting gas expands until reaching the superpressure needed for the envelope to achieve its designed shape. The envelope becomes fully inflated at an altitude of about 10–12 km, where the attitude transitions towards approximately horizontal. Up to this point, the HHAA ascent is similar to that of the SwRI HiSentinel stratospheric airship [1]. Preliminary CIRA material illustrates this phase with indicative climb rates on the order of 300 ft/min.

- **Phase 2: Aerodynamic powered ascent (airship mode).** Above roughly 12 km altitude, the hybrid airship enters an aerodynamic ascent phase. Electric motor-driven propellers provide the thrust needed to climb from the transition layer up to the operating altitude of 18–20 km. Exact segment durations are still under study and should be regarded as indicative.

- **Phase 3: On-station mission at 18–20 km.** At cruise altitude, the HHAA operates

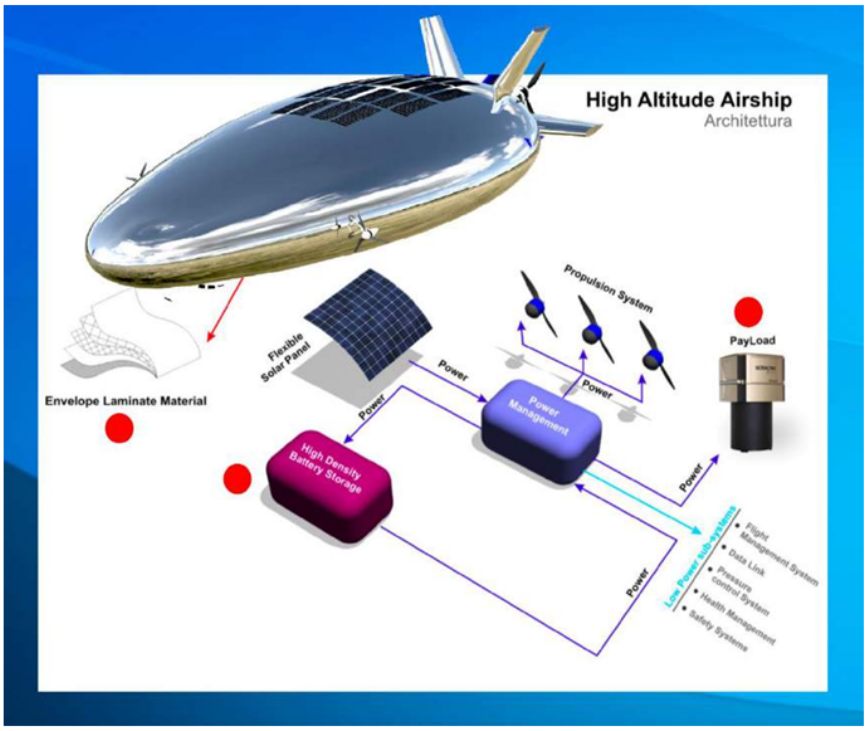

Figure 2.2: Power-system architecture of the CIRA hybrid high-altitude airship (HHAA), featuring a hull-mounted flexible solar array with high-density battery storage, power-management unit, electric propulsion system, and payload interface (courtesy: CIRA, [1]).

on-station using an all-electric hybrid power system (hull-mounted solar array plus batteries), designed for multi-month endurance and night-time autonomy. Recovery is achieved through a controlled descent and a short aerodynamic landing, enabling reuse of the platform [1].

Figure 2.2 summarizes the HHAA power-system architecture, highlighting the solar–battery hybrid design and the all-electric propulsion and payload supply.

2.3.3 Station-keeping and hovering strategies

Several studies suggest that keeping a high-altitude platform (HAPS) quasi-stationary over its service area is often preferable to wide-area patrol patterns. The main reasons are related to energy efficiency, night-time autonomy, and the limited performance gains obtained by moving the platform laterally.

First, holding position in the calm stratospheric air requires relatively low propulsion

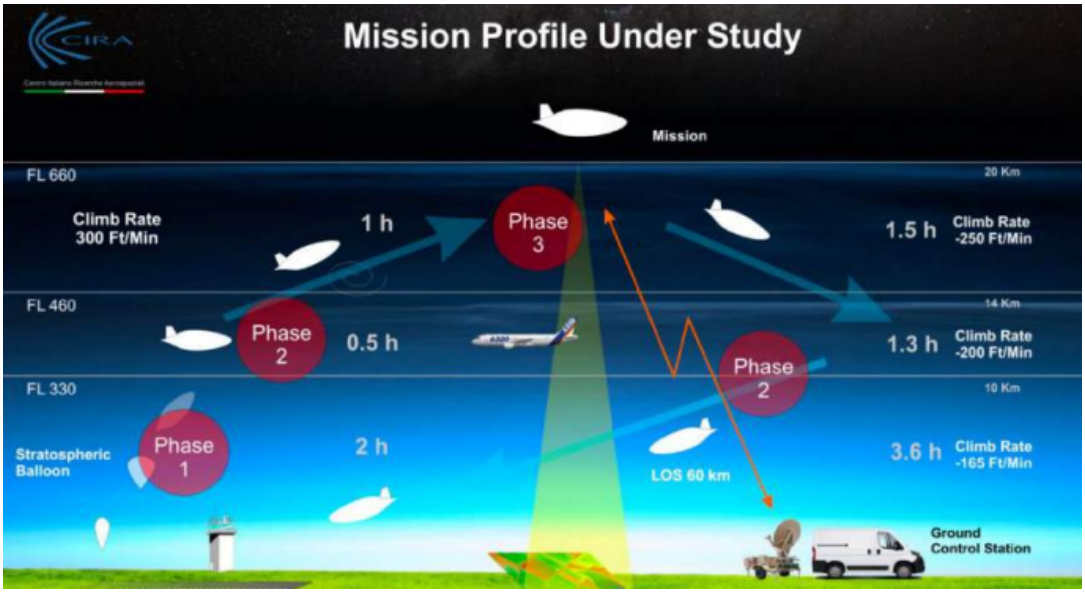

Figure 2.3: Mission profile under study: vertical launch, buoyant ascent, powered ascent, on-station at 18–20 km, and recovery (courtesy: CIRA) [1].

power, since winds at 18–20 km altitude are usually light and steady [8]. Manoeuvring or flying long horizontal legs can almost double the propulsion budget, leaving less energy available for the communications payload. Some analyses show that minimising horizontal motion and exploiting gliding phases at night can reduce daily energy consumption from about 13.5 kWh to 6.4 kWh [8].

Second, during night-time the platform is entirely powered by its batteries. Reducing unnecessary motion helps stretch the stored energy over the 10–12 hours of darkness, which is essential to avoid service interruptions. Estimates indicate that a typical night requires around 35 kWh to keep both airframe and payload operational [9].

Third, HAPS operate above most weather systems, where turbulence and convective phenomena are limited [10]. Remaining within this smooth layer simplifies flight control and reduces fast SINR fluctuations on the radio link [11]. Large horizontal excursions may expose the platform to stronger winds or less favourable regions.

Finally, from an altitude of roughly 20 km a single platform can illuminate a radius of 50–100 km. Small lateral shifts therefore change the geometry only marginally. With directional antennas featuring beamwidths of a few degrees, slow drift can be compensated by electronic beam steering, and the benefits of deliberate roaming are often limited

Figure 2.4: Vertical launch configuration of the HHAA (partially inflated) (courtesy: CIRA) [1].

compared with the added complexity in handover and pointing control.

For these reasons, static or slowly drifting hovering configurations are commonly adopted as a baseline in the HAPS literature, and the same assumption is used in the simulation framework of this thesis.

2.3.4 Size, mass, and payload class

CIRA documentation and patent disclosures indicate a medium-size, reusable airship class designed for payloads of about 25–100 kg, with MTOW on the order of 25–450 kg, and scalable geometric envelopes (representative ranges: length ∼10–40 m; width ∼8–35 m), depending on the embodiment and mission configuration [1]. These values guide

Figure 2.5: Rendering of a CIRA tactical hybrid stratospheric airship (courtesy: CIRA) [1].

feasible payload power budgets, station-keeping margins, and integration constraints in the simulator.

2.3.5 Representative mission roles

Under recent European demonstration efforts, the HHAA class targets ISR and communications roles, including (as representative payload families) *lidar* payloads for maritime/land target detection and classification, *COMINT/ELINT* packages, and *meshed broadband networking* nodes. These roles translate into distinct payload duty cycles and link requirements, which are later mapped to traffic categories and priority classes in Subsec. 3.2.3.

2.4 Dynamic Bandwidth Allocation and Quality of Service (QoS)

Efficient resource management is central to HAPS operation. Common scheduling strategies include the ones summarised in Table 2.1.

Strategy	Objective	Strengths	Caveats
Proportional Fair (PF)	Balance throughput and fairness	High spectral efficiency; simple to implement	May starve critical low-SINR users
Max–Min Fairness	Maximize worst-user rate	Strong minimum-rate guarantees	Lower overall throughput; sensitive to outliers
Weighted PF	Favor priority tiers	Tunable QoS; supports emergency services	Needs careful weight design and policing

Table 2.1: Scheduling strategies for HAPS-based overlays: objectives, strengths, and caveats.

In this work, these objectives are introduced as standard formulations from the HAPS/NTN literature. The simulator does not implement a full MAC-level proportional-fair or max–min scheduler; instead, it focuses on the Best-HAPS association and on the admission policies in Sec. 3.4.2, which decide which services are actually admitted on each beam under blackout conditions. The corresponding formulations are reported below.

Max–min fairness

$$\max_{\mathbf{x}} \min_{u} R_u \quad \text{s.t. resource and power constraints.} \tag{2.1}$$

Proportional fairness (PF)

$$\max_{\mathbf{x}} \sum_{u} \log\big(\bar{R}_u\big), \qquad \bar{R}_u \leftarrow (1-\alpha)\,\bar{R}_u + \alpha\, R_u, \tag{2.2}$$

where $\bar{R}_u$ is the smoothed past throughput (EWMA, $0 < \alpha \ll 1$).

Priority-weighted utility

$$\max_{\mathbf{x}} \sum_{u} w_u \log\big(\bar{R}_u\big) \quad \text{or} \quad \max_{\mathbf{x}} \sum_{u} w_u\, R_u, \tag{2.3}$$

with $w_u > 1$ for high-priority services/areas and $w_u = 1$ otherwise.

In emergency scenarios such as blackout events, it is crucial to adopt priority-aware scheduling, ensuring that critical infrastructures (e.g., hospitals, SCADA systems, emergency response centres) receive guaranteed QoS. While priority mechanisms are well

established in terrestrial networks, their adaptation to HAPS-based NTN architectures for disaster resilience remains underexplored.

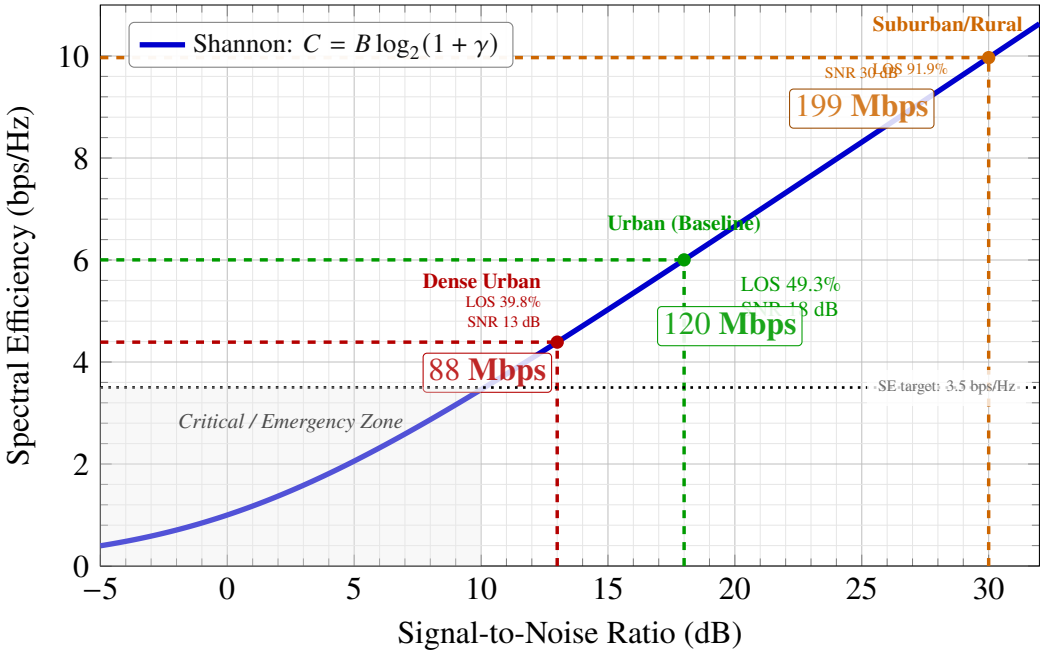

Figure 2.6: Shannon capacity (20 MHz @ 2 GHz) at 30° elevation for three representative urbanicity classes. The dotted line marks an operational target of SE = 3.5 *bps/Hz and the shaded region highlights the critical low-SNR regime.*

Figure 2.6 illustrates the Shannon capacity for three representative operating points corresponding to dense urban, urban, and suburban/rural conditions at 30° elevation, over a 20 MHz channel at 2 GHz. Using 3GPP TR 38.811 parameters (LOS probability and clutter losses encoded in `haps_parameters.py`), the Dense Urban case (red, SNR = 13 dB, LOS = 39.8%) yields a spectral efficiency of $\eta \approx 4.39$ bps/Hz, i.e., about 88 Mbps. The Urban baseline (green, SNR = 18 dB, LOS = 49.3%) reaches $\eta \approx 6.00$ bps/Hz (120 Mbps), while the Suburban/Rural point (orange, SNR = 30 dB, LOS = 91.9%) attains $\eta \approx 9.97$ bps/Hz, i.e., nearly 199 Mbps. The shaded area on the left highlights the critical regime where SNR < 10 dB, and the dotted horizontal line marks an illustrative operational target of SE = 3.5 bps/Hz.

The Shannon spectral efficiency (in bps/Hz) for an AWGN channel is given by:

$$\eta_{\text{bps/Hz}} = \log_2(1+\gamma), \tag{2.4}$$

where γ is the *linear* signal-to-noise ratio (SNR), defined as the ratio between received

signal power P_{s} and noise power P_{n}:

$$\gamma \;=\; \frac{P_{\mathrm{s}}}{P_{\mathrm{n}}}. \tag{2.5}$$

In practice, SNR is often expressed in dB:

$$\mathrm{SNR}_{\mathrm{dB}} \;=\; 10\log_{10}(\gamma), \qquad \gamma \;=\; 10^{\mathrm{SNR}_{\mathrm{dB}}/10}. \tag{2.6}$$

The corresponding link capacity C over a bandwidth B (in Hz) is:

$$C \;=\; \eta_{\mathrm{bps/Hz}} \times B_{\mathrm{Hz}} \;=\; B_{\mathrm{Hz}}\,\log_2\!\big(1+\gamma\big). \tag{2.7}$$

In the numerical examples, bandwidth is expressed in MHz and capacity in Mbps:

$$C_{\mathrm{Mbps}} \;=\; \eta_{\mathrm{bps/Hz}} \times B_{\mathrm{MHz}}. \tag{2.8}$$

For example, the Urban scenario at SNR = 18 dB yields:

- Spectral efficiency: $\eta = 6.0$ bps/Hz;
- Bandwidth: $B = 20$ MHz;
- Total capacity: $C = 6.0 \times 20 = 120\ Mbps$.

This capacity is sufficient to serve 12 concurrent telemedicine sessions (10 Mbps each) or 60 emergency two-way voice/video sessions (2 Mbps each) from a single HAPS beam.

Most contributions focus on generic throughput and coverage rather than on network resilience under extreme conditions. In particular, limited attention has been devoted to scenarios in which terrestrial infrastructures are severely impaired—such as large-scale blackouts—where inter-HAPS cooperation and service prioritization could provide decisive advantages. This research gap motivates the simulation framework developed in this thesis, which integrates realistic propagation models (including path loss, directivity gain, and fading), real geographic traces, and priority-based selection mechanisms.

2.5 Multi-HAPS Simulations and Models in the Literature

Several strands of work study cooperation among multiple HAPS:

- **Nearest/strongest-cell association (baseline).** The simplest rule assigns each user to the geographically nearest or strongest HAPS. It is widely used as a baseline but becomes suboptimal under interference and congestion, as shown by stochastic-geometry analyses and modern cell-association studies for 5G/6G [12] [13]. ,

- **Capacity-based association.** Users select the HAPS that maximizes an instantaneous (Shannon-like) rate or spectral-efficiency metric—often extended to account for load and limited backhaul. This is standard in massive-MIMO-style modeling and has been adapted in NTN contexts [7, 14, 15].

- **Cooperative relaying across HAPS.** Inter-HAPS links (FSO/RF) enable multi-hop relaying to extend coverage, bypass obstructed areas, or improve secrecy/robustness (e.g., selective DF relays, mixed FSO/RF chains, satellite–HAPS–ground paths) [16, 17, 18].

- **Priority-aware scheduling/offloading.** Weights w_i (service/area priority) are combined with capacity terms C_i (e.g., $s_i = w_i C_i$) to favor critical infrastructures or low-latency tasks in HAP/NTN edge-cloud systems [19, 20].

These directions align with the deployment scenarios and modeling assumptions collected by 3GPP for NR over Non-Terrestrial Networks (including HAPS and inter-HAPS links) [6, 14].

2.5.1 Regulatory landscape and spectral masks for HAPS (and HIBS/IMT)

Recent ITU World Radiocommunication Conferences have clarified the spectrum availability for high-altitude platforms. WRC-19 globally identified HAPS operation in the fixed service within *31–31.3 GHz* and *38–39.5 GHz*, and confirmed *47.2–47.5 GHz* and *47.9–48.2 GHz* for worldwide use under specified technical conditions. WRC-23 further established regulations enabling high-altitude platform stations as IMT base stations (HIBS) in the *2 GHz* and *2.6 GHz* mobile bands in certain regions, supporting direct-to-device use cases where licensed by administrations.

In this thesis, the HAPS access link is modeled in the S-band around $f_c \approx 2\,\text{GHz}$ with a bandwidth of $B = 20\,\text{MHz}$, consistent with FR1 NTN studies in 3GPP (where S-band is one of the main candidate bands for NR over NTN) and with the WRC-23 identification of 2 GHz and 2.6 GHz IMT bands for HAPS as IMT base stations (HIBS) [6, 14]. The focus is on system-level performance and resilience rather than on detailed coexistence analysis; therefore, the adopted carrier frequency and bandwidth are chosen to be representative of IMT-like deployments in S-band, while exact spectral-mask definitions and emission limits are left to future work.

At the same time, several studies and regulatory discussions indicate Ka-band as a more realistic option for high-capacity HAPS/NTN deployments. Translating the proposed architecture into an operational spectrum plan will therefore require revisiting the link budget and beam layout under Ka-band-specific regulatory and coordination constraints, which is left for future work.

2.6 Research gap

The existing literature on HAPS and NTN has established the feasibility of stratospheric platforms for broadband access, clarified their regulatory framework, and proposed a variety of scheduling and association strategies. Most contributions, however, focus on generic throughput and coverage objectives under nominal conditions. Scenarios in which terrestrial infrastructures are severely impaired — such as large-scale blackouts — have received comparatively less attention, especially when inter-HAPS cooperation and explicit service prioritisation are required to protect critical infrastructures.

This gap motivates the simulation framework developed in the following chapters, which combines realistic propagation models (including path loss, directivity gain, and fading), real traffic traces, and priority-aware resource-management policies to assess the resilience of multi-HAPS overlays in blackout scenarios.

Chapter 3

Materials and Methods

This chapter describes the traffic dataset and its preprocessing, the multi-HAPS system model, the allocation policies, and the simulation framework.

3.1 Traffic dataset: structure and dimensions

The analysis relies on urban mobile-traffic traces from the *Orange* network in the Paris region, with a time resolution of **15-minute slots**. After preprocessing, a three-dimensional NumPy array named `big_matrix_MB_bs1_108_service1_68_time1_672.npy` is obtained, storing per-zone, per-service, per-slot variables.

The study area is partitioned into 108 zones covering approximately 107 km^2 ($\approx 9 \times 12$ km) and spans dense-urban to suburban contexts.

For each zone, the dataset includes:

- the spatial information on service areas, with polygonal boundaries and centroid coordinates (latitude, longitude);
- an urbanicity label (dense-urban / urban / suburban; see Fig. 3.4), obtained via the POI-based clustering pipeline described in Sec. 3.2.1;
- a geographical priority flag (High/Low; see Fig. 3.6), derived from the presence of critical facilities such as hospitals, emergency centres, police stations, and major

transport hubs;

- a service-priority label for the 68 traffic services in the dataset (23 high-priority and 45 non-priority), depending on the type of digital service (e.g., emergency communications, cloud services, video streaming).

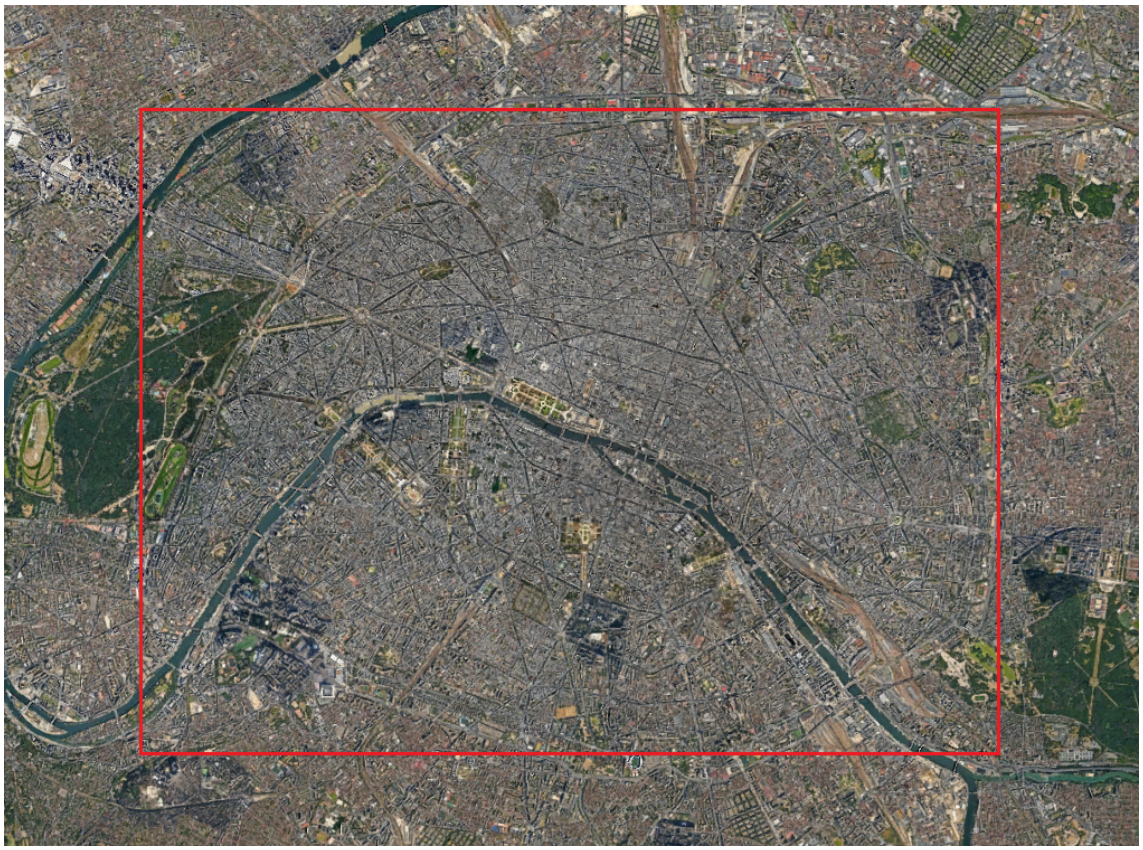

Figure 3.1: Area of interest over Paris (approximately 107 km^2*).*

3.2 Traffic Traces: Preprocessing and Construction

The traffic traces are obtained from two complementary *Orange* datasets that are first merged: (i) a per-service traffic dataset indexed by `station_id`, and (ii) a metadata dataset providing base-station coordinates (latitude, longitude). By joining these sources, each base station (BS) is enriched with geographic coordinates and its associated services.

A planning coverage radius is then assigned to each BS according to the antenna band/-type (field `emr_lb_system`). This radius is used to construct circular coverage footprints

serving as geometric inputs for mapping and subsequent coverage analyses. Starting from WGS84 coordinates (EPSG:4326), data are projected to a metric CRS suitable for Paris (Lambert–93, EPSG:2154); buffers are computed in meters; and geometries are finally reprojected to WGS84 for downstream mapping and spatial joins. The band-to-radius mapping is a pragmatic planning heuristic for spatial analysis and visualization, not a detailed RF propagation model. The nominal band (emr lb system) is mapped to a reference radius:

Frequency Band	Radius (km)
LTE 700	1.5
LTE 800	1.2
LTE 1800	0.8
LTE 2100	0.5
LTE 2600	0.3

Table 3.1: Reference radius by band (km) used for BS coverage buffers.

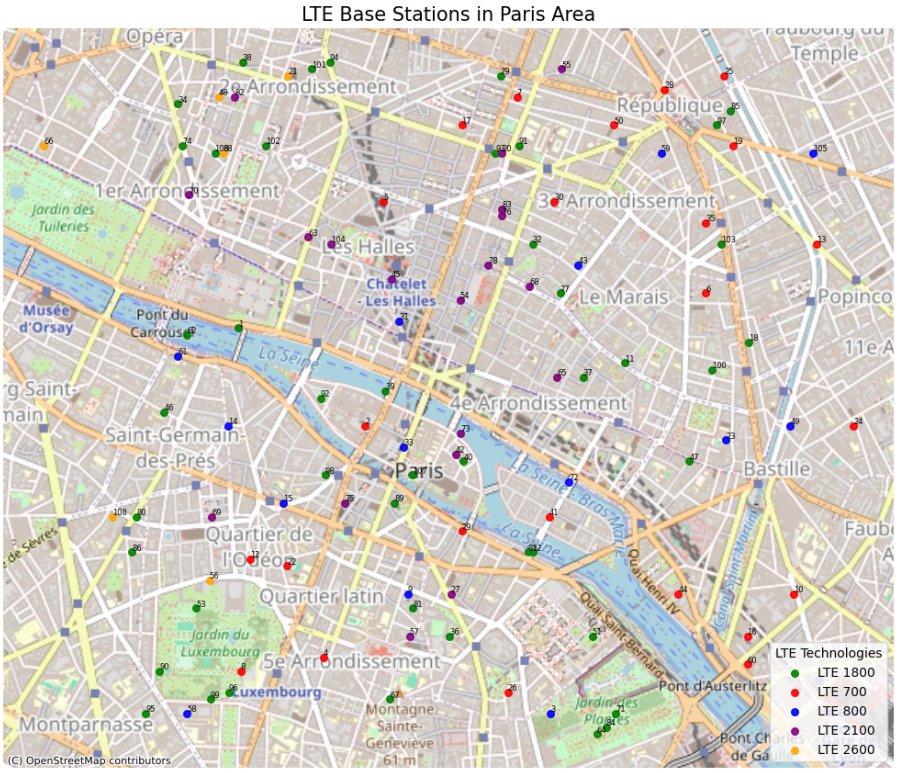

Figure 3.2: Location of LTE base stations by frequency band.

When the band is unknown or unmatched, a default radius of 1.0 km is applied. This results in a per-row field "radius - km" used to build the coverage polygons. The area is computed directly from the buffered geometry in the metric CRS and converted to km^2. The centroid (a point geometry) is also stored for fast spatial joins. A consistency check ensures that the final set includes 108 BS with complete metadata for simulation. Fig. 3.2 shows the location of LTE BS by frequency band, while Fig. 3.3 also shows the LTE BS coverage.

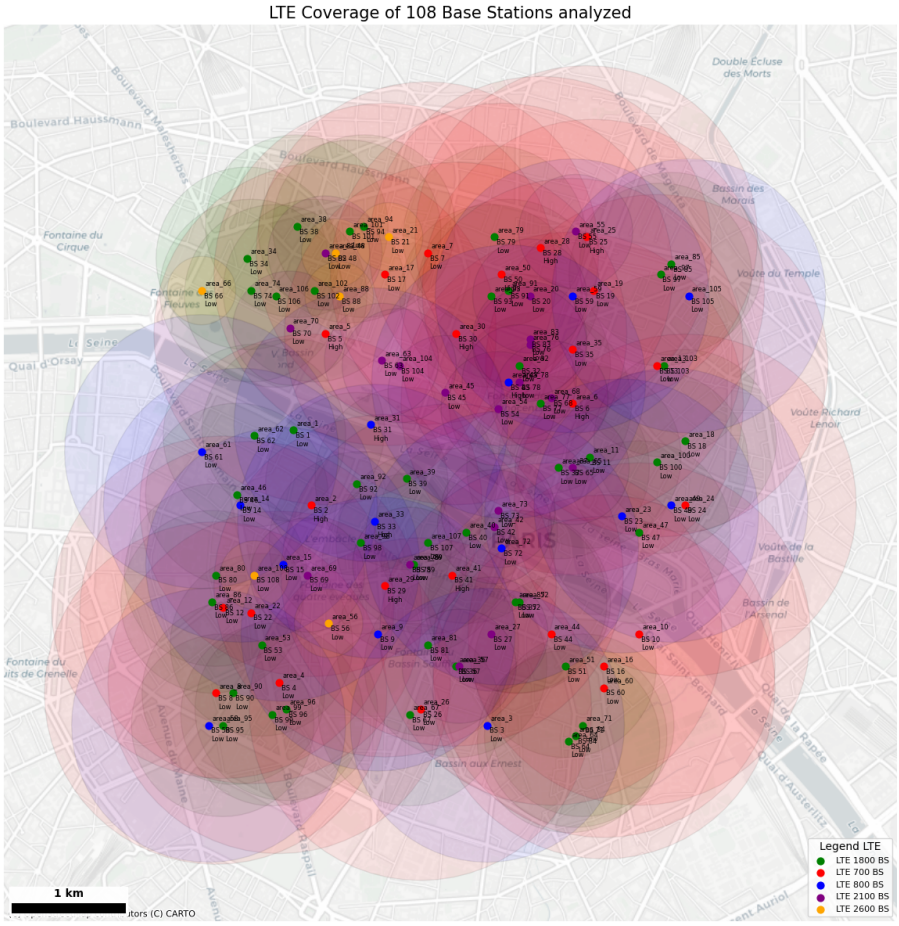

Figure 3.3: Coverage and location of LTE base stations by frequency band.

3.2.1 Urbanicity Labelling (Dense Urban / Urban / Suburban) through Overpass Turbo

In this work, *Points of Interest* (POIs) from OpenStreetMap are queried via Overpass Turbo within the study area and used to construct a simple notion of urbanicity. Each

base-station polygon is assigned to one of three macro-classes, namely *dense_urban*, *urban*, or *suburban_rural*, depending on the local concentration of critical services and transport facilities. To this end, three families of tags are considered:

- **critical services** (e.g., `hospital`, `police`, `fire_station`)
- **health & emergency** (e.g., `lifeguard`, `assembly_point`)
- **air transport** (e.g., `helipad`, `airport`)

Queries return OSM nodes as point features via `out center`, which simplifies downstream processing.

```
(
  nwr["amenity"~"^(hospital|police|fire_station)$"
   ](48.80,2.25,48.92,2.42);
  nwr["emergency"~"^(lifeguard|assembly_point)$"
   ](48.80,2.25,48.92,2.42);
  nwr["aeroway"~"^(helipad|airport)$"](48.80,2.25,48.92,2.42);
);
```

Codice 3.1: Overpass Turbo query with fixed bbox.

All POIs are spatially joined to traffic circles. For circle i, a POI feature vector $\mathbf{p}_i \in \mathbb{R}^K$ is built, where K is the number of POI categories:

$$\mathbf{p}_i = \left[\text{count}_i^{(1)}, \ldots, \text{count}_i^{(K)} \right].$$

Then, a *service-weighted* POI intensity is computed:

$$\phi_i = \sum_{k=1}^{K} \alpha_k \, \text{count}_i^{(k)}, \quad \alpha_k \geq 0 \text{ (higher for critical categories).}$$

The density normalization is done to mitigate polygon size effects:

$$\delta_i = \frac{\phi_i}{\text{area}(i)}.$$

A $\log(1 + \delta_i)$ transform can be applied to handle heavy tails.

Algorithm 1 POI → urbanicity labelling pipeline (DBSCAN-based)

1: **Input:** polygons $\{\mathcal{A}_i\}$, Overpass tag set $\mathcal{T}$, DBSCAN parameters $(\varepsilon, \texttt{min\_samples})$
2: OSM POIs in the AOI are downloaded with tags $\mathcal{T}$
3: Spatial join assigns POIs to polygons $\mathcal{A}_i$
4: $\mathbf{p}_i$ (counts per category) is built; $\phi_i = \sum_k \alpha_k \, \text{count}_i^{(k)}$ is computed
5: Normalization: $\delta_i = \phi_i / \text{area}(i)$ (optionally with a $\log(1 + \delta_i)$ transform)
6: Run DBSCAN on $\{\delta_i\}$ (and associated BS positions) with $(\varepsilon, \texttt{min\_samples})$
7: Each polygon is assigned a cluster label $\texttt{cluster}_i \in \{-1, 0, 1, \dots\}$; label -1 denotes noise
8: **Output:** cluster labels, features CSV, plots

Regarding cluster statistics, for each discovered cluster c ($c \geq 0$) returned by DBSCAN, the following quantities are computed:

$$\texttt{num\_stations}(c) = \#\{\text{BS in } c\}, \qquad \overline{r}(c) = \text{mean}(\texttt{radius\_km} \text{ of the BSs in } c).$$

According to the area classification rules, each station j is assigned an area type, yielding three macro-categories based on its cluster:

$$\texttt{area\_type}_j = \begin{cases} \text{dense\_urban}, & \text{if } \texttt{cluster}_j \neq -1 \wedge \texttt{num\_stations}(c_j) \geq 10 \wedge \overline{r}(c_j) \leq 0.900 \text{ km}, \\ \text{urban}, & \text{if } \texttt{cluster}_j \neq -1 \wedge \texttt{num\_stations}(c_j) \geq 5, \\ \text{suburban\_rural}, & \text{if } \texttt{cluster}_j = -1 \text{ or } \texttt{num\_stations}(c_j) < 5. \end{cases}$$

Here c_j is the cluster label of station j, with $\texttt{cluster}_j = -1$ denoting noise.

The `area_type` column is inserted right after `area_km2`, and the file is saved to `BS_coord_radius_priority_filtered_with108bs_classified_with_services` with the `_classified.csv` suffix. Each run prints a short summary of the processed rows per file.

- **Dense urban**: high POI intensity, corresponding to critical downtown cores (shadow fading $\sigma_{\text{LOS}} \approx 3$–$3.5$ dB; clutter loss >30 dB);
- **Urban**: medium density, with mixed residential and commercial zones ($\sigma_{\text{LOS}} \approx 4$ dB, $\sigma_{\text{NLOS}} \approx 6$ dB);

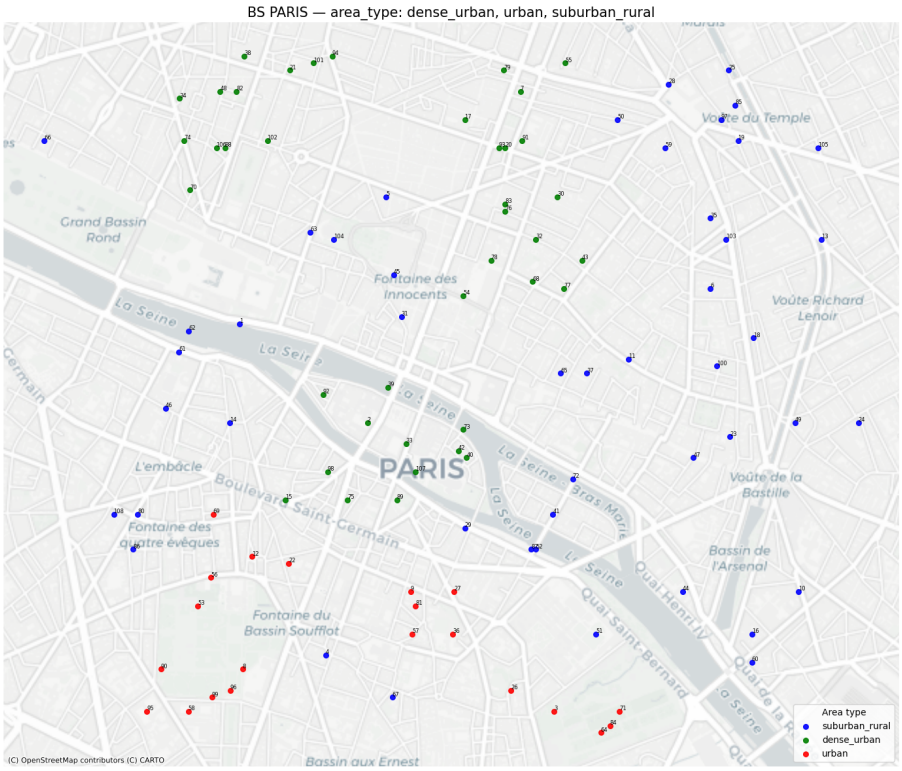

Figure 3.4: Dense-urban, urban, and suburban labelling from OSM POI density and DBSCAN clustering.

- **Suburban/rural**: low POI density, larger areas with weaker infrastructure ($\sigma_{\text{LOS}} \approx$ 0.7–1.8 dB; lower clutter loss, ≈ 16–19 dB).

These values follow the S-band propagation models from 3GPP TR 38.811, encoded in `haps_parameters.py` and summarised in Table 3.2, where LOS probabilities also increase with elevation angle (e.g., ~28% at 10° in dense-urban, >90% at 70°–80° in suburban).

Table 3.2: LOS probability and S-band path-loss parameters by elevation and urbanicity class.

Elev.	Class	LOS (%)	σ_{LOS} (dB)	σ_{NLOS} (dB)	CL_{NLOS} (dB)
10°	Dense-urban	28.2	3.5	15.5	34.3
10°	Urban	24.6	4.0	6.0	34.3

Continues on next page

Elev.	Class	LOS (%)	σ_{LOS} (dB)	σ_{NLOS} (dB)	CL_{NLOS} (dB)
10°	Suburban/rural	78.2	1.79	8.93	19.52
20°	Dense-urban	33.1	3.4	13.9	30.9
20°	Urban	38.6	4.0	6.0	30.9
20°	Suburban/rural	86.9	1.14	9.08	18.17
30°	Dense-urban	39.8	2.9	12.4	29.0
30°	Urban	49.3	4.0	6.0	29.0
30°	Suburban/rural	91.9	1.14	8.78	18.42
40°	Dense-urban	46.8	3.0	11.7	27.7
40°	Urban	61.3	4.0	6.0	27.7
40°	Suburban/rural	92.9	0.92	10.25	18.28
50°	Dense-urban	53.7	3.1	10.6	26.8
50°	Urban	72.6	4.0	6.0	26.8
50°	Suburban/rural	93.5	1.42	10.56	18.63
60°	Dense-urban	61.2	2.7	10.5	26.2
60°	Urban	80.5	4.0	6.0	26.2
60°	Suburban/rural	94.0	1.56	10.74	17.68
70°	Dense-urban	73.8	2.5	10.1	25.8
70°	Urban	91.9	4.0	6.0	25.8
70°	Suburban/rural	94.9	0.85	10.17	16.50
80°	Dense-urban	82.0	2.3	9.2	25.5
80°	Urban	96.8	4.0	6.0	25.5
80°	Suburban/rural	95.2	0.72	11.52	16.30
90°	Dense-urban	98.1	1.2	9.2	25.5
90°	Urban	99.2	4.0	6.0	25.5
90°	Suburban/rural	99.8	0.72	11.52	16.30

Note: LOS probabilities are expressed in percent. In the simulation code they are converted to fractions in [0,1].

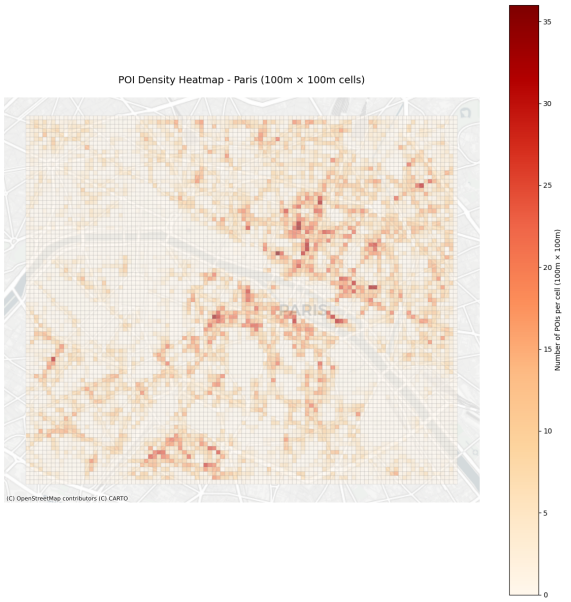

Figure 3.5: Heatmap of POI counts over a 100×100 *regular grid covering the study area. Darker cells correspond to higher POI density and highlight the main urban cores prior to aggregating POIs on BS polygons and running the clustering step.*

Figure 3.5 provides an initial visualization of POI density on a regular 100×100 grid. Each cell stores the number of POIs returned by the Overpass query within its footprint, revealing the main urban corridors and dense cores before any aggregation on BS polygons. This intermediate view is useful to verify that the query captures the expected spatial patterns over Paris and to tune the tag set and bounding box if needed.

3.2.2 Geographical priority (BS classification)

As mentioned above, the areas are categorised by criticality:

- **High priority**: areas hosting critical facilities such as hospitals and clinics, emergency call centres, police stations, SCADA sites, fire and rescue stations, and major transport hubs;
- **Low priority**: residential and commercial areas without specific strategic relevance.

Each class is assigned a weight w_i, used in the allocation metric $s_i = w_i \cdot C_i$.

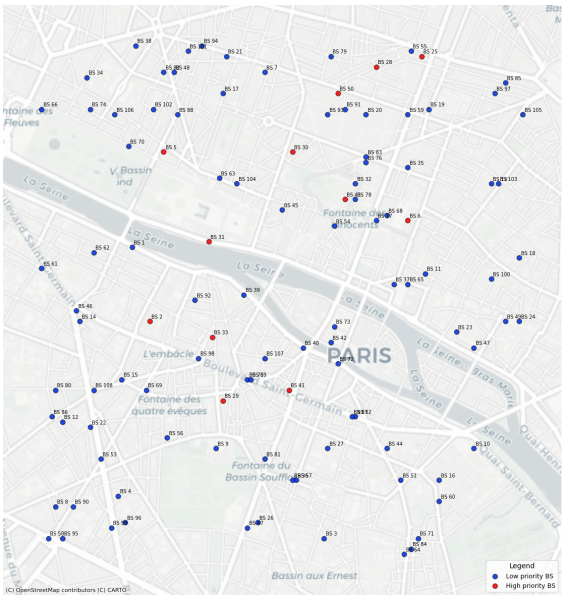

Figure 3.6: Geographical High/Low priority map for the 108 areas.

As introduced in Sec. 3.2.1, POIs from OSM provide a proxy for the spatial distribution of human activities across the Paris region. The study area is discretised into a regular grid of 100 m × 100 m cells, and for each cell the number of POIs falling within its boundaries is counted. The resulting density map, shown in Fig. 3.5, highlights several hot-spots corresponding to transport hubs, commercial areas, and touristic locations.

To obtain a binary classification of the 108 areas into *High* and *Low* priority, a simple scoring mechanism based on OpenStreetMap POIs is applied. Starting from the `PriorityArea.geojson` file, the following service types are extracted:

- hospitals and clinics;
- fire stations and rescue units;
- police stations and gendarmerie posts.

Each POI is mapped to a service type and contributes to an integer `priority_score` for the area whose polygon it falls into. The contribution of each service type is summarised in Table 3.3.

For each area, the `priority_score` is obtained by summing the contributions of all POIs intersecting its polygon; areas with no POIs receive a score equal to zero. The binary priority label is then assigned as follows:

Table 3.3: Score contribution of each service type to the area `priority_score`.

Service type	Score contribution
Hospital / clinic	3
Fire station / rescue	2
Police / gendarmerie	1

- all areas are initially marked as `Low`;
- among those with `priority_score` > 0, the 12 areas with the highest score are relabelled as `High`;
- all remaining areas keep the `Low` label.

The resulting `priority` column (`High`/`Low`) and the associated `priority_score` are used throughout the thesis to distinguish Tier-A (critical) from non-critical areas in the allocation strategies and in the evaluation of priority-aware KPIs.

3.2.3 Service Classification: Priority (23) vs. Non-Priority (45)

The service catalogue (N_S=68) is partitioned into *priority* (23 services) and *non-priority* (45 services) to guide capacity allocation under constrained scenarios. The mapping is version-locked in the simulator configuration to ensure reproducibility. Priority levels are defined based on the *service type*, with emergency, public-safety, and productivity applications given precedence over entertainment and best-effort traffic. The spatial distribution of Tier-A facilities (e.g., hospitals, public safety, command centres) is instead captured by the area-priority labels introduced in Subsec. 3.2.2, and combined with service priority in the admission and scheduling policies.

The resulting labels are used in the analysis pipeline in two main ways. First, they enable the aggregation of traffic per BS into priority and non-priority components (Sec. 3.2.4). Second, they provide a natural basis for defining weights w_i in the allocation score $s_i = w_i \cdot C_i$, which can be further modulated by urbanicity class or POI intensity. These

23 Priority services					
#	**Service**	**Motivation**	**#**	**Service**	**Motivation**
3	Apple_Mail	Medical communications	7	Apple_Web_Service	Backend for Apple clinical apps
8	Apple_iCloud	Patient records synchronization	14	Dropbox	Hosting of medical files
20	Google_Docs	Collaborative urgent documents	21	Google_Drive	Repository of clinical data
22	Google_Mail	Institutional hospital email	23	Google_Maps	Ambulance navigation and transfers
24	Google_Meet	Telemedicine and remote briefings	26	Google_Web_Services	Backend for hospital applications
28	LinkedIn	Staffing coordination and external expertise	29	Microsoft_Azure	Hosting of hospital software systems
30	Microsoft_Mail	Institutional communication	31	Microsoft_Office	Clinical records and documentation
32	Microsoft_Skydrive	Storage of medical files	34	Microsoft_Web_Services	Backend for management software
46	TeamViewer	Remote technical support	47	Telegram	Emergency communication channels
52	Waze	Ambulance routing and logistics	62	Web_Weather	Forecasts for medical transport
64	WhatsApp	Fast clinical group communication	65	Wikipedia	Quick medical reference
66	Yahoo_Mail	Residual personal communication			

Table 3.4: List of 23 Priority services in the Paris dataset with motivations.

weights are linked to radio-layer parameters such as directivity gain (Sec. 3.3.4, computed via `Directivity_Gain.py`), path-loss tables (from `haps_parameters.py`), and the small-scale fading models implemented in `small_scale_fading.py`.

3.2.4 Traffic Exploration per BS

Given a selected BS index $b \in \{1, \ldots, 108\}$, the slice $\mathbf{X}[b, :, :] \in \mathbb{R}^{68\times672}$ is extracted. Two diagnostic time series are produced: (i) the total traffic $T_{\text{tot}}(b,t) = \sum_{s=1}^{68} X(b,s,t)$, and (ii) the per-service trajectories $X(b,s,t)$, visualized with the human-readable labels from `services.csv`. The legend maps each curve to its corresponding service.

The catalogue is pre-labelled into *priority* (23 services) and *non-priority* (45 services) (see *Service Classification*). This mapping enables per-BS aggregation into the two complements:

$$T_{\text{prio}}(b,t) = \sum_{s\in\mathcal{S}_{\text{prio}}} X(b,s,t), \qquad T_{\text{non}}(b,t) = \sum_{s\in\mathcal{S}_{\text{non}}} X(b,s,t),$$

45 Non-priority services			
#	**Service**	**#**	**Service**
1	Paris_Amazon_Web_Services	25	Paris_Google_Play_Store
2	Paris_Apple_App_Store	27	Paris_Instagram
4	Paris_Apple_Music	33	Paris_Microsoft_Store
5	Paris_Apple_Siri	35	Paris_Molotov
6	Paris_Apple_Video	36	Paris_Netflix
9	Paris_Apple_iMessage	37	Paris_Orange_TV
10	Paris_Apple_iTunes	38	Paris_Periscope
11	Paris_Clash_of_Clans	39	Paris_Pinterest
12	Paris_DailyMotion	40	Paris_PlayStation
13	Paris_Deezer	41	Paris_Pokemon_GO
15	Paris_EA_Games	42	Paris_Skype
16	Paris_Facebook_Live	43	Paris_Snapchat
17	Paris_Facebook_Messenger	44	Paris_SoundCloud
18	Paris_Facebook	45	Paris_Spotify
19	Paris_Fortnite	48	Paris_Tor
49	Paris_Twitch	50	Paris_Twitter
51	Paris_Uber	53	Paris_Web_Ads
54	Paris_Web_Adult	55	Paris_Web_Clothes
56	Paris_Web_Downloads	57	Paris_Web_Finance
58	Paris_Web_Food	59	Paris_Web_Games
60	Paris_Web_Streaming	61	Paris_Web_Transportation
63	Paris_Web_e-Commerce	67	Paris_Yahoo
68	Paris_YouTube		

Table 3.5: List of 45 Non-Priority services in the Paris dataset.

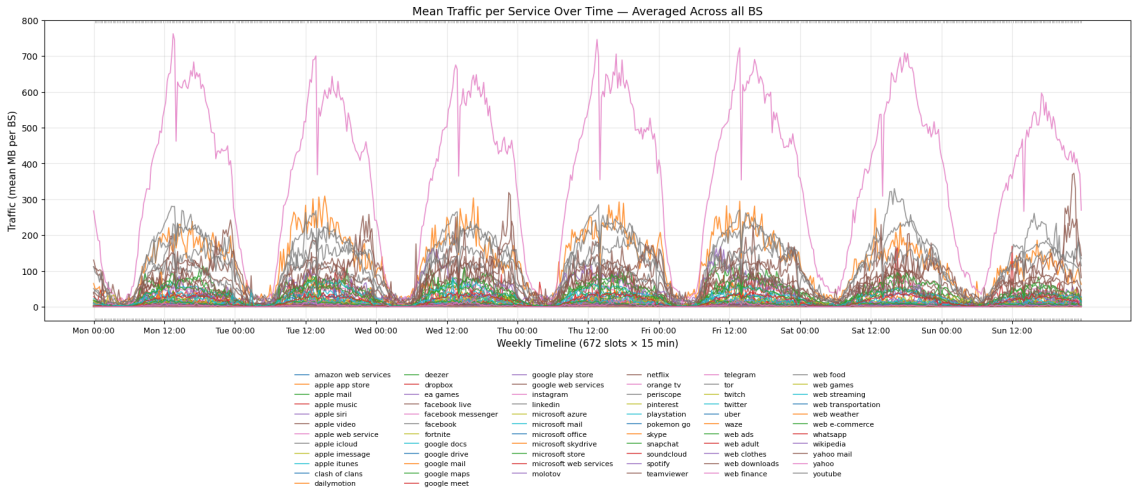

Figure 3.7: Mean traffic per service over time, averaged across all base stations. Each curve corresponds to one service; the y-axis reports the mean traffic (MB) per BS at 15-minute resolution over one week.

and the priority share

$$\rho(b,t) = \frac{T_{\text{prio}}(b,t)}{T_{\text{prio}}(b,t) + T_{\text{non}}(b,t)},$$

used downstream by the allocation metric $s_i = w_i \cdot C_i$ (Sec. 3).

The per-service chart with 68 curves can be visually dense; in practical use it is advisable to filter a subset (e.g., top-k contributors by weekly volume) or to overlay only the priority set. The service order in `services.csv` must match the second axis of X to avoid label mismatch. Given fixed dataset versions and a chosen BS index, results are deterministic.

3.2.5 Traffic concentration across base stations and services

To characterise the temporal dynamics of the dataset, the mean traffic over all base stations is computed for each service s at every 15-minute slot. The resulting time series, shown in Fig. 3.7, spans one full week and highlights both diurnal cycles and differences across services.

To better visualise how traffic concentrates across the network, the five base stations with the highest weekly load and the five most demanding services are selected. Figure 3.8 shows a heatmap of the total traffic (MB) generated by each BS–service pair.

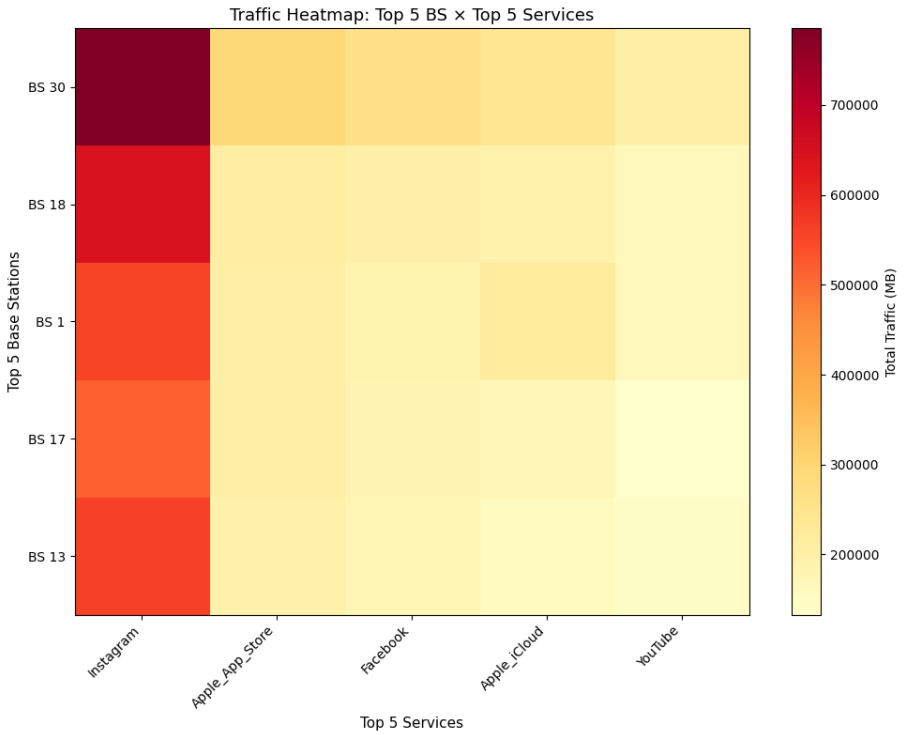

Figure 3.8: Traffic heatmap for the top five base stations and top five services. Colours encode the total weekly traffic per BS–service pair (MB).

The plot reveals that a single social-media platform (*Instagram*) dominates the load across all high-traffic sites, while other services such as *Apple App Store*, *Facebook*, *Apple iCloud*, and *YouTube* contribute smaller yet non-negligible volumes.

All services exhibit a clear daily pattern, with low traffic during night hours, a morning ramp-up, and evening peaks. A few applications (notably *Apple iCloud* and *Netflix*) dominate the overall load and show well-defined evening maxima, while most other services remain approximately one order of magnitude below. Weekend days display slightly different profiles, with flatter curves and peaks shifted towards later hours. Taken together, these spatial and temporal trends motivate the introduction of service-level priority classes and capacity-allocation policies that explicitly account for the heterogeneous and time-varying behaviour of different services.

3.3 Multi-HAPS System Model

The system comprises multiple HAPS operating at stratospheric altitude, each equipped with multi-beam antennas and sharing a common radio configuration. At a high level, the model is characterised by:

- number of HAPS N_{HAPS} and beams per platform N_{beams};
- transmit power P_{tx}, noise floor P_n, and carrier frequency f_c;
- aperture efficiency η and beam geometry (directivity gain vs. beam radius, altitude, user offset);
- HAPS altitude H within the CIRA HHAA operating range ($\approx$ 18–20 km).

3.3.1 Channel Modeling

On top of the geometric layout, the radio channel model combines large-scale attenuation, antenna directivity, small-scale fading and, when enabled, weather-induced losses. In particular:

- **Small-scale fading**: Rician or Rayleigh fading with configurable K-factor;
- **Path loss, shadowing and LOS probability**: large-scale terms derived from the S-band NTN profile in 3GPP TR 38.811;
- **Directivity gain**: antenna gain as a function of beam radius, altitude, and user offset, computed by the parabolic-antenna model;
- **Environmental attenuation** L_{env}: optional weather-induced loss (gases, rain, cloud/fog) applied along the slant path as detailed in Sec. 3.3.7.

These components jointly feed the SINR computation and capacity model in Sec. 3.4.

3.3.2 Operating Frequency Bands (S vs. Ka)

In line with 3GPP TR 38.811 (Study on NR to support NTN), S- and Ka-band are regarded as suitable options for HAPS/NTN operation. In this thesis, S-band is adopted as the baseline to remain consistent with a large portion of the HAPS literature, while Ka-band is acknowledged as a realistic alternative from a regulatory standpoint, enabling wider bandwidths at the cost of stronger atmospheric attenuation and narrower beams.

In qualitative terms:

- **S-band** offers wider coverage per beam for a given aperture and improved robustness to rain and cloud/fog attenuation, at the price of more limited spectrum availability.
- **Ka-band** provides larger feasible bandwidth and closer alignment with foreseen HAPS/NTN deployments, but entails higher free-space path loss and significantly more severe weather-related attenuation, which tighten EIRP and link-margin requirements.

The achievable rate scales approximately as

$$C = B \log_2\big(1 + \mathrm{SNR}\big),$$

so any migration from S- to Ka-band can increase capacity through a larger bandwidth B, provided that the resulting SNR remains acceptable after accounting for path loss and environmental attenuation. Within the simulator, the band choice affects:

1. the carrier frequency and, consequently, the free–space path loss term;
2. the feasible bandwidth B per beam;
3. the weather-loss component L_{env} (if enabled), which grows with frequency;
4. the beam design trade-off between coverage and spatial reuse.

A full quantitative assessment of an S→Ka migration falls outside the scope of this thesis and is left as future work, but the framework is band-agnostic and can be reused for Ka-band configurations by updating the link-budget parameters.

3.3.3 S-band Operation and Beamforming Capabilities

The S-band (2–4 GHz) represents a favourable compromise between link robustness and bandwidth availability for high-altitude platform systems. Compared to Ka-band, S-band links experience significantly lower atmospheric attenuation, enabling more reliable communication during adverse weather events. However, the lower carrier frequency implies a longer wavelength ($\lambda \approx 15$ cm at 2 GHz), which directly affects the achievable beamwidth and, in turn, the number of independent beams that can be formed.

At a HAPS altitude of approximately 21 km, the half-power beamwidth $\theta_{3\text{dB}}$ is related to the antenna aperture D by

$$\theta_{3\text{dB}} \approx \frac{70\lambda}{D}\ [^\circ].$$

For example, an aperture of $D = 1.0$ m yields $\theta_{3\text{dB}} \approx 10.5^\circ$, corresponding to a ground-footprint radius of roughly 1.8 km. Increasing the aperture to $D = 1.5$ m narrows the beamwidth to about 7°, reducing the footprint to approximately 1.2 km and increasing the directivity gain according to

$$G_m^0 = \eta \left(\frac{70\pi}{\theta_{3\text{dB}}}\right)^2,$$

where $\eta \approx 0.95$ is the aperture efficiency. This relation, implemented in the Directivity_Gain.py module, captures how smaller beamwidths translate into higher antenna gains at the expense of more RF chains and higher EIRP per beam.

In S-band operation, practical constraints on array size, power budget, and isolation typically limit the number of simultaneously active beams. A realistic configuration for a 4×4 MIMO payload operating at 2 GHz supports approximately 4–12 concurrent beams, with 8 beams representing a power-balanced operating point. Each beam covers a few square kilometres and can dynamically prioritise either critical or best-effort ground clusters.

The per-beam capacity is constrained by the 20 MHz bandwidth configured in MIMOconfig.py, the transmit power $P_{\text{tx}} = 41$ dBm, and the channel model implemented in small_scale_fading.py. Considering a path loss on the order of 135–140 dB, LOS probabilities from 3GPP TR 38.811, and typical SNRs between 15 and 25 dB, the achievable data rate per beam can be estimated via the Shannon equation,

$$C = B \log_2\big(1 + \text{SNR}\big),$$

yielding approximately 120–160 Mbit/s (single stream) and up to 220–280 Mbit/s under favourable LOS conditions with 2×2 MIMO.

Assuming eight beams active at the same time, the total system throughput of a single S-band HAPS platform lies between 1.0 and 1.3 Gbit/s in conservative conditions and can reach 2.0–2.2 Gbit/s in optimised scenarios. These values account for link-level overhead, beamforming losses, and moderate weather attenuation, and provide a realistic estimate of the S-band downlink capacity in the simulations.

3.3.4 Antenna modeling (parabolic approximation vs UPA/MIMO-style)

Unless stated otherwise, the antenna directivity pattern is modeled with a *parabolic-antenna approximation* that captures the main-lobe gain and an approximate roll-off. This choice is widespread and adequate for the scope of the analysis focused on coverage, link budget, and scheduling. A migration to a *uniform planar array (UPA)* / MIMO-style pattern is planned as a targeted improvement; the expected impact on the *main-lobe* gain is minor for the configurations considered, but side-lobe structure and spatial selectivity would be represented more faithfully in future revisions.

3.3.5 3D Beamforming Patterns for Representative Areas

To illustrate how the directivity model is instantiated for different ground locations, Fig. 3.9 shows two 3D beamforming patterns corresponding to two representative areas in the Paris dataset (`area_1` and `area_43`), both affected by the considered blackout scenario.

Each surface represents the antenna gain as a function of the angular displacement from boresight, computed through the parabolic approximation implemented in Directivity_Gain.py. The two areas are illuminated with slightly different boresight angles because they correspond to different off–nadir directions and slant ranges. Although the maximum

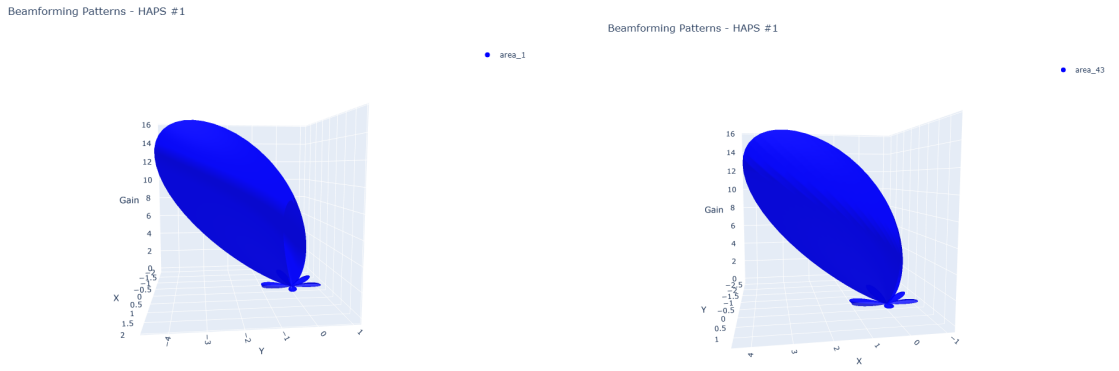

(a) Beamforming pattern for `area_1`. *(b) Beamforming pattern for* `area_43`.

Figure 3.9: 3D beamforming patterns of HAPS #1 towards two representative areas in the blackout region. Each surface shows the antenna gain as a function of angular displacement from boresight, according to the parabolic directivity model.

gain is identical (same antenna and configuration), the main lobes appear similar and only slightly rotated, reflecting the fact that the two blackout zones are geographically close.

This visualisation clarifies how the HAPS beam is steered towards different ground polygons in the footprint and provides an intuitive link between the abstract gain model and the concrete spatial geometry of the served areas. These beam patterns are then fed into the SINR computation and capacity model of (3.3).

3.3.6 Spectral Masks and Emission Regulations

In addition to propagation and fading, the radio model enforces spectral constraints to ensure coexistence with incumbent services. In-band limits (EIRP spectral density) and out-of-band (OOB) / spurious masks at specified frequency offsets (e.g., ±1 MHz, ±5 MHz, ±10 MHz) from the S-band carrier are considered, with stringent suppression levels beyond the nominal channel edges. Where applicable, ground *power flux density (PFD)* constraints are applied on a per-beam basis as a function of elevation angle to limit unwanted illumination of terrestrial areas. Such assumptions align with the regulatory framework summarised in Chapter 2 and with the use of radiation masks consistent with

Resolution 221 practices for sidelobe control [21].

Assumptions for the S-band case used in the simulator are summarised below:

- **EIRP spectral density (per beam)** limited to comply with national licensing and to avoid harmful interference to adjacent services in the 2 GHz range.
- **OOB/spurious mask**: aggressive roll-off at band edges; high rejection (e.g., −60 to −100 dB) outside the authorized bandwidth; harmonic suppression included in the PA/filter chain.
- **Antenna radiation mask**: phased-array pattern with strong sidelobe/backlobe attenuation (pattern constraints inspired by Resolution 221 practices) to reduce energy outside the useful sector [21].
- **PFD constraint**: per-beam check at ground for off-nadir angles, used as a feasibility guard in stress scenarios (e.g., high EIRP or narrow beams).

The mask is applied in the link budget as an additional attenuation for spectral leakage (OOB/spurious) and as a directional constraint via the radiation pattern. Capacity and scheduler outputs (Sec. 3.7) reflect these limits by capping per-beam spectral density and adjusting the beam layout to respect sidelobe constraints.

3.3.7 Weather–aware Path Loss Model

The end-to-end HAPS–ground path loss is modeled as the sum (in dB) of multiple components:

$$L_{\text{tot}}(t, \mathbf{x}) = L_{\text{FSPL}}(f, d) + L_{\text{shadow/clutter}}(\theta, \text{scenario}) + L_{\text{env}}(t, \theta, f) - G_{\text{tx}}(\theta) - G_{\text{rx}}. \tag{3.1}$$

Free–space path loss is computed as

$$L_{\text{FSPL}} = 32.45 + 20\log_{10}(f_{\text{MHz}}) + 20\log_{10}(d_{\text{km}}),$$

where d is the slant range and f the carrier (here ∼2 GHz).

Scenario-dependent shadowing and clutter penalties depend on elevation θ and environment (dense–urban, urban, suburban/rural), with LoS probabilities and NLoS penalties derived from the NTN profile [6]. Parameters are instantiated per scenario and applied area-by-area.

Although at 2 GHz average atmospheric attenuation is moderate, intense events (convective rain, dense fog, liquid water clouds) can introduce non-negligible degradation along the slant path—especially at low elevations—impacting SNR and capacity. For this reason a time-varying term $L_{\text{env}}(t, \theta, f)$ consistent with the link elevation is added to the model.

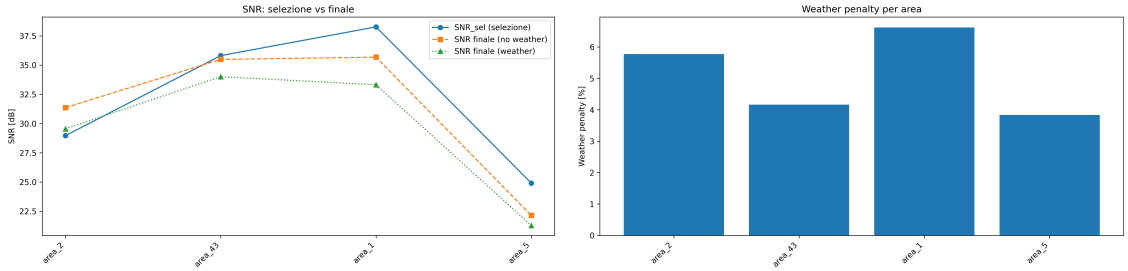

(a) SNR: selection vs final (weather on/off). *(b) Weather penalty by area (%).*

Figure 3.10: Impact of HAPS selection and weather conditions on link performance across representative areas.

The environmental term L_{env} combines (in dB) the main contributions:

$$L_{\text{env}}(t, \theta, f) \;=\; A_{\text{gas}}(t, \theta, f) \;+\; A_{\text{rain}}(t, \theta, f) \;+\; A_{\text{cloud/fog}}(t, \theta, f). \tag{3.2}$$

- **Gaseous absorption** (A_{gas}): function of pressure, temperature, and humidity (O_2/H_2O lines/continua). It is computed as specific attenuation integrated over the effective slant path according to ITU-R P.676 [22].

- **Rain** (A_{rain}): specific attenuation $\gamma_R = k(f)\, R^{\alpha}$ [dB/km], with R the rain rate (mm/h), per ITU-R P.838 [23]. The path attenuation is $A_{\text{rain}} \approx \gamma_R \cdot s_{\text{eff}}(\theta)$, where s_{eff} includes the slant projection ($\propto 1/\sin\theta$) and the rainy-layer thickness actually crossed.

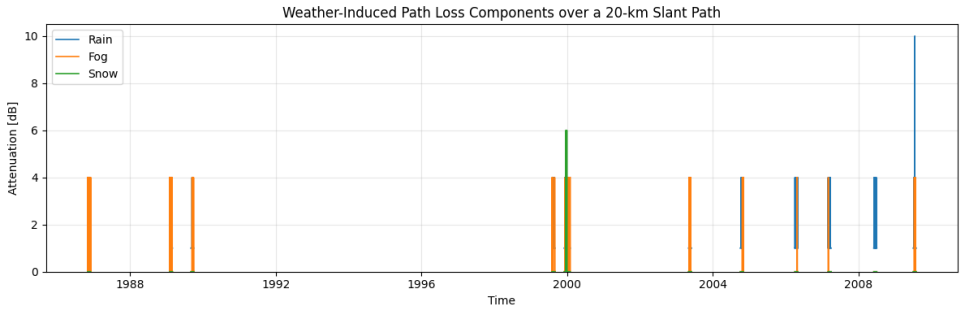

Figure 3.11: Rendering of a CIRA tactical hybrid stratospheric airship. Courtesy: CIRA [1].

Table 3.6: Best vs. worst day at the hourly peak of A_{env} (20 km slant path).

Case	Timestamp	A_{env} [dB]	Rain [dB]	Fog [dB]	Snow [dB]
Worst-day (peak)	1999-12-27 22:00	10.00	4.00	0.00	6.00
Best-day (peak)	1986-11-02 00:00	1.00	1.00	0.00	0.00

- **Cloud/fog** ($A_{\text{cloud/fog}}$): derived from specific attenuation proportional to the liquid water content (LWC) and optical thickness, integrated on the cloud/fog segment per ITU-R P.840 [24].

To validate the implementation, an empirical check was carried out on the EPW (TMY) weather file adopted in the simulator. The best and worst days were extracted according to the daily *peak* of A_{env} on a 20 km slant path. The worst-day peak occurs on 1999-12-27 22:00, with a total weather loss of 10.00 dB, composed of 4.00 dB rain, 0.00 dB fog, and 6.00 dB snow. Conversely, the best-day peak (least severe) occurs on 1986-11-02 00:00, with 1.00 dB total, entirely due to rain (1.00 dB), while fog and snow are 0.00 dB. These values are consistent with the models in [22, 23, 24] and confirm that, at S-band, severe events are episodic and elevation-dependent.

For each time t and area/user, the simulator follows a simple operational pipeline:

1. compute d and θ from HAPS–user geometry;
2. evaluate $L_{\text{FSPL}}(f, d)$ and apply shadowing/clutter by scenario and elevation as in [6];
3. build $L_{\text{env}}(t, \theta, f)$ from hourly weather fields mapped to the nearest cell (rain $R(t)$, LWC/visibility, thermo–hygrometric state) using [23, 24, 22];

4. subtract antenna gains $G_{\text{tx}}(\theta)$ and G_{rx} to obtain L_{tot};

5. superimpose small-scale fading (MU–MIMO channel) for SNR/capacity computation.

The transmit gain $G_{\text{tx}}(\theta)$ follows the HAPS beam directivity and decays with angular offset from boresight; it is computed by the antenna/directivity module and combined with small-scale fading (Rician in LoS, Rayleigh in NLoS) at MU–MIMO channel level.

At S-band (~2 GHz), gaseous attenuation A_{gas} is small but non-zero [22]; rain attenuation A_{rain} becomes relevant under intense downpours, especially at low elevations [23]; cloud/fog attenuation $A_{\text{cloud/fog}}$ increases for high-LWC fogs [24]. Including L_{env} allows the simulator to reproduce realistic capacity dips during severe events while preserving accuracy under fair-weather conditions.

3.4 Association and Capacity Allocation Policies

This section formalises user–HAPS association and per-slot capacity allocation with explicit objectives and constraints. A modular link budget is adopted in which the end-to-end attenuation is expressed as the product of free-space loss, 3GPP TR 38.811 shadow/clutter terms, and a weather-induced loss L_{env} computed from EPW/WEA data (rain, fog, snow). The signal-to-interference-plus-noise ratio (SINR) for user u associated to HAPS h and served on beam/resource b is

$$\text{SINR}_{u,h,b} = \frac{P_{h,b}\, G_{h,b}(\theta_u)\, G_u}{L_{\text{FSPL}}\, L_{\text{shadow}}\, L_{\text{clutter}}\, L_{\text{env}}\, \big(1 + I_{u,h,b}/(N_0 B)\big)}, \tag{3.3}$$

where $G_{h,b}(\theta_u)$ is the transmit directivity towards the user elevation θ_u, G_u the user receive gain, $I_{u,h,b}$ the in-band interference, N_0 the thermal noise spectral density, and B the bandwidth. The achievable rate on resource b is

$$R_{u,h,b} = B_b \log_2\big(1 + \text{SINR}_{u,h,b}\big). \tag{3.4}$$

Given a set $\mathcal{B}_h$ of resources (PRBs/streams) on h, the instantaneous user rate is

$$R_{u,h} = \sum_{b \in \mathcal{B}_h} x_{u,h,b} R_{u,h,b}$$

with $x_{u,h,b} \in \{0, 1\}$ and $\sum_u x_{u,h,b} \leq 1$.

3.4.1 Best-HAPS Association Strategy

At each slot t, each user u selects the HAPS $h^\star$ that maximises a composite score balancing link quality and load/backhaul constraints:

$$S_{u,h} \;=\; w_u \cdot \underbrace{\log\big(1 + \mathrm{SINR}_{u,h}\big)}_{\text{link quality}} \cdot \underbrace{(1 - \rho_h)}_{\text{load headroom}} \cdot \underbrace{\eta_h^{\mathrm{BH}}}_{\text{backhaul factor}}, \tag{3.5}$$

where $\rho_h \in [0, 1)$ is the instantaneous load of HAPS h (resource occupancy) and $\eta_h^{\mathrm{BH}} \in (0, 1]$ captures the usable backhaul headroom for h (1 denotes unconstrained backhaul, values < 1 penalise saturated backhaul). The association rule is

$$h^\star(u) \;=\; \arg\max_{h \in \mathcal{H}_u} S_{u,h}, \tag{3.6}$$

with $\mathcal{H}_u$ the set of visible HAPS above a minimum elevation. Ties are broken by the largest headroom $(1 - \rho_h)$, then by the highest $R_{u,h}$.

Expression (3.5) ensures that high-priority users ($w_u > 1$) favour links with strong SINR but are discouraged from joining overloaded HAPS (large ρ_h) or backhaul-limited nodes (small η_h^{BH}). This stabilises the network against myopic "best-SINR only" decisions that would tend to collapse load onto a single platform.

The Python simulator executes the selection of the *best HAPS* through a sequential decision pipeline that evaluates all available platforms and assigns each area (or user cluster) to the HAPS offering the highest composite score $S_{u,h}$. This flow integrates radio conditions, platform load, and backhaul constraints, providing a dynamic and priority-aware association in each simulation slot.

Note. Figure 3.12 follows the canonical structure of a HAPS/NTN scheduler. The implemented simulator executes the blocks up to the Best-HAPS association and the update of ρ_h and η_h^{BH}.

The "weighted PF allocation" block is a conceptual placeholder: in practice, capacity is assigned using the admission strategies described in Sec. 3.4.2.

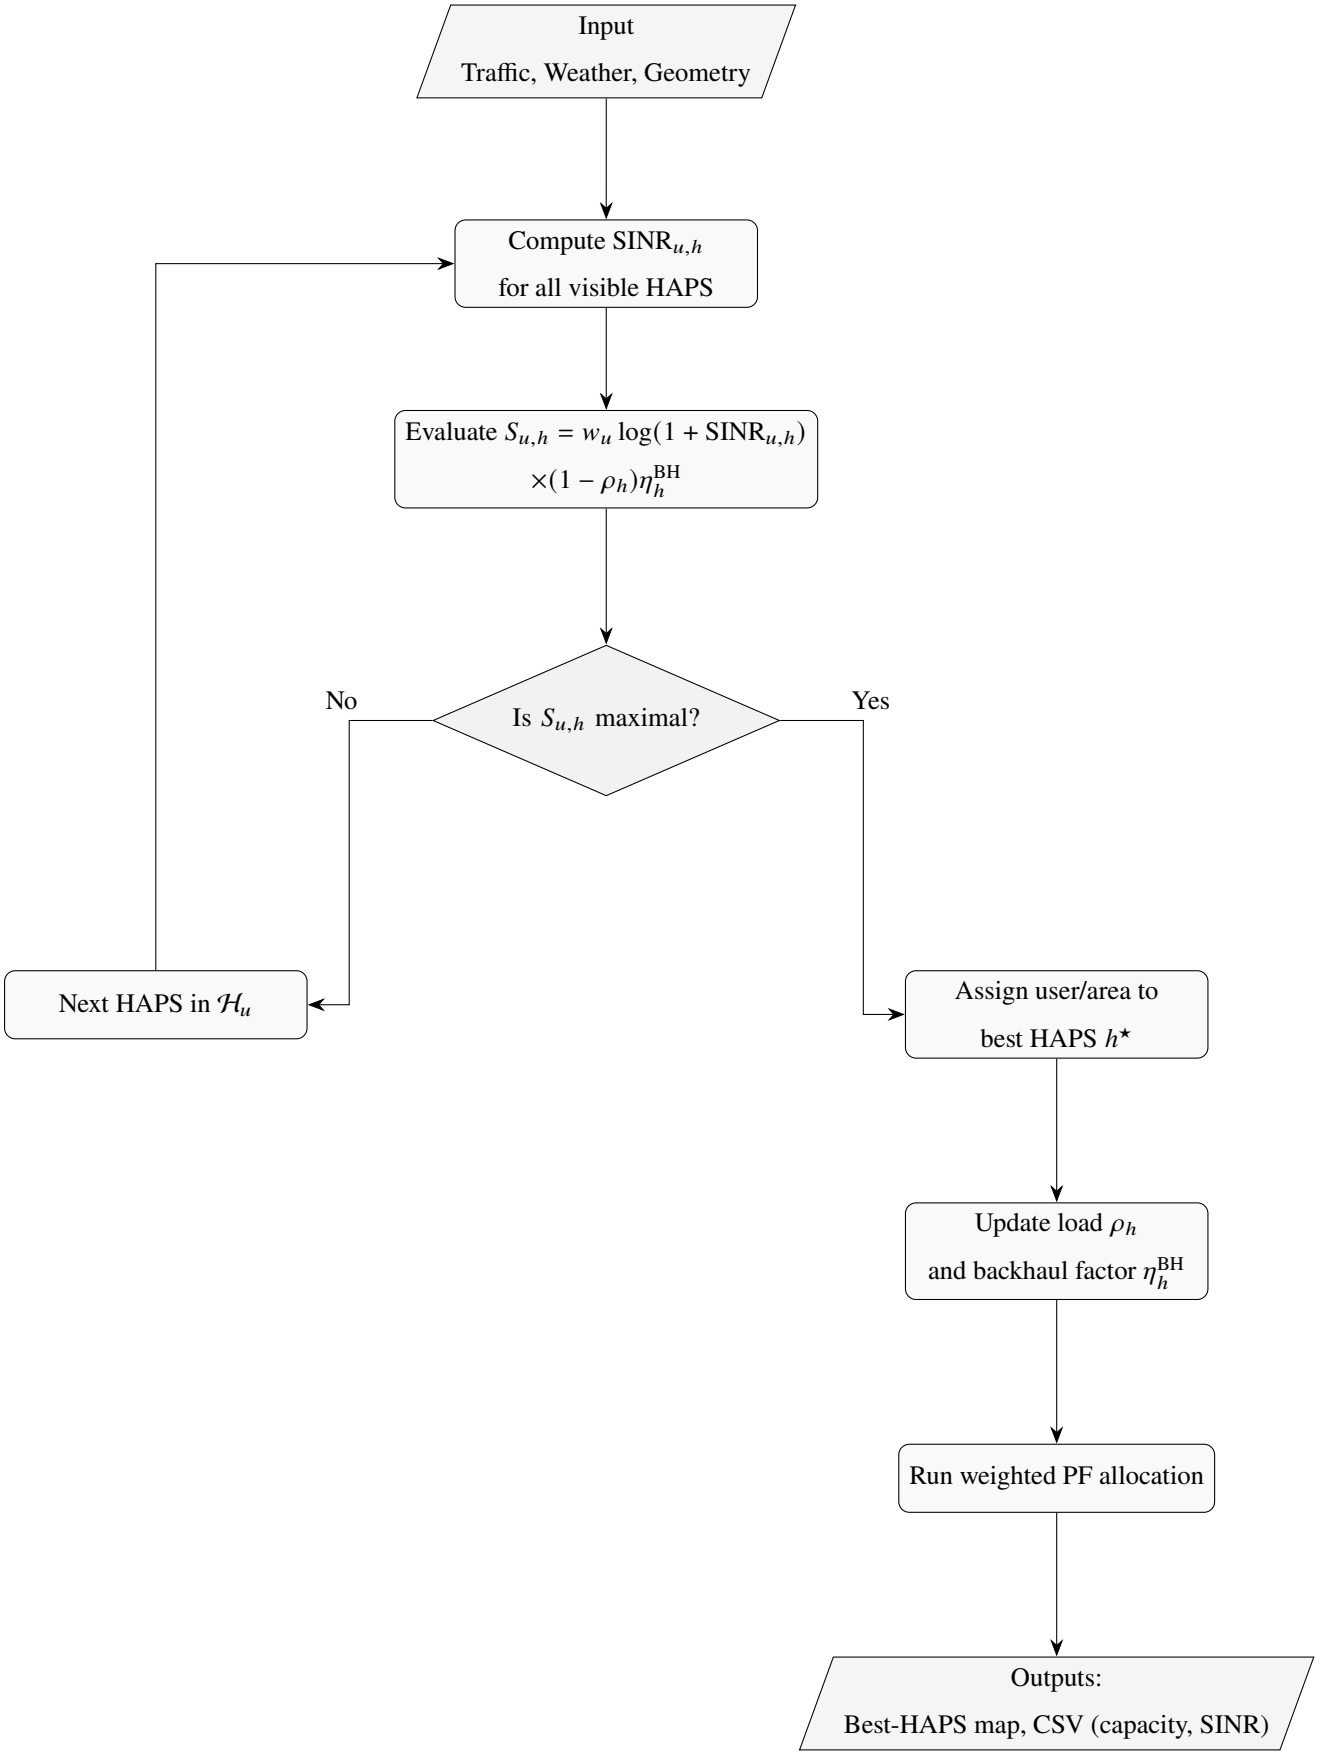

Figure 3.12: Operational flow of the Best-HAPS selection. Each user/area evaluates all visible HAPS, computes SINR, applies the composite score of Equation (3.5) *and associates to the platform with maximum value. The final "weighted PF allocation" block is conceptual, as capacity is actually distributed by the admission policies of Sec. 3.4.2.*

Algorithm 2 Best-HAPS association and weighted PF allocation per slot

Input: Users $\mathcal{U}$, HAPS set $\mathcal{H}$, priorities $\{w_u\}$, backhaul factors $\{\eta_h^{\mathrm{BH}}\}$, past throughputs $\{\bar{R}_u\}$.
Association:
for $u \in \mathcal{U}$ **do**
    Compute candidate set $\mathcal{H}_u$ (elevation $\geq \theta_{\min}$).
    for $h \in \mathcal{H}_u$ **do**
        estimate $\mathrm{SINR}_{u,h}$ and $\hat{R}_{u,h}$ with current L_{env}.
        compute score $S_{u,h}$ via (3.5).
    end for
    associate u to $h^\star(u) = \arg\max_h S_{u,h}$; update provisional loads ρ_h.
end for
Allocation (per HAPS):
for $h \in \mathcal{H}$ **do**
    $\mathcal{U}_h \leftarrow$ users associated to h; $\mathcal{B}_h \leftarrow$ available resources.
    for $b \in \mathcal{B}_h$ **do**
        compute PF metrics $M_{u,h,b} = w_u \cdot R_{u,h,b}/(\bar{R}_u + \epsilon)$ for $u \in \mathcal{U}_h$;
        allocate b to $u^\star = \arg\max_u M_{u,h,b}$;
    end for
end for
Update: For all u, update $\bar{R}_u \leftarrow (1-\alpha)\bar{R}_u + \alpha R_u$.

This schematic corresponds to the logic implemented in `SIMULATOR_PARIS.ipynb`: for each polygonal area, the program computes SINR and capacity from channel parameters, derives a score $S_{u,h}$ that merges physical and load-related aspects, selects the best HAPS, and finally stores the outputs in `capacities_output.csv`. The structure guarantees reproducibility and allows post-processing such as ranking of alternative HAPS, resilience evaluation, and capacity heatmap visualisation.

At each time slot, every user (or demand pixel) therefore selects the hosting platform that maximises

$$S_{u,h} = w_u \cdot \log\bigl(1 + \mathrm{SINR}_{u,h}\bigr) \cdot (1 - \rho_h) \cdot \eta_h^{\mathrm{BH}}.$$

The term $\log(1 + \mathrm{SINR}_{u,h})$ captures instantaneous link quality, w_u encodes mission priori-

ties, $(1-\rho_h)$ prevents piling up users on saturated platforms, and η_h^{BH} penalises HAPS with constrained backhaul. Association is greedy on $S_{u,h}$. In standard HAPS/NTN schedulers, the association step is typically followed by a weighted proportional-fair allocator. In the present simulator, this block is represented conceptually in Fig. 3.12, but it is not executed: once users/areas are attached, capacity is assigned through the admission strategies of Sec. 3.4.2.

3.4.2 Capacity Allocation Policies

In order to decide which services are admitted on each HAPS beam, three admission strategies are considered. All policies operate on the same inputs (beam capacity, per–area traffic demand, and service priority labels), but they implement different optimisation criteria and degrees of flexibility.

Priority-for-area-and-services (manual rule): The first strategy is a rule-based policy that enforces strict prioritisation both at area and service level. Areas labelled as *High* priority are always served before *Medium* and *Low* ones. Within each area, services are ordered according to their class (priority vs. non-priority) and to their offered traffic load. Traffic is admitted greedily while residual capacity is available on the serving beam. Simple thresholds on the per-service demand are used to avoid allocating very small residual fragments of capacity. As a consequence, best-effort services may be systematically dropped in order to preserve capacity for mission-critical traffic.

Knapsack-greedy admission policy: The second strategy casts the admission problem as a knapsack-like optimisation. Each candidate service is modelled as an "item" with a size equal to its requested bit rate and a profit that depends on the service and area priority. The beam capacity plays the role of the knapsack size. A greedy algorithm sorts all items according to their profit-to-size ratio and admits them sequentially until the residual capacity becomes insufficient. This policy does not enforce hard constraints on priority services: instead, it aims at maximising the total profit, which typically leads to a larger number of admitted services, at the cost of a weaker protection for critical traffic.

Reliability-aware Q-learning policy: The third strategy is a learning-based policy

that uses reliability-aware Q-learning. The environment state encodes the current beam utilisation, the number of active services in each priority class, and the presence of blackout areas. Actions correspond to discrete admission decisions (e.g., admit/reject a candidate service, or select which area to serve next). After each episode, the agent receives a reward that increases with the number of priority services successfully admitted in critical areas and penalises outages or blackouts. The Q-table is updated using the standard temporal-difference rule with learning rate and discount factor tuned empirically. Once training converges, the learned policy is frozen and used in the simulator to take admission decisions without further exploration.

3.5 Simulator Architecture

The simulator runs in discrete time slots. At each slot it combines traffic demand, channel conditions, and scheduling decisions to produce per-area capacities and performance indicators.

Figure 3.13 summarises the main loop. The inputs to this loop are the traffic and priority information per area (derived from the tensor in Sec. 3.5.1), the weather-induced attenuation L_{env} described in Sec. 3.3.7, the channel parameters and directivity gains of Sec. 3.3.1, and the overall system configuration, including the `MIMOConfig`, HAPS geometry, and beam layout.

For each simulation slot the following steps are executed:

1. per-area traffic demand is read from the tensor and scaled;
2. channel quantities (slant range, elevation, path loss, directivity, L_{env}) are updated;
3. the Best-HAPS score $S_{u,h}$ in (3.5) is evaluated and users/areas are associated (Sec. 3.4.1);
4. the selected admission policy (manual, knapsack, or Q-learning; Sec. 3.4.2) allocates capacity on each beam;

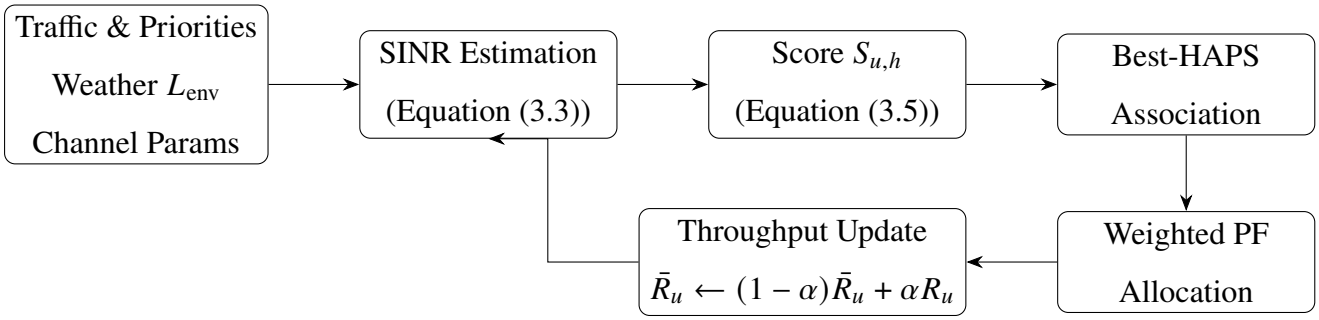

Figure 3.13: Slot-level loop: inputs → SINR → Best-HAPS association. The PF allocation and throughput-update blocks reflect the canonical structure of multi-user schedulers, but in the current simulator capacity is allocated through the admission policies of Sec. 3.4.2 rather than via a PRB-level PF scheduler.

5. optionally, long-term throughputs $\bar{R}_u$ can be updated in PF-type extensions; this step is not executed in the current implementation.

The main outputs of each run are the capacity per area and per service class, the percentage of priority areas that can be kept served, and the SINR and rate distributions (for both direct and relay links, when present), together with aggregate fairness and utilisation indicators and log files for post-processing (maps, tables, resilience metrics).

Unless otherwise stated, all scenarios are run under the three admission strategies defined in Sec. 3.4.2.

3.5.1 Traffic tensor construction, unit conversion, and scaling

The enriched trace files are consolidated into a single `NumPy` tensor. Let $N_{BS} = 108$, $N_S = 68$, and $N_T = 672$. The array $\mathbf{X}^{(\text{bytes})} \in \mathbb{R}^{N_{BS} \times N_S \times N_T}$ stores traffic values per BS, per service, and per 15-minute slot (one week of data). Both raw and converted versions are saved:

- `big_matrix_bs1_108_service1_68_time1_672.npy` (bytes),
- `big_matrix_MB_bs1_108_service1_68_time1_672.npy` (MB).

The temporal dimension consists of $N_T = 672$ consecutive 15-minute slots, aligned across all services and BS. Missing samples are set to zero to preserve time consistency.

Traffic volumes are converted from bytes to megabytes (MB) as

$$\mathbf{X}^{(\mathrm{MB})} = \frac{\mathbf{X}^{(\mathrm{bytes})}}{1024^2}.$$

To obtain a more realistic load, the traces are scaled by a constant factor $\alpha = 40$, reflecting the fact that the measurement campaign captures only a fraction of active users per BS:

$$\widetilde{\mathbf{X}} = \alpha\, \mathbf{X}^{(\mathrm{MB})}.$$

The scaled tensor, `big_matrix_MB_scaled40x_bs1_108_service1_68_time1_672.npy`, is used as input to the HAPS simulator. Per-BS totals

$$T_{\mathrm{BS}}(b,t) = \sum_s \widetilde{X}(b,s,t)$$

are inspected to check consistency. After scaling, slot-level totals fall in the $\sim$ 150–300 MB range, in line with typical LTE dense-urban cell loads.

3.6 Model Validation

The simulation framework brings together traffic modelling, geometry, antenna directivity, propagation (3GPP + ITU-R), weather attenuation, and scheduling. Several consistency checks are performed on these components before running full-scale scenarios.

3.6.1 Link-budget consistency

The end-to-end attenuation model is checked by comparing the free–space loss used in the simulator with the analytical expression

$$L_{\mathrm{FSPL}}(f,d) = 32.45 + 20\log_{10}(f_{\mathrm{MHz}}) + 20\log_{10}(d_{\mathrm{km}}).$$

Across representative slant ranges (10–30 km), the difference remains below 0.1 dB, indicating that unit conversions and distance calculations are handled correctly.

When shadow–clutter terms (3GPP TR 38.811), antenna gain from the parabolic model, and thermal noise are combined, the resulting SINR values fall in the expected range for S-band HAPS deployments.

3.6.2 Directivity–geometry alignment

To test the antenna and steering logic, the boresight of each HAPS beam is matched against the ground position of every polygon. The examples in Fig. 3.9 show 3D patterns for two representative areas (area_1 and area_43) served by HAPS #1. The main lobes rotate and stretch according to off–nadir angle and slant range, which confirms that the parabolic directivity model is applied consistently over the footprint and that beam steering reacts correctly to the underlying geometry. The blackout configuration in Fig. 4.1 is used as a reference testbed for these checks, since the affected areas form a compact cluster that requires moderate but non-negligible beam steering from the serving HAPS.

3.6.3 Weather-induced attenuation

The weather module is validated by extracting hourly maxima of L_{env} from the EPW (TMY) dataset and comparing them with the attenuation predicted by ITU-R P.676, P.838, and P.840 for the same temperature, humidity, rain rate, and cloud liquid water content. Differences remain within 0.1–0.3 dB on a 20 km slant path.

The comparison between SNR at selection time and final SNR (with and without weather), shown in Fig. 3.10a, indicates that weather losses shift all SNR values downward but preserve the relative ranking across areas. This behaviour suggests that the Best-HAPS association, which is based on pre-weather estimates, remains stable when the full weather-aware link budget is applied.

3.6.4 End-to-end pipeline check

A final check is carried out on the complete pipeline by running the simulator on a minimal setting (five areas, deterministic traffic) and comparing internal logs against analytical calculations:

- geometric distances and elevation angles match GeoPandas outputs;
- antenna gains match `Directivity_Gain.py`;
- path loss agrees with 3GPP and ITU-R references;
- final SINR matches the analytical $P/(L\,N)$ formulation;
- the beam capacities and admitted services match the outcome of the admission policies implemented in the code (manual, knapsack, Q-learning), given the same inputs.

Across the pipeline, the maximum deviation observed is below 0.5 dB at SINR level, corresponding to less than 2% relative difference in achievable rate. Overall, the simulator combines geometric, radiative, and atmospheric components in a numerically coherent way, with behaviour consistent with the underlying 3GPP and ITU-R models.

Chapter 4

Results

This chapter presents the main findings obtained with the simulation framework. Starting from a baseline configuration, the analysis examines how performance changes under different conditions: area coverage, service prioritisation, stress scenarios, and scaling the number of HAPS. Alternative association and admission strategies are compared, with particular attention to their impact on critical infrastructures in blackout conditions.

4.1 Blackout Scenario and Simulation Setup

The analysis follows the same logic adopted in the simulator pipeline: (i) blackout zones are defined and classified as High or Low priority; (ii) the capacity that the HAPS network can deliver to these zones is estimated; (iii) Best-HAPS association is computed via the composite score $S_{u,h}$; (iv) the three admission strategies are applied to distribute per-beam capacity across services; and (v) the coverage envelope is studied when moving from two to three HAPS over a Paris-like footprint.

4.1.1 Definition of blackout zones

To assess the framework under stress conditions, a large-scale *blackout scenario* is considered, affecting a subset of the 108 zones in the Paris area. The outage is assumed to impact both power and terrestrial communications, so that only the HAPS overlay can provide

connectivity to the selected zones. The choice of which polygons are in blackout is guided by the geographical priority map introduced in Chapter 3 (Fig. 3.6). In the baseline configuration, the blackout primarily involves *High*-priority zones (hospitals, emergency call centres, major transport hubs), while extended scenarios progressively add surrounding *Low*-priority residential and commercial areas to emulate cascade effects and spill-over of the crisis.

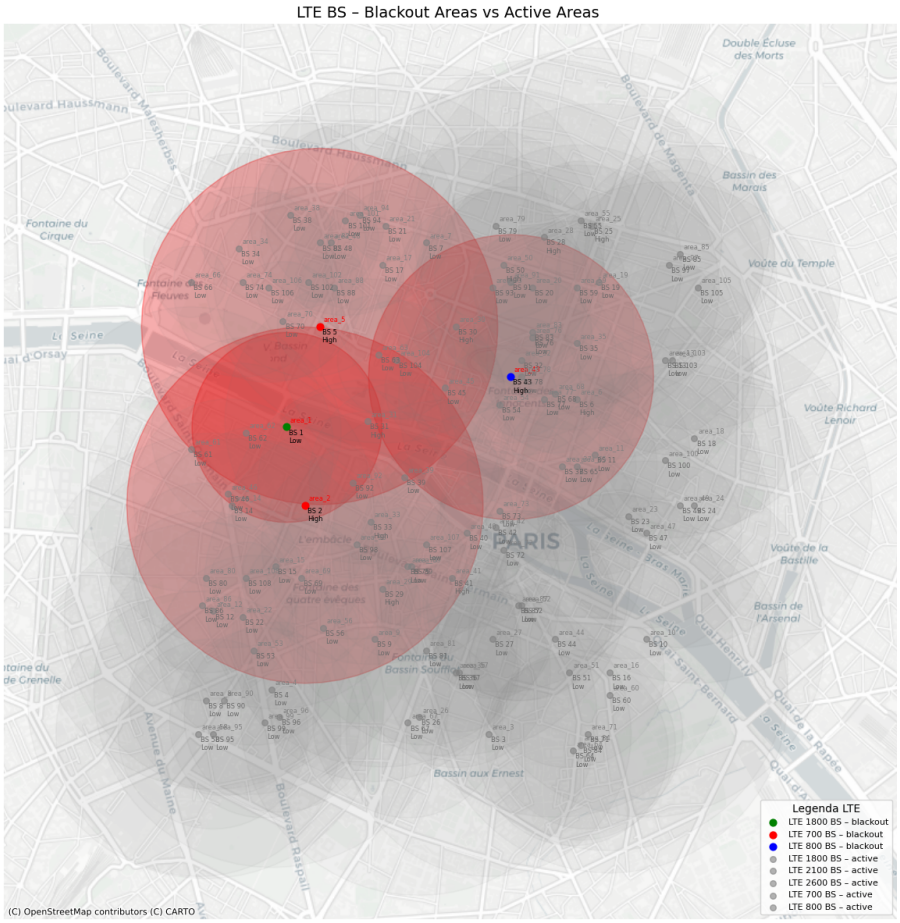

Figure 4.1: Blackout zones considered in the simulation.

In the simulator, each polygon can be independently flagged as (i) High or Low priority and (ii) affected or not by the blackout. This parametrisation controls both the *severity* of the outage (number of zones in blackout) and its *composition* in terms of critical vs. non-critical areas. By tuning these two levers, a range of realistic stress scenarios can be explored and the capability of the HAPS layer to sustain essential services while preserving fairness across the footprint can be quantified.

4.2 HAPS-generated Capacity over Blackout Zones

For a given blackout configuration, the first step consists in estimating the capacity that the HAPS network can effectively deliver to the affected zones. For each time slot and for every candidate HAPS, the simulator computes the weather-aware link budget, the resulting SINR, and the achievable rate toward each blackout polygon, using the channel and propagation models defined in Chapter 3.

Figure 4.2 reports the per-beam capacities obtained in the reference blackout scenario, while Fig. 4.3 shows the aggregate capacity per HAPS. Together, these plots provide a first indication of how much traffic the aerial overlay can sustain over the blackout footprint and how capacity is distributed across beams and platforms.

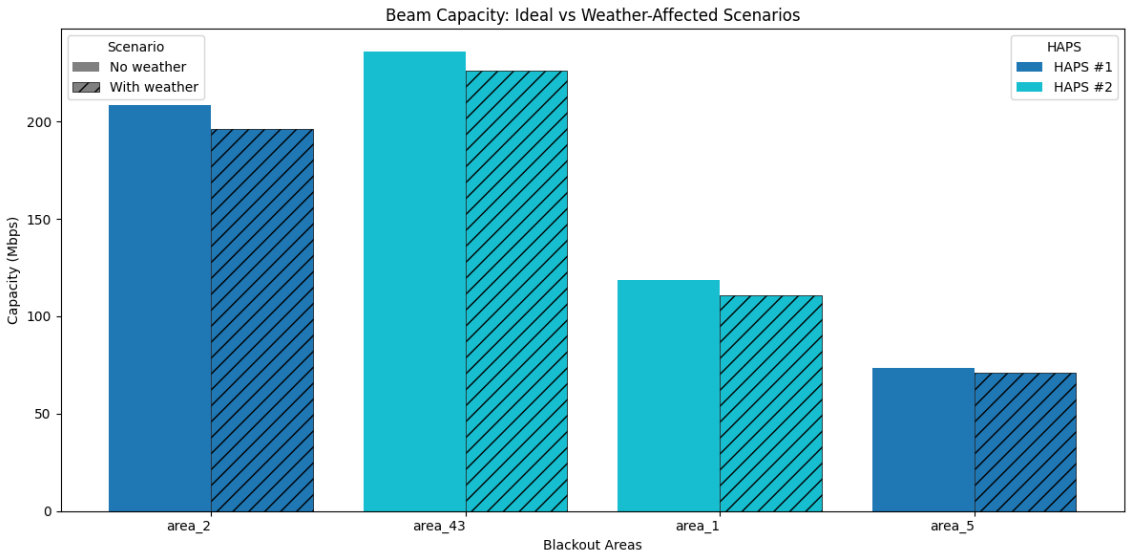

Figure 4.2: Per-beam capacities in the blackout scenario.

These capacity figures serve as the baseline for the subsequent association and admission steps: they represent the resource budget available to serve priority and non-priority services in the affected zones. In the reference configuration with two S-band HAPS and four blackout zones (two High-priority cores and two mixed-priority surrounding areas), the per-beam capacities of Fig. 4.2 translate into an aggregate capacity of approximately 636 Mbit s^{-1} without weather and 604 Mbit s^{-1} with weather. In this setting, HAPS #1 serves `area_2` and `area_5` with a total of about 282 Mbit s^{-1} (reduced to 267 Mbit s^{-1} when weather losses are included), while HAPS #2 serves `area_43` and `area_1` with

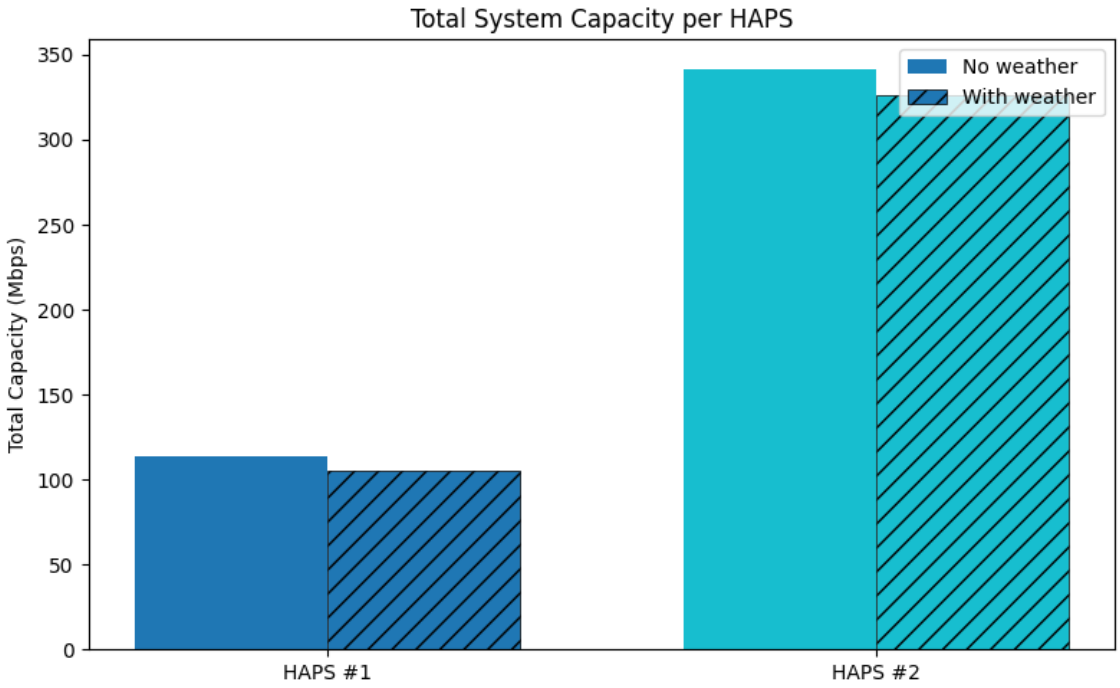

Figure 4.3: Per-HAPS aggregate capacities in the blackout scenario.

roughly 354 Mbit s^{-1} (down to 337 Mbit s^{-1} with weather). At beam level, the resulting rates fall in the range 70–235 Mbit s^{-1}, consistently with the SNR values (about 22–35 dB) produced by the link-budget and small-scale fading model of Chapter 3. These numbers correspond to spectral efficiencies on the order of 6–12 bit/s/Hz over a 20 MHz channel, which is compatible with advanced coded modulation schemes and therefore physically plausible for a high-gain S-band HAPS link. The "with weather" curves exhibit a moderate degradation (about 5% in the aggregate), as expected at S-band, where rain and clouds increase the path loss but do not dominate the budget as in higher-frequency bands.

All capacity values reported in this chapter should be interpreted as upper bounds at the physical layer. The simulator computes

$$C = B \log_2(1 + \mathrm{SNR})$$

from the weather-aware link budget, antenna directivity and small-scale fading, but does not explicitly model protocol overheads, coding gaps with respect to Shannon, HARQ retransmissions, scheduling inefficiencies or transport-layer dynamics. The curves in Figs. 4.2–4.3 therefore indicate how much throughput the HAPS layer could in principle sustain under the assumed conditions, rather than the exact user-perceived rates of a

specific air-interface implementation. This abstraction is consistent with the objective of the study, which is to compare association and admission strategies under a common and physically grounded capacity budget, rather than reproducing the performance of a particular standard.

4.3 User–HAPS Association in Blackout Conditions

The association between blackout zones and platforms is performed by the *Best-HAPS* selection routine (Alg. A.1). For each blackout area and each candidate HAPS h, the simulator computes a weather-aware link budget and evaluates the achievable rate over the effective bandwidth

$$B_{\mathrm{eff},h} \;=\; \frac{B}{n_h + 1},$$

where B is the total S-band bandwidth and n_h is the number of beams already active on HAPS h. The channel gain combines path loss (with and without weather), antenna directivity, and small-scale fading, as described in Chapter 3. Thermal noise is modelled via a noise power density P_n in dBm/Hz, which is converted to noise power over $B_{\mathrm{eff},h}$.

The routine assigns a *selection score* to each candidate according to the chosen metric. In the blackout experiments discussed here, the selection metric is capacity-based (`selection_metric = "capacity"`), so that each area is attached in first instance to the HAPS that offers the highest predicted rate under the current load (through $B_{\mathrm{eff},h}$). Weather-affected capacities are still computed and logged, but are not used directly as selection scores in this configuration.

The selection step therefore translates the "raw" capacity that each platform could offer into a concrete per-area assignment, taking into account: (i) propagation conditions and directivity; (ii) thermal noise; and (iii) the number of beams already active on each HAPS (via $B_{\mathrm{eff},h}$). The subsequent admission and scheduling stages then refine this allocation under global budget constraints and per-beam caps.

Table 4.1 illustrates the behaviour of the `select_best_haps_for_area` routine in a four-zone blackout configuration. For each area, both HAPS candidates are evaluated and the platform with the largest capacity-based score is selected, unless beam/budget

Table 4.1: Best-HAPS candidate capacities in the four-zone blackout run.

Area ID	Priority	HAPS #	$C^{(\mathrm{sel})}$ [Mbps]	$C_{\mathrm{w}}^{(\mathrm{sel})}$ [Mbps]	Selected
area_2	High	1	157.5	148.1	
		2	210.98	195.49	✓
area_43	High	1	235.39	225.29	✓
		2	158.6	150.2	
area_1	Low	1	112.7	105.4	
		2	120.65	110.81	✓
area_5	Low	1	75.56	71.57	✓
		2	19.6	17.1	

constraints prevent its allocation. For instance, in area_2 the routine compares approximately 157.5 Mbps from HAPS #1 with 210.98 Mbps from HAPS #2 and assigns area_2 to HAPS #2. In area_43, HAPS #1 offers the largest capacity (about 235.39 Mbps versus 158.6 Mbps from HAPS #2) and is therefore selected. Similarly, area_1 is attached to HAPS #2, which provides slightly higher capacity than HAPS #1 (120.65 Mbps versus 112.7 Mbps), while area_5 is served by HAPS #1, since the alternative link from HAPS #2 would be strongly capacity-limited. Overall, the example confirms that the selection routine consistently favours the HAPS that offers the largest per-area capacity, while also accounting for weather-induced losses through the effective capacity $C_{\mathrm{w}}^{(\mathrm{sel})}$.

In the final allocation, the capacities reported in Figs. 4.2–4.3 are adjusted to respect the per-beam cap (280 Mbps in S-band) and the global 2 Gbps budget per HAPS. In the configuration of Table 4.1, the resulting beams lead to total per-HAPS loads of approximately 115 Mbps for HAPS #1 and 340 Mbps for HAPS #2 in ideal conditions, decreasing to about 105 Mbps and 320 Mbps, respectively, under worst-day weather. These values remain well below the 2 Gbps platform limit and can be regarded as realistic *upper bounds* for a lightly loaded two-HAPS deployment over a Paris-like blackout footprint, where ample capacity headroom remains available for additional blackout areas and subsequent time slots.

4.4 Admission Strategies in Blackout Conditions

Once the Best-HAPS association is fixed for a given time slot, the per-beam capacity is distributed among services according to the admission strategies defined in Sec. 3.4.2. All strategies operate on the same inputs (beam capacity, per-area traffic demand, and service-priority labels), but implement different trade-offs between protection of critical services and overall resource utilisation.

The comparison focuses on the number of services that can be kept active in the blackout areas. For each area, the simulator counts how many priority and non-priority services are admitted over the considered time window, and the results are then aggregated across areas and strategies.

4.4.1 Strategy 1: Priority-for-area-and-services

The first strategy is a rule-based policy that enforces strict prioritisation at both area and service level: High-priority areas are served before Low-priority ones, and within each area priority services are always admitted before non-priority traffic.

Figure 4.4 shows the behaviour of *Strategy 1: Priority-for-area-and-services*. In all blackout areas only priority services are admitted, while best-effort traffic is systematically rejected. The dense-urban beams serving `area_2` and `area_43` accommodate, respectively, 19 and 22 out of the 23 available priority services, using almost the entire beam capacity (about 191–225 Mbit/s out of 196–226 Mbit/s). The suburban beams (`area_5` and `area_1`) host 10 and 7 priority services, respectively, reflecting both lower local demand and smaller beam capacities. Overall, each blackout area can host between roughly 7 and 22 priority services, and all residual capacity that could be used for non-priority sessions is deliberately left unused in order to provide hard protection for critical services.

4.4.2 Strategy 2: Knapsack-greedy

The second strategy casts the admission problem as a knapsack-like optimisation. Each candidate service is modelled as an item with a size equal to its requested bit rate and a

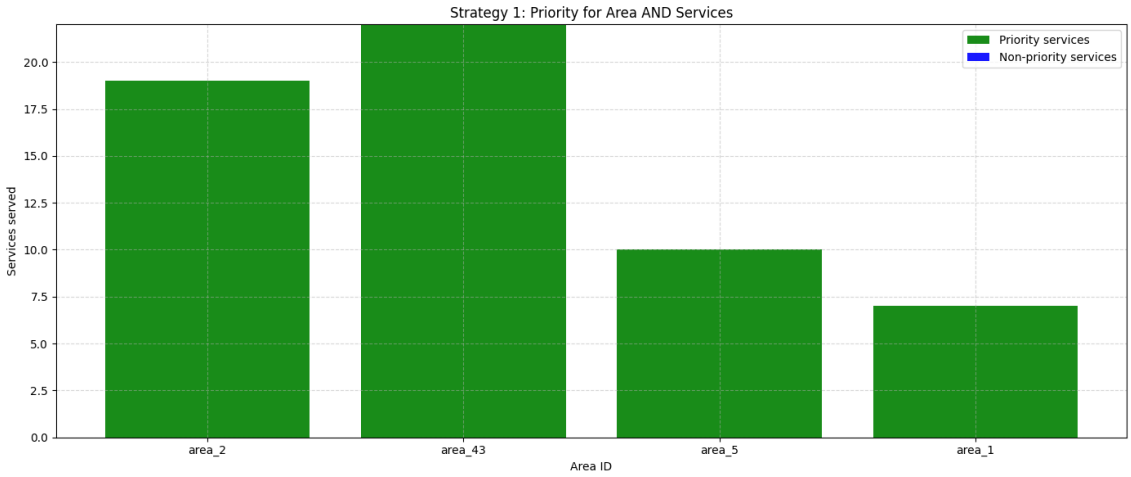

Figure 4.4: Per-area number of admitted services under Strategy 1: Priority-for-area-and-services. *Only priority services are admitted, whereas non-priority traffic is systematically dropped.*

profit that depends on the service and area priority. The beam capacity plays the role of the knapsack size, and a greedy heuristic packs items in increasing demand order so as to maximise the number of admitted services.

Figure 4.5 shows the per-area breakdown for the knapsack-greedy policy. In this case all services are treated homogeneously by the allocator, and priority is only used a posteriori to classify the outcome. As a result, all blackout areas host many more active services than under Strategy 1: `area_2` admits 41 services (15 priority and 26 non-priority), `area_43` admits 54 (20 priority and 34 non-priority), while the suburban `area_5` and `area_1` admit 33 and 31 services, respectively (about one third of which are priority). Beam capacities are almost fully exploited in all cases (around 183–211 Mbit/s out of 196–226 Mbit/s in dense-urban areas and 69–103 Mbit/s out of 71–111 Mbit/s in suburban areas), but a substantial fraction of the admitted sessions is non-priority traffic. This behaviour reflects the objective of the greedy knapsack formulation, which tends to pack as many sessions as possible without enforcing strict protection for critical services; compared to Strategy 1, several priority services in dense-urban blackout areas are sacrificed in order to accommodate a large amount of best-effort traffic.

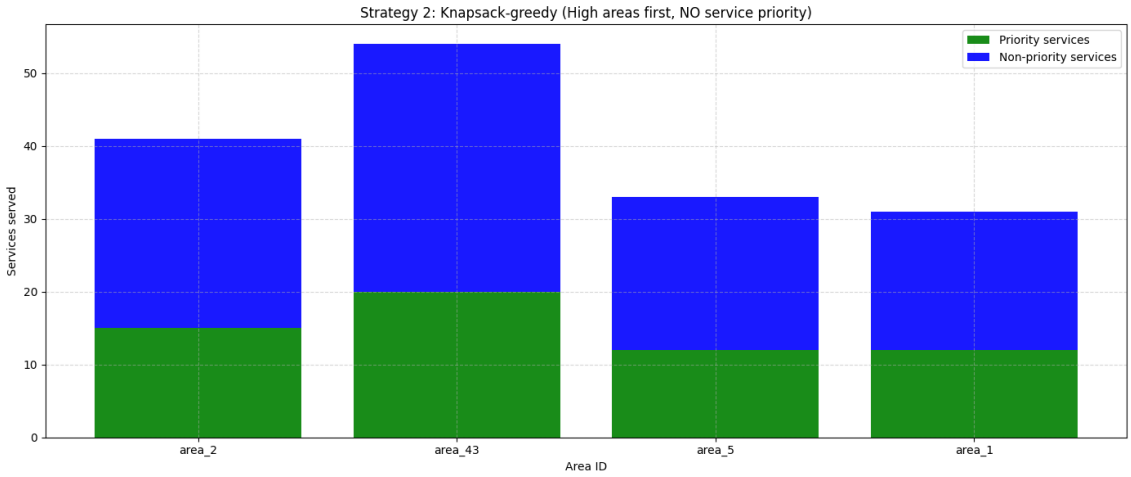

Figure 4.5: Per-area number of admitted services under Strategy 2: Knapsack-greedy. *The policy maximises the total number of admitted services, leading to a significant fraction of non-priority traffic in high-capacity areas.*

4.4.3 Strategy 3: Reliability-aware Q-learning

The third strategy is a learning-based admission policy that uses a tabular, reliability-aware Q-learning agent. The environment state aggregates a small number of discrete indicators:

- **Area priority** (2 levels): high- vs. low-priority area;
- **Service priority** (2 levels): priority vs. non-priority service;
- **Offered load** (3 levels): low/medium/high demand, normalised to the nominal clear-sky capacity of the area;
- **Residual capacity** (3 levels): low/medium/high remaining capacity on the beam serving the area;
- **Time band** (7 levels): day-of-week index (only daily variability is kept).

Actions are binary admission decisions (*allocate* vs. *reject*) for each service request. This leads to a compact Q-table with $2 \times 2 \times 3 \times 3 \times 7 \times 2 = 504$ state–action pairs, which is small enough to be learned reliably while still encoding the main dimensions that affect robustness in blackout conditions.

The reward function is explicitly shaped to favour long-term reliability of priority traffic. Invalid allocations (attempting to admit a service whose demand exceeds the

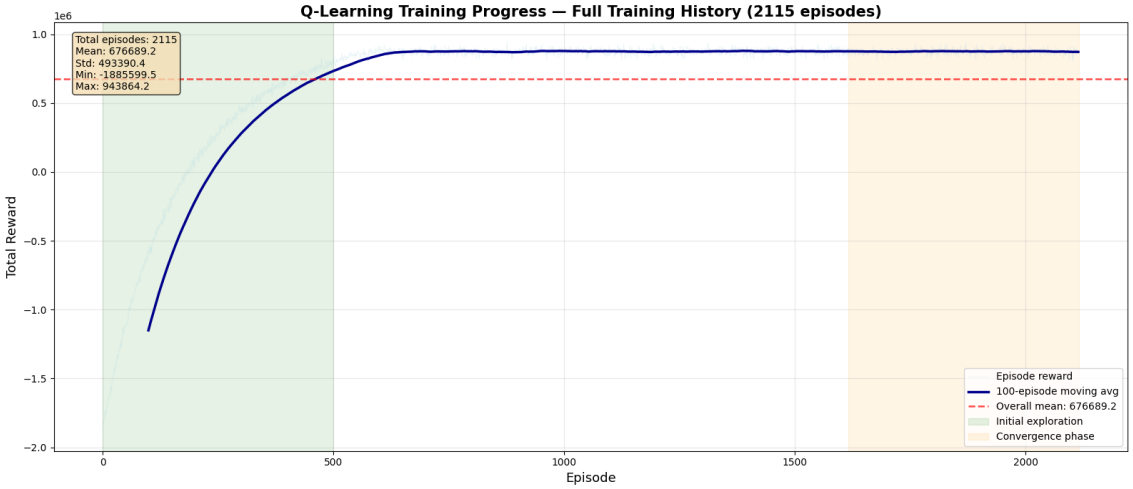

Figure 4.6: Training curve of the reliability-aware Q-learning policy in the blackout scenario (2115 episodes). The light blue line shows the total reward per episode, the dark blue line is the 100-episode moving average, and the red dashed line marks the overall mean reward. The shaded regions highlight the initial exploration phase (green) and a convergence window (orange).

residual capacity) incur a large negative reward, while rejections are mildly penalised, especially for priority services. When a request is admitted, the agent receives a positive reward that is increased if (i) the area is high-priority, (ii) the service belongs to the priority class, (iii) the allocation uses capacity efficiently (intermediate utilisation levels are rewarded), and (iv) the link survival probability p_{survive} is high. A small time-of-day bonus is also included to encourage protection of evening peak hours. Overall, this design pushes the agent towards decisions that keep priority services active over long horizons, rather than greedily maximising short-term throughput.

Training is carried out over several blackout weeks using an ϵ-greedy exploration policy with decaying ϵ, learning rate $\alpha = 0.1$ and discount factor $\gamma = 0.95$. The learning process is monitored through the total reward per episode. The resulting curve is shown in Fig. 4.6, which reports both the raw episode rewards and their 100-episode moving average over a training history of 2115 episodes. The shaded green region highlights the initial exploration phase, where the agent frequently takes sub-optimal actions and the reward is highly variable. After a few hundred episodes the moving average increases monotonically and approaches a plateau. The orange shaded area marks a convergence window in which the moving average remains essentially flat, indicating that the Q-values have stabilised around a consistent policy.

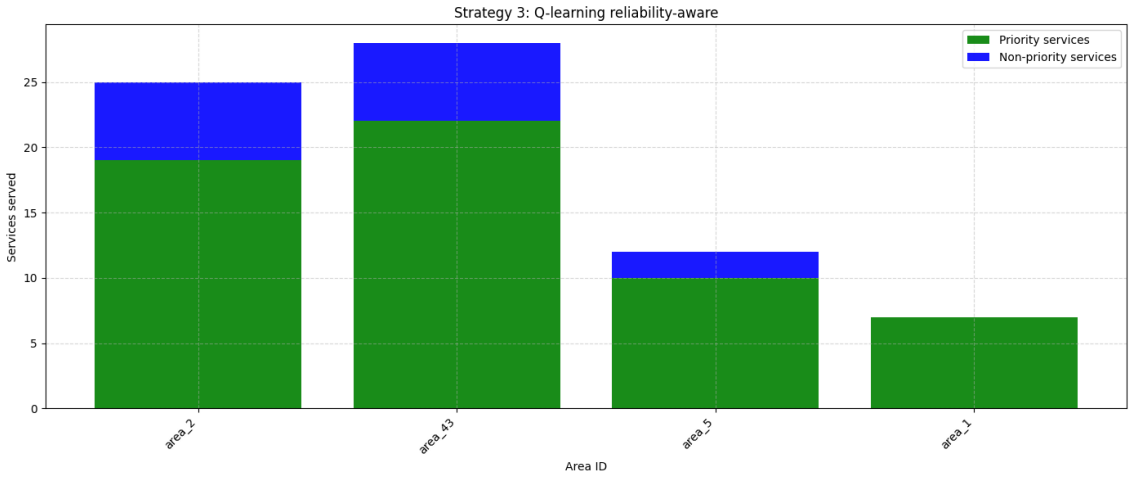

Figure 4.7: Per-area number of admitted services under Strategy 3: Reliability-aware Q-learning. *Green bars indicate priority services, blue bars non-priority ones. The learned policy strongly favours priority traffic and admits non-priority services only when sufficient residual capacity is available.*

Once training has converged, the learned policy is frozen and evaluated on the blackout areas. Figure 4.7 reports, for each area, the number of admitted services split into priority and non-priority classes. Compared with the knapsack–greedy baseline, almost all admitted services belong to the priority class; non-priority traffic is only served when the agent detects clear residual capacity, and disappears entirely in the most congested area. This confirms that the reliability-aware reward design is effective in steering the agent towards solutions that consistently protect mission-critical traffic, while still exploiting spare capacity for non-priority services when available.

4.4.4 Strategy comparison

The overall comparison between the three strategies is summarised in Fig. 4.8. The knapsack-greedy policy admits the largest number of services (about 160 in this run), but only around one third of them belong to the priority class. The reliability-aware Q-learning strategy admits fewer services overall (around 70), yet almost all of them are priority services, while using a comparable amount of capacity. The manual priority-only policy remains the most conservative option and may admit no services when the combination of demand and reliability constraints becomes too strict (around 60). These

results highlight the trade-off between maximising the number of active sessions and protecting mission-critical traffic.

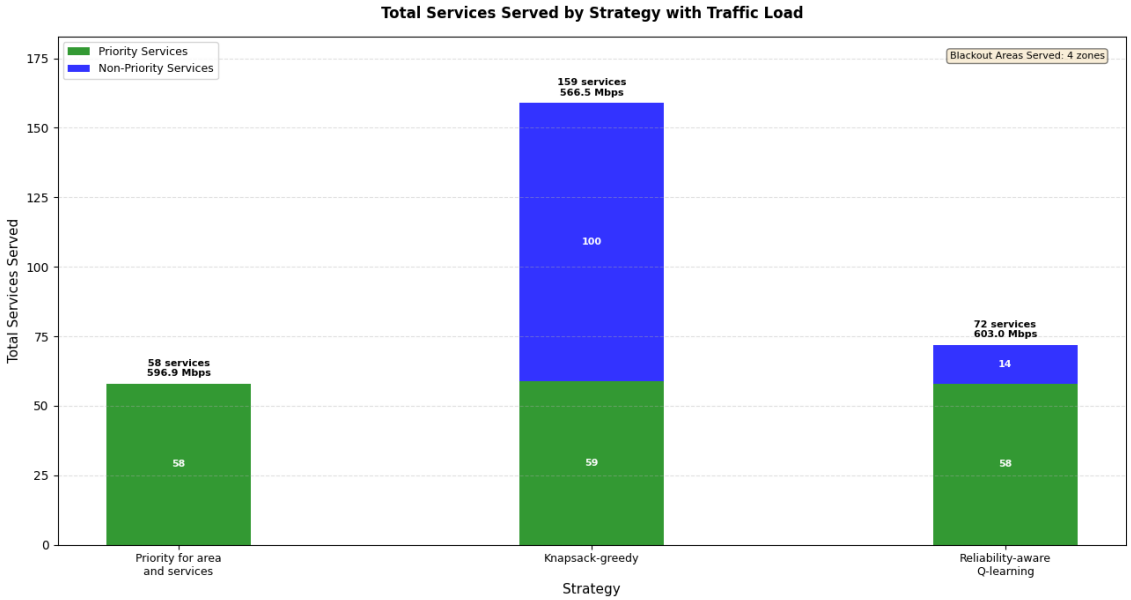

Figure 4.8: Total number of admitted services (priority and non-priority) and corresponding traffic load for the three admission strategies in the blackout scenario. The reliability-aware Q-learning policy achieves a higher share of priority services while exploiting a similar capacity budget.

4.5 Coverage Envelope with Two HAPS

To gauge the practical feasibility of HAPS support over a dense metropolitan area, a minimal overlay composed of *two* S-band HAPS is considered, and the number of zones in blackout that they are required to sustain is progressively increased.

The analysis starts from a compact configuration with four blackout zones (two High-priority cores and two surrounding areas with mixed priority), and then iteratively adds additional polygons. For each step, the simulator:

1. recomputes the Best-HAPS association using the routine of Alg. A.1;

2. applies the three admission strategies discussed above;

3. records per-zone capacity, priority coverage, and blocking metrics.

The process continues until adding further blackout zones either yields negligible marginal benefit (in terms of additional served critical traffic) or leads to unacceptable degradation of High-priority coverage. This parametric sweep is used to identify an approximate *maximum coverage envelope* for a two-HAPS deployment over a Paris-like city, given the available capacity and the constraint of at most sixteen simultaneous beams (eight per HAPS).

In practice, for the traffic and channel conditions considered in this thesis, two HAPS can sustain a compact blackout cluster while preserving acceptable QoS for High-priority areas. As the blackout footprint grows, additional affected zones must either share already saturated beams or accept reduced service levels, signalling the need for further platforms or relaxed performance targets.

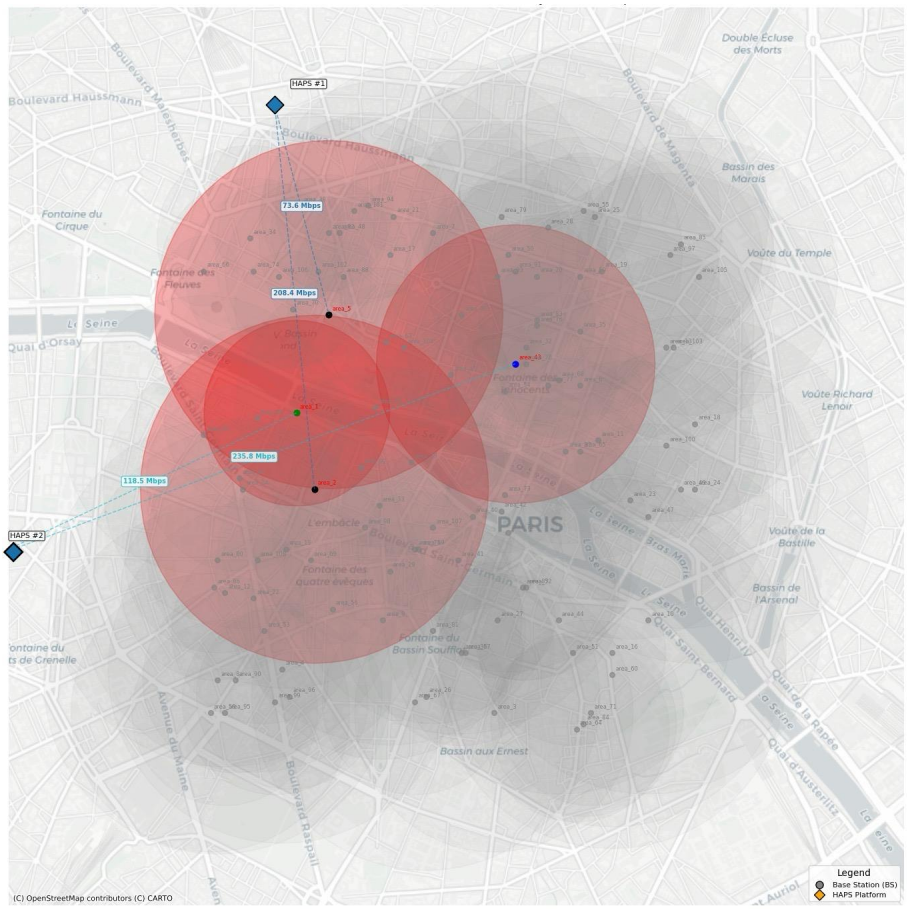

Figure 4.9: HAPS beam footprint over the blackout zones in the two-HAPS configuration.

4.5.1 Analysis by Area Priority

Areas labelled as critical (High-priority) tend to achieve coverage and throughput close to their configured weights. In the blackout scenarios considered in this chapter, almost all High-priority zones maintain near-complete service as long as the blackout footprint remains within the operating regime where the two-HAPS overlay has sufficient headroom. In all these experiments, the physical configuration is kept fixed to two HAPS, each equipped with eight beams (16 beams in total), and all beams are pointed at blackout-affected areas. Under this constraint, priority-aware admission policies ensure that hospitals and other critical infrastructures are consistently prioritised over residential or commercial zones, confirming that the proposed allocation strategy is effective in preserving resilience where it is most needed.

Non-critical areas absorb most of the residual imbalance when resources are scarce. As the system is progressively stressed (reduced capacity, increased traffic, larger blackout footprint up to the extreme case where all 108 areas are in blackout), the first visible effect is a sharp reduction in the throughput and coverage experienced by Low-priority zones, while High-priority demand remains largely protected until headroom is exhausted. This behaviour is consistent with the design objective of shielding High-priority areas at the expense of non-priority demand.

4.6 Stress Test

This section evaluates system robustness under degraded conditions, focusing on two key dimensions: reduced capacity (e.g., power loss, backhaul disruption) and unexpected traffic surges (e.g., emergencies, large events). The aim is to verify whether critical services are preserved and how the system degrades under pressure. In all stress tests, the same two-HAPS, 8-beam-per-HAPS configuration is retained, and all beams remain dedicated to the blackout footprint.

In the reference two-HAPS configuration considered here, the aggregate capacity delivered to the blackout areas under clear-sky conditions is on the order of a few hundred Mbit/s (around 0.45 Gbps). Under worst-day meteorology, the simulator reports only a

modest reduction (about 5–10%) in aggregate capacity. Moreover, at S-band, weather-induced attenuation is not the main limiting factor: network-level resilience is driven more by how capacity is distributed across beams and areas, by backhaul and power constraints, and by inter-HAPS coordination than by meteorological losses alone. For this reason, the stress tests focus on stronger perturbations (a 30% capacity reduction and a 2× traffic surge) to probe the limits of the overlay.

4.6.1 Capacity reduction

To emulate limited backhaul or onboard power constraints, the available capacity on the beams serving the blackout footprint is reduced by 30%. The number of HAPS is kept fixed to two, and the same set of blackout zones is served under reduced resources, still with 16 beams dedicated to blackout areas. In this regime, the aggregate traffic admitted by the overlay decreases markedly: in the considered snapshot, the total served traffic drops from roughly 0.45 Gbps in the reference case to about 0.31 Gbps under reduced-capacity conditions. The blocking of priority services increases noticeably, with the fraction of priority sessions that cannot be admitted rising from about 40% to almost 60%. Non-priority services, which are already largely rejected in the baseline configuration, become almost completely blocked when capacity is reduced.

These results confirm that the admission logic effectively protects critical traffic by sacrificing non-essential demand. As capacity shrinks, the system continues to allocate the available resources primarily to priority services, at the cost of further reducing the already limited support for non-priority flows. In other words, a substantial loss of capacity mainly translates into a higher blocking probability for priority sessions and the near-complete exclusion of non-priority services, rather than into an unstructured collapse of service across all classes.

4.6.2 Traffic increase

When traffic demand is doubled, the system rapidly becomes capacity-limited. The total amount of traffic that can be served remains almost unchanged (around 0.45 Gbps in both the baseline and 2×-demand cases), but the share of the offered load that is actually

admitted is roughly halved: the throughput ratio decreases from about 5.9% to 2.9%. In this scenario, the blocking probability of priority services grows significantly (from roughly 40% to more than 60%), while non-priority services are almost entirely discarded.

The traffic-surge test therefore highlights how explicit priority ordering shapes the degradation pattern under extreme load: as demand increases, the admission rule first removes support for non-priority flows and then progressively sheds a portion of the priority sessions once the overlay reaches its capacity limit. Critical services are not immune to degradation—their blocking probability does increase—but they remain the only class that is systematically admitted, whereas non-priority traffic absorbs virtually all of the additional stress.

4.7 Scaling with the Blackout Footprint

This section studies how system performance evolves as the *offered load per HAPS* increases under a fixed two-HAPS deployment with eight beams per HAPS, all beams being dedicated to blackout-affected areas. The number of HAPS is kept constant, and the stress is applied by progressively enlarging the blackout footprint, i.e., by increasing the number of affected zones that must be served by the same pair of platforms. This setting is equivalent to asking how far a two-HAPS overlay with about 0.85 Gbps of aggregate capacity can be stretched before losing its effectiveness.

4.7.1 Coverage percentage vs. blackout extent

As the number of blackout zones grows, the fraction of zones that can be fully served by the two HAPS decreases. For small blackout sets, the available capacity is sufficient to clear almost all demand, and the Best-HAPS association plus priority-aware admission can preserve service even under worst-day meteorology.

In the extreme case where *all* 108 areas in the Paris-like map are assumed to be in blackout simultaneously and all 16 beams are dedicated to this footprint, the two-HAPS overlay is clearly under-dimensioned if the goal is to guarantee full service everywhere: it can still sustain High-priority demand in a subset of areas, but several Low-priority zones

experience strong rate reductions or are not served at all. Beyond a certain blackout extent, adding more affected areas without increasing the number of platforms necessarily leads to larger uncovered demand.

In other words, for a fixed pair of HAPS, there exists a practical operating regime in which the overlay can effectively sustain a limited-size blackout cluster. Beyond this regime, the two-HAPS overlay must either drop non-priority flows entirely in several areas or accept reduced service levels even for some High-priority zones. This behaviour highlights the need to either deploy additional HAPS or relax QoS targets if the blackout footprint becomes too large.

4.8 Discussion of Results

The results highlighted in this chapter can be summarised as follows:

- **Priority resilience.** The combination of Best-HAPS association and priority-aware admission policies ensures that critical infrastructures retain service even under stress. Both in capacity-reduction and traffic-surge scenarios, High-priority areas maintain coverage and throughput close to their targets, while non-priority zones absorb most of the degradation.

- **HAPS cooperation.** Inter-HAPS coordination extends coverage and mitigates shadowed or poorly served areas, especially in the presence of clustered High-priority demand. In the two-HAPS configuration with roughly 0.8–0.9 Gbps of aggregate capacity and 16 beams dedicated to the blackout footprint, this cooperation allows a non-trivial set of blackout areas to be supported without losing Tier-A services.

- **Selection accuracy.** Naive nearest-based allocation is inadequate under realistic channel conditions. Score-based Best-HAPS association, which accounts for SINR, load, and backhaul factors, combined with explicit priority weights in the admission stage, captures fading and directivity effects and produces significantly better outcomes for priority traffic than purely distance-based rules.

- **Scalability.** For a fixed pair of HAPS, there exists a practical upper bound on

the blackout footprint that can be supported while preserving acceptable QoS for High-priority zones; beyond that point, either additional HAPS are required or QoS targets must be relaxed.

It is also useful to clarify the geometric scale of the simulated scenario. Each area in the Paris-like map corresponds to approximately 1 km^2, so the 108 active areas considered in this chapter cover roughly 108 km^2. This is much smaller than the theoretical footprint of a stratospheric HAPS, which can illuminate several thousands of square kilometres. The choice of a relatively small, dense urban scenario is deliberate: it allows us to work with a fine-grained traffic map while keeping the computational cost of the simulations manageable.

A similar remark applies to the beam layout. In the present model, each HAPS uses eight beams, and all of them are pointed at the same urban region and exclusively allocated to blackout-affected areas. In a real deployment, the same set of beams would typically be distributed across different towns, rather than being fully concentrated on a single city. As a result, the two-HAPS configuration studied here is conservative: the platforms are forced to spend all their beam resources on the blackout-affected area, and cannot offload part of the traffic to neighbouring regions. The scalability limits identified in this chapter should therefore be interpreted as a lower bound on what could be achieved with a larger footprint and a more flexible multi-region beam allocation.

Chapter 5

Applications to Civil Protection and Critical Infrastructures

5.1 Introduction

High Altitude Platform Stations (HAPS) combine long endurance, wide-area coverage, and relatively fast deployment. These features make them interesting for civil protection, continuity of service, and support to critical infrastructures.

Although the main part of this thesis focuses on urban blackout scenarios, the same association and scheduling mechanisms can be reused in other contexts. In particular, the *Best-HAPS* association rule and the weighted proportional-fair (PF) scheduler introduced in the previous chapters behave consistently under both fair-weather and worst-day meteorology, once the environmental loss L_{env} (rain/fog/snow) is included in the link budget. This is in line with current HAPS reference architectures and NTN standardisation work [25, 26].

5.2 Civil Applications

In civil scenarios, HAPS are mainly used as an additional layer on top of existing terrestrial networks, not as a replacement. The following subsections describe some example use

cases where the mechanisms studied in this thesis can be applied in practice.

5.2.1 Disaster Response and Public Safety

Natural disasters (earthquakes, floods, wildfires) can damage or overload terrestrial infrastructure. A HAPS deployed over the affected area can restore basic connectivity for first responders and the population, provide backhaul between surviving terrestrial "islands", and support emergency broadcast services [26].

The simulation results suggest that the *Best-HAPS* score $S_{u,h}$ helps to avoid purely SINR-based attachments that ignore load. Instead, it spreads traffic across platforms and takes into account both radio quality and backhaul headroom. The priority weights w_u allow emergency calls, e-health flows, and public-safety voice traffic to pre-empt non-essential services during peaks. Including L_{env} in the link budget also prevents optimistic assumptions on the available capacity at the cell edge in case of heavy rain or dense fog.

5.2.2 Critical Infrastructure Resilience

During large-scale power outages, HAPS can act as a communication lifeline for hospitals, emergency-control centres, and transport hubs. The weighted PF scheduler gives higher protection to critical flows and uses EWMA (Exponentially Weighted Moving Average) smoothing to avoid starving non-priority users completely. This type of behaviour is coherent with recent guidelines on acceptable risk and certification for HAPS-based services [27].

5.2.3 Rural and Remote Connectivity

In low-density areas, a small number of HAPS can complement fixed wireless access and LEO systems by providing a sort of "tower in the sky", with lower latency than satellites and flexible targeting of specific regions. In this case, the *Best-HAPS* association, with the backhaul headroom factor η_h^{BH}, helps to avoid overloading the "best-looking" platform and stabilises the attach success rate [26].

5.2.4 Temporary Events and Surge Capacity

For large events (concerts, sports events, fairs), one or two HAPS can be used to add temporary capacity above the venue. Time-varying priorities w_u (for example for safety channels and crowd notifications) and a weighted PF scheduler updated at each time slot can keep blocking probability and 95th-percentile latency within acceptable limits, even in the presence of strong diurnal variations.

5.2.5 Environmental Monitoring and Smart Cities

HAPS can also host payloads for environmental monitoring and smart-city services. Examples include EO/IR cameras for wildfire detection, hot-spot imaging of industrial areas, AIS/ADS-B relays, RF interference hunting, and IoT data aggregation. The multi-day or multi-week persistence already demonstrated by current platforms [28, 29] makes it possible to observe the same area for long periods. The same priority-aware resource management used in blackout scenarios can be reused here to ensure that high-value sensing and aggregation flows receive the necessary capacity.

5.3 Economic Considerations and Complementarity with LEO Constellations

From an economic point of view, stratospheric HAPS and LEO constellations sit on very different investment scales, even when they serve partially similar services. Large LEO broadband systems require multi-billion-dollar capital expenditure (CAPEX): SpaceX estimated the total cost of designing, building, and deploying *Starlink* at no less than $10 billion, while recent analyses place the investment for large LEO constellations in the $10–30 billion range, with annual operating costs of the order of $1–2 billion for replenishment and operations [30, 31, 32]. The European IRIS secure constellation is budgeted at about €10.5 billion for roughly 280 satellites across LEO/MEO [33]. The first generation of OneWeb raised on the order of $3.4 billion in equity, with individual satellites reported around $1 million each before launch costs [34, 35].

HAPS overlays have a much lower entry cost and can be scaled incrementally. For high-end solar-powered platforms such as the Airbus *Zephyr*, open-source estimates indicate a unit cost in the $10–20 million range per airframe [28, 36]. Market studies suggest annual operating costs around $1.2 million per platform for persistent communication services, including maintenance and ground-segment operations [26, 37]. At the other extreme, balloon-based systems such as Google Loon have much lower unit costs: a material and equipment audit estimated a CAPEX of about $17,870 per balloon and a five-year total cost of roughly $40,000 per unit [38, 39], while Google reported operating costs of "hundreds of dollars per day" per balloon for connectivity and operations [40]. These numbers are only indicative and depend strongly on design choices (payload, autonomy, regulation), but they show that HAPS can be fielded with much lower up-front investment than a global LEO constellation.

Qualitatively, a multi-HAPS overlay for a regional civil-protection mission (disaster-affected area, critical corridor, maritime region) might require a few platforms and tens of millions of dollars in CAPEX, plus single-digit millions per year in OPEX. A comparable solution based on LEO would normally use existing global constellations rather than a dedicated system, because building a custom LEO network for a single region would not be economically justified. This helps to explain why LEO constellations are optimised for long-term, global services, whereas HAPS are more suitable for regional and time-bounded missions [26, 5, 37].

HAPS should therefore not be seen as a one-to-one alternative to LEO. Instead, they fit better as a complementary layer:

- **LEO constellations** provide global coverage, multi-year continuity, and large capacity where demand and investment allow it.
- **HAPS overlays** offer a relatively fast option to create or densify coverage over specific regions of interest, including areas with limited commercial appeal or temporary relevance (disasters, critical corridors, remote communities).

In this view, a HAPS network looks more like a regional infrastructure investment, with costs comparable to a small fleet of specialised UAVs and a modest ground segment, while a LEO constellation resembles a strategic, multi-decade programme. The simulation results

in this thesis suggest that the proposed Best-HAPS association rule and priority-aware scheduling can turn a small number of HAPS into a useful overlay for civil protection and critical infrastructures, leaving global connectivity and long-term backhaul to existing or planned LEO systems.

5.4 From Simulation Results to Deployment Guidelines

Several practical indications emerge from the simulations.

First, the priority-aware scheduler is effective in preserving critical services. Setting $w_u > 1$ for emergency and high-priority flows allows them to keep acceptable throughput and delay even during traffic surges, while EWMA throughput averaging avoids excessive starvation of other users.

Second, the Best-HAPS association rule helps to avoid overloading a single platform and reduces handover storms. By combining SINR, current load, and backhaul headroom, it gives a simple rule that is still consistent with worst-day meteorology and can be implemented in a realistic controller.

Third, weather-aware planning, obtained by including L_{env} (gases, rain, clouds/fog) via ITU-R models, avoids overestimating the capacity available at the edge of the coverage area and helps to tune elevation masks, beam tilts, and launch windows.

The simulation campaign does not point to a single "optimal" deployment. However, it suggests a set of design rules that can help planners when dimensioning multi-HAPS overlays for civil protection and critical infrastructures.

5.5 Future Perspectives

Beyond the specific scenarios studied here, multi-HAPS networks are a candidate component for future 6G/NTN architectures. The same priority-aware association and scheduling mechanisms can, in principle, be extended to richer radio interfaces, denser constellations, and tighter integration with satellite and terrestrial layers, within civil-protection and critical-infrastructure frameworks.

Chapter 6 outlines a more detailed research agenda, including improved channel models, mobility- and energy-aware control, advanced inter-HAPS cooperation, and learning-based resource allocation. Overall, multi-HAPS overlays could evolve from a simulation concept into an operational tool, provided that technological progress is matched by progress in governance, certification, and spectrum coordination.

Chapter 6

Conclusions and Future Work

6.1 Interpretation of the Results

The simulations carried out in this thesis show that cooperative HAPS deployments, combined with priority-aware scheduling, can improve the resilience of communication services during urban blackouts. Tier-A infrastructures (for example hospitals and emergency services) are consistently prioritised, even when capacity is reduced or when sudden traffic surges occur.

The results also highlight that simple nearest-attachment policies are not adequate in realistic urban conditions. Load- and backhaul-aware Best-HAPS association and SINR-/capacity-aware strategies, when coupled with priority weights, lead to measurable gains in attach success, blocking probability, and tail latency. The inclusion of realistic propagation components (path loss, shadowing due to clutter, antenna directivity) and weather-related losses L_{env} supports the conclusion that the proposed allocation stack remains effective under both fair-weather and worst-day meteorology.

Overall, the case study confirms that multi-HAPS overlays can provide a meaningful "resilience layer" on top of terrestrial networks. Under blackout conditions and worst-day meteorology, the proposed stack keeps Tier-A blocking probabilities markedly lower than those of non-priority traffic, and it maintains service continuity in a subset of Tier-B areas, whereas nearest-attachment baselines either overload a few platforms or leave entire blackout zones underserved. At the same time, the use of weighted scheduling and

admission control allows the system to exploit residual capacity for non-priority services whenever available, rather than dedicating all resources exclusively to emergency traffic.

A further result concerns the role of learning-based control. The single-beam Q-learning admission strategy shows that it is possible to adapt resource allocation to local demand and capacity conditions without an explicit model of the traffic dynamics. In the considered setting, Q-learning reduces the blocking of priority services and improves beam utilisation compared with static rule-based policies, while preserving the desired ordering between traffic classes. This suggests that reinforcement learning can be a viable complement to rule-based and optimisation-based schemes in future multi-HAPS architectures, especially in scenarios characterised by high uncertainty and rapidly varying traffic.

6.2 Future Work

The present model has several limitations, which naturally suggest directions for future development. A first step concerns the refinement of the channel model: incorporating rain fading, turbulence, Doppler shifts, and time-correlated small-scale fading would enable the simulator to capture fast temporal dynamics and impairments that become especially relevant at higher frequencies. Another important extension involves a more explicit modelling of user mobility and of the spatial and temporal variability of traffic. Allowing demand to evolve at sub-cell granularity would make it possible to reproduce realistic usage patterns and handover behaviour, especially in dense urban environments.

Energy-awareness also remains an open aspect. Integrating solar harvesting, battery degradation models, and adaptive beam management would link platform endurance to capacity and coverage decisions, enabling the study of long-duration missions. In parallel, the framework could benefit from more advanced forms of inter-HAPS cooperation, such as multi-hop relaying, cooperative beamforming, and dynamic role switching, which would exploit spatial diversity at constellation level.

The use of learning-based control mechanisms is another promising direction. While this thesis explored a single-beam Q-learning admission strategy, extending the approach to

multi-agent reinforcement learning or model-based predictive control could allow dynamic resource adaptation under uncertainty. Security and resilience also require dedicated study: incorporating explicit anti-jam logic, strengthening the control plane, and defining controlled degradation modes would make the system more robust against intentional interference.

Finally, moving towards field validation is essential. Small-scale pilots, hardware-in-the-loop experiments, and the use of real operational traffic would help bridge the gap between simulation and deployment, providing quantitative calibration of the models and confirming the feasibility of multi-HAPS overlays for resilient communications.

Appendix A

Source Code (Extracts)

This appendix reports selected excerpts of the Python source code used in the multi-HAPS simulation framework. The full implementation is available in the accompanying repository and includes, among others:

- channel and propagation models;
- antenna and MIMO configuration modules;
- geometry and mapping utilities ;
- high-level routines for HAPS selection and capacity allocation.

A.1 HAPS selection routine (full version)

```
def select_best_haps_for_area(self,
                              area: dict,
                              haps_positions: list,
                              haps_altitude: float,
                              selection_metric: str = "capacity") ->
    tuple:
    """
␣␣␣␣Select␣the␣best␣HAPS␣for␣a␣given␣area.

␣␣␣␣Returns
```

```
    -------
    tuple
        (best_haps_coord, best_index, best_score, all_scores)

        all_scores contains, for each candidate:
        - capacity (bps)
        - capacity_Mbps (Mbps)
        - capacity_with_weather (bps)
        - snr (dB)
        - PL_dB, G_dBi, etc.
    """

    EPS = 1e-18

    # --- System parameters (raw read) ---
    B = float(getattr(self, "bandwidth", 0.0))  # may be in MHz in the config
    Ptx_dBm = float(getattr(self, "P_tx", 0.0))
    use_weather = bool(getattr(self, "selection_use_weather", True) or
                       area.get("selection_use_weather", True))
    alpha = float(getattr(self, "selection_alpha", 0.6) if selection_metric == "hybrid" else 1.0)

    # --- Normalize bandwidth: ensure it is in Hz ---
    # If B is given in MHz (e.g., 20), convert to Hz.
    # Heuristic: if < 1e5, interpret as MHz.
    if B <= 0:
        raise ValueError(f"Bandwidth must be > 0. Got {B}.")
    if B < 1e5:
        B *= 1e6  # from MHz to Hz

    # --- Noise density N0 in dBm/Hz (robust to config errors) ---
    # Priority: self.noise_power_density (expected already in dBm/Hz)
    if hasattr(self, "noise_power_density"):
        N0_dBm_per_Hz = float(getattr(self, "noise_power_density"))
    else:
        # Fallback: if P_n exists, it may be total noise in dBm
        if hasattr(self, "P_n"):
```

```
46              Pn_val = float(getattr(self, "P_n"))
47              # If it looks like a "total" noise (-130..-70 dBm), convert
    to density
48              if -130.0 < Pn_val < -70.0:
49                  N0_dBm_per_Hz = Pn_val - 10.0 * np.log10(B)
50              else:
51                  # Otherwise, use standard physical constant
52                  N0_dBm_per_Hz = -174.0
53          else:
54              N0_dBm_per_Hz = -174.0
55
56      # If a non-density value (> -160 dBm/Hz) was provided by mistake,
57      # interpret it as total noise over B and convert to density.
58      if N0_dBm_per_Hz > -160.0:
59          N0_dBm_per_Hz = N0_dBm_per_Hz - 10.0 * np.log10(B)
60
61      # --- Power and noise conversions ---
62      Ptx_mW = 10.0 ** (Ptx_dBm / 10.0)
63      N0_mW_per_Hz = 10.0 ** (N0_dBm_per_Hz / 10.0)
64
65      # Per-HAPS loads (shared bandwidth)
66      H = len(haps_positions)
67      haps_loads = list(area.get("haps_loads", [0] * H))
68
69      centroid = area["area_centroid"]
70
71      # Resolve channel API name (typo-safe)
72      rsrp_fn = None
73      for cand in ["RSRP_cluster_parameters", "RSPR_cluster_parameters",
    "RSPR_cluster_paramters"]:
74          if hasattr(self, cand):
75              rsrp_fn = getattr(self, cand)
76              break
77      if rsrp_fn is None:
78          raise AttributeError("Missing method: RSRP_cluster_parameters")
79
80      candidates = []
81
```

```
    # ============================================================
    # STEP 1: compute metrics for all candidates
    # ============================================================
    for i, hp in enumerate(haps_positions):
        if len(hp) == 2:
            haps_coord = (hp[0], hp[1], haps_altitude)
        else:
            haps_coord = (hp[0], hp[1], hp[2])

        try:
            # returns (fading, PL_no_weather, PL_with_weather, gain)
            fading, path_loss_dB, path_loss_w_dB, gain_dBi = rsrp_fn(
area, haps_coord)
        except Exception as e:
            candidates.append({
                "index": i,
                "valid": False,
                "reason": str(e),
                "haps_coord": haps_coord
            })
            continue

        PL_no_weather = float(path_loss_dB)
        PL_with_weather = float(path_loss_w_dB)
        G_dBi = float(gain_dBi)

        # Path loss used for selection
        PL_dB = PL_with_weather if use_weather else PL_no_weather

        # Channel vector (at least 1D), squared norm
        fading = np.atleast_1d(fading).astype(np.complex128).reshape
(-1, 1)

        # Effective bandwidth after sharing
        n_after = int(haps_loads[i]) + 1
        B_eff = max(B / n_after, 1.0)
        N_mW = N0_mW_per_Hz * B_eff
        N_dBm = 10.0 * np.log10(max(N_mW, EPS))
```

```
118
119         # ---- Capacity NO WEATHER ----
120         amp_lin_no_w = 10.0 ** ((G_dBi - PL_no_weather) / 10.0)
121         h_no_w = (amp_lin_no_w ** 0.5) * fading
122         h_norm2_no_w = float(np.vdot(h_no_w, h_no_w).real) + EPS
123         signal_mW_no_w = h_norm2_no_w * Ptx_mW
124         sinr_lin_no_w = signal_mW_no_w / (N_mW + EPS)
125         capacity_no_weather = B_eff * np.log2(1.0 + sinr_lin_no_w)
126
127         # ---- Capacity WITH WEATHER ----
128         amp_lin_w = 10.0 ** ((G_dBi - PL_with_weather) / 10.0)
129         h_w = (amp_lin_w ** 0.5) * fading
130         h_norm2_w = float(np.vdot(h_w, h_w).real) + EPS
131         signal_mW_w = h_norm2_w * Ptx_mW
132         sinr_lin_w = signal_mW_w / (N_mW + EPS)
133         capacity_with_weather = B_eff * np.log2(1.0 + sinr_lin_w)
134
135         # SNR for selection (coherent with chosen PL_dB)
136         amp_lin_sel = 10.0 ** ((G_dBi - PL_dB) / 10.0)
137         h_sel = (amp_lin_sel ** 0.5) * fading
138         h_norm2_sel = float(np.vdot(h_sel, h_sel).real) + EPS
139         signal_mW_sel = h_norm2_sel * Ptx_mW
140         sinr_lin_sel = signal_mW_sel / (N_mW + EPS)
141         snr_dB = 10.0 * np.log10(max(sinr_lin_sel, EPS))
142
143         # Optional debug on first candidate
144         if i == 0 and getattr(self, "verbose_selection_debug", False):
145             print(f"[DEBUG] B_eff={B_eff:.1f} Hz | N0={N0_dBm_per_Hz:.1
    f} dBm/Hz | "
146                   f"N={N_dBm:.1f} dBm | SNR_sel={snr_dB:.2f} dB")
147
148         candidates.append({
149             "index": i,
150             "valid": True,
151             "haps_coord": haps_coord,
152             "capacity_bps": capacity_no_weather,                # NO
    weather
153             "capacity_with_weather_bps": capacity_with_weather,  # WITH
```

```
    weather
154             "snr_dB": snr_dB,
155             "B_eff_Hz": B_eff,
156             "noise_mW": N_mW,
157             "noise_dBm": N_dBm,
158             "PL_dB": PL_dB,
159             "PL_no_weather": PL_no_weather,
160             "PL_with_weather": PL_with_weather,
161             "G_dBi": G_dBi,
162             "n_after": n_after
163         })
164 
165     # ============================================================
166     # STEP 2: select best candidate
167     # ============================================================
168     valid = [c for c in candidates if c["valid"]]
169     if valid:
170         if selection_metric == "capacity":
171             for c in valid:
172                 c["score"] = c["capacity_bps"]
173         elif selection_metric == "snr":
174             for c in valid:
175                 c["score"] = c["snr_dB"]
176         elif selection_metric == "hybrid":
177             cap_vals = np.array([c["capacity_bps"] for c in valid],
    dtype=float)
178             snr_vals = np.array([c["snr_dB"] for c in valid], dtype=
    float)
179             cap_min, cap_max = float(cap_vals.min()), float(cap_vals.
    max())
180             snr_min, snr_max = float(snr_vals.min()), float(snr_vals.
    max())
181             for c in valid:
182                 cap_n = 0.0 if cap_max <= cap_min + EPS else (c["
    capacity_bps"] - cap_min) / (cap_max - cap_min)
183                 snr_n = 0.0 if snr_max <= snr_min + EPS else (c["snr_dB
    "] - snr_min) / (snr_max - snr_min)
184                 c["score"] = alpha * cap_n + (1.0 - alpha) * snr_n
```

```
185         else:
186             raise ValueError(f"Unknown selection_metric: {
    selection_metric}")
187 
188         def _key(c):
189             x, y, _ = c["haps_coord"]
190             dx, dy = x - centroid[0], y - centroid[1]
191             dist_xy = (dx*dx + dy*dy) ** 0.5
192             return (-float(c["score"]), float(c["PL_dB"]), float(
    dist_xy))
193 
194         best = sorted(valid, key=_key)[0]
195         best_index = best["index"]
196         best_haps = best["haps_coord"]
197         best_score = float(best["score"])
198 
199         all_scores = []
200         for c in candidates:
201             if not c["valid"]:
202                 all_scores.append({
203                     "index": c["index"],
204                     "score": None,
205                     "reason": c.get("reason", "invalid")
206                 })
207             else:
208                 all_scores.append({
209                     "index": c["index"],
210                     "score": c.get("score"),
211                     "capacity": c["capacity_bps"],
212                     "capacity_Mbps": c["capacity_bps"] / 1e6,
213                     "capacity_with_weather": c["
    capacity_with_weather_bps"],
214                     "capacity_with_weather_Mbps": c["
    capacity_with_weather_bps"] / 1e6,
215                     "snr": c["snr_dB"],
216                     "B_eff_Hz": c["B_eff_Hz"],
217                     "n_after": c["n_after"],
218                     "noise_mW": c["noise_mW"],
```

```
                    "noise_dBm": c["noise_dBm"],
                    "PL_dB": c["PL_dB"],
                    "PL_no_weather": c["PL_no_weather"],
                    "PL_with_weather": c["PL_with_weather"],
                    "G_dBi": c["G_dBi"]
                })

        return best_haps, best_index, best_score, all_scores

    # ============================================================
    # FALLBACK: Nearest HAPS
    # ============================================================
    print(f"[FALLBACK] No valid signal for area {area.get('area_id','?')}    using nearest")

    min_dist = float("inf")
    nearest_haps = None
    nearest_index = None

    for i, hp in enumerate(haps_positions):
        x, y = hp[0], hp[1]
        dx, dy = x - centroid[0], y - centroid[1]
        dist = (dx*dx + dy*dy) ** 0.5
        if dist < min_dist:
            min_dist = dist
            nearest_haps = (x, y, haps_altitude)
            nearest_index = i

    n_after = int(haps_loads[nearest_index]) + 1
    B_eff = max(B / n_after, 1.0)
    N_mW = N0_mW_per_Hz * B_eff
    N_dBm = 10.0 * np.log10(max(N_mW, EPS))

    try:
        fading, path_loss_dB, path_loss_w_dB, gain_dBi = rsrp_fn(area, nearest_haps)
        PL_no_w = float(path_loss_dB)
        PL_w = float(path_loss_w_dB)
```

```
255         G_dBi = float(gain_dBi)
256         fading = np.atleast_1d(fading).astype(np.complex128).reshape
    (-1, 1)
257
258         # NO weather
259         amp_no_w = 10.0 ** ((G_dBi - PL_no_w) / 10.0)
260         h_no_w = (amp_no_w ** 0.5) * fading
261         h_norm2_no_w = float(np.vdot(h_no_w, h_no_w).real) + EPS
262         signal_no_w = h_norm2_no_w * Ptx_mW
263         sinr_no_w = signal_no_w / (N_mW + EPS)
264         capacity_bps = B_eff * np.log2(1.0 + sinr_no_w)
265
266         # WITH weather
267         amp_w = 10.0 ** ((G_dBi - PL_w) / 10.0)
268         h_w = (amp_w ** 0.5) * fading
269         h_norm2_w = float(np.vdot(h_w, h_w).real) + EPS
270         signal_w = h_norm2_w * Ptx_mW
271         sinr_w = signal_w / (N_mW + EPS)
272         capacity_with_weather_bps = B_eff * np.log2(1.0 + sinr_w)
273         snr_dB = 10.0 * np.log10(max(sinr_w, EPS))
274
275     except Exception:
276         PL_no_w = float(getattr(self, "fallback_pathloss_dB", 150.0))
277         PL_w = PL_no_w
278         G_dBi = float(getattr(self, "fallback_gain_dBi", 0.0))
279         h_norm2 = 10.0 ** ((G_dBi - PL_no_w) / 10.0)
280         signal_mW = h_norm2 * Ptx_mW
281         sinr_lin = signal_mW / (N_mW + EPS)
282         capacity_bps = B_eff * np.log2(1.0 + sinr_lin)
283         capacity_with_weather_bps = capacity_bps * 0.9
284         snr_dB = 10.0 * np.log10(max(sinr_lin, EPS))
285
286     all_scores = [{
287         "index": nearest_index,
288         "score": None,
289         "capacity": capacity_bps,
290         "capacity_Mbps": capacity_bps / 1e6,
291         "capacity_with_weather": capacity_with_weather_bps,
```

```
292         "capacity_with_weather_Mbps": capacity_with_weather_bps / 1e6,
293         "snr": snr_dB,
294         "B_eff_Hz": B_eff,
295         "n_after": n_after,
296         "noise_mW": N_mW,
297         "noise_dBm": N_dBm,
298         "PL_dB": PL_w if use_weather else PL_no_w,
299         "PL_no_weather": PL_no_w,
300         "PL_with_weather": PL_w,
301         "G_dBi": G_dBi,
302         "fallback": True,
303         "reason": "nearest_assignment"
304     }]
305 
306     return nearest_haps, nearest_index, -1.0, all_scores
```

Codice A.1: Best-HAPS selection routine used in the simulator

Appendix B

Jupyter Notebooks and Simulation Scripts

This appendix describes the structure of the main Jupyter notebooks used to run the simulations and generate the figures and tables presented in Chapter 4. The most relevant notebooks are:

- `SIMULATOR_PARIS.ipynb`: end-to-end multi-HAPS simulation;
- `Traces.ipynb`: traffic preprocessing and exploratory analysis.

B.1 `SIMULATOR_PARIS.ipynb`

This notebook implements the full simulation loop used to obtain the main results:

1. **Input and preprocessing**
    - load the traffic tensor `big_matrix_MB_bs1_108_service1_68_time1_672.npy`;
    - load area geometry and classification (`BS_108classified_without_services.csv`);
    - load HAPS configuration and beam parameters (CSV files used in Chapter 3).
2. **Time-slot-based simulation loop**
    - for each 15-minute slot, update traffic demand per area;

- compute link budget and SINR for all visible HAPS–area pairs;
- apply the allocation policy (priority-based, knapsack-greedy, Q-learning, etc.);
- record assigned capacity, blocked traffic, and coverage indicators.

3. **Output generation**

- generate throughput and coverage plots for each strategy;
- compute summary KPIs (served traffic, priority satisfaction, per-area statistics);
- export CSV files for further analysis and figure generation.

The figures showing coverage percentages, served traffic by strategy, and capacity histograms in Chapter 4 are directly produced or post-processed starting from this notebook.

B.2 `Traces.ipynb`

The `Traces.ipynb` notebook focuses on the Orange traffic dataset:

- parsing and cleaning of raw CSV/NetMob traces;
- construction of the weekly traffic tensor $X \in \mathbb{R}^{N_{BS} \times N_S \times N_T}$;
- unit conversion from bytes to MB and scaling to match typical LTE cell loads;
- exploratory plots (e.g., top services by volume, diurnal patterns, spatial heterogeneity across the 108 areas).

These analyses support the dataset description reported in Chapter 3 and motivate the choice of the scaling factor and the service-priority mapping.

Bibliography

[1] P. Lobner, "Cira high altitude platform systems (haps) — modern airships," tech. rep., Lyncean Group, February 2024. Overview of CIRA HHAA design, mission profile, and EuroHAPS context.

[2] Reuters, "Portugal's ren says no sign blackout caused by cyberattack," 4 2025.

[3] T. Guardian, "Electricity restored to most of portugal and 90% of spain after massive power outage," 4 2025.

[4] ENTSO-E, "28 april 2025 iberian blackout," 4 2025. Factual description of the system event and restoration.

[5] International Telecommunication Union, "Haps – high-altitude platform systems," 2023. Definition of HAPS (20–50 km altitude range).

[6] "Study on new radio (nr) to support non-terrestrial networks," Technical Report TR 38.811, 3rd Generation Partnership Project (3GPP), June 2019. Release 15 study on NTN scenarios, incl. HAPS.

[7] E. Björnson, J. Hoydis, and L. Sanguinetti, "Massive mimo networks: Spectral, energy, and hardware efficiency," *Foundations and Trends in Signal Processing*, vol. 11, no. 3-4, pp. 154–655, 2017.

[8] A. Delgado, D. Domínguez, J. Gonzalo, and A. Escapa, "Station-keeping haps mission through optimal sprint and drift trajectories," *Aerospace Science and Technology*, vol. 152, p. 109365, 2024.

[9] S. C. Arum, D. Grace, P. D. Mitchell, M. D. Zakaria, and N. Morozs, "Energy management of solar-powered aircraft-based high altitude platform for wireless communications," *Electronics*, vol. 9, no. 1, p. 179, 2020.

[10] M. Weiner, "High-altitude platforms: Flying beyond the tropopause." `https://aeroreport.de/en/innovation/high-altitude-platforms-flying-beyond-the-tropopause`, 2024. Accessed Oct. 2025.

[11] Electronic Communications Committee, "Ecc report 156: Conditions for possible co-existence between haps gateway links and other services/systems in the 5850–7075 mhz band," tech. rep., CEPT ECC, Cardiff, UK, 2011.

[12] J. G. Andrews, A. K. Gupta, and H. S. Dhillon, "A primer on cellular network analysis using stochastic geometry," *arXiv preprint arXiv:1604.03183*, 2016.

[13] J. Xu, L. Duan, and R. Zhang, "User association in massive mimo heterogeneous networks," *arXiv preprint arXiv:1501.03407*, 2015.

[14] X. Lin, S. Rommer, S. Euler, E. A. Yavuz, and R. S. Karlsson, "5g from space: An overview of 3gpp non-terrestrial networks," *IEEE Communications Standards Magazine*, 2021. preprint version available on arXiv:2103.09156.

[15] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Non-terrestrial networks in the 6g era: Challenges and opportunities," *IEEE Communications Magazine*, vol. 59, no. 11, pp. 95–101, 2021.

[16] P. K. Singya, P. Bansal, S. N. Merchant, and U. B. Desai, "Mixed fso/rf based multiple haps assisted multiuser relaying system," *Frontiers in Communications and Networks*, vol. 3, p. 746201, 2022.

[17] K. O. Odeyemi and P. A. Owolawi, "A mixed fso/rf integrated satellite–high altitude platform relaying networks for multiple terrestrial users with presence of eavesdropper: A secrecy performance," *Photonics*, vol. 9, no. 1, p. 32, 2022.

[18] Y. J. Hyun *et al.*, "Multidirectional all-optical precompensation and forward relay for haps-assisted free-space optical communications," *Optics & Laser Technology*, 2025. In press.

[19] V. Yajnanarayana *et al.*, "Priority-aware scheduling in non-terrestrial networks: Ensuring qos for critical services," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 10, pp. 10577–10589, 2021.

[20] X. Dai and X. Chen, "Priority-aware task offloading and resource allocation in satellite and hap assisted edge-cloud collaborative networks," in *Proc. 2023 15th International Conference on Communication Software and Networks (ICCSN)*, 2023. IEEE.

[21] C. . ECC, "Working document on haps / radiation mask (resolution 221)," tech. rep., CEPT, —. Working document for radiation mask in WRC/HAPS studies.

[22] "Attenuation by atmospheric gases," Tech. Rep. Recommendation P.676-13, ITU-R, Geneva, Switzerland, 2023.

[23] "Specific attenuation model for rain for use in prediction methods," Tech. Rep. Recommendation P.838-3, ITU-R, Geneva, Switzerland, 2005.

[24] "Attenuation due to clouds and fog," Tech. Rep. Recommendation P.840-8, ITU-R, Geneva, Switzerland, 2019.

[25] "Release 17 highlights." https://www.3gpp.org/specifications-technologies/releases/release-17, 2024. Accessed Oct. 4, 2025.

[26] HAPS Alliance, "Haps reference architecture series: Haps advantages in an era of satellite connectivity." https://hapsalliance.org/wp-content/uploads/formidable/12/2025_HAPSAlliance_Reference_Architecture_Advantages_Satellite_Connectivity_TWG_Whitepaper.pdf, 2025. Accessed Oct. 4, 2025.

[27] HAPS Alliance, "Haps certification pathways." https://hapsalliance.org/publications/, 2024. HAPS Alliance Aviation Working Group white paper, accessed Oct. 4, 2025.

[28] Airbus, "Zephyr high altitude platform station (haps)." https://www.airbus.com/en/products-services/defence/uas/zephyr, 2024. Accessed Oct. 4, 2025.

[29] Amprius, "Aalto zephyr achieves world-record 67-day flight powered by amprius ultra-high-energy batteries." https://ir.amprius.com/news-events/press-releases/detail/133/aalto-zephyr-achieves-world-record-67-day-flight-powered-by-amprius-ultra-high- 2025. Press release, May 21, 2025. Accessed Oct. 4, 2025.

[30] SpaceX, "Spacex's starlink: Frequently asked questions." https://www.starlink.com/faq, 2023. Includes official estimates of Starlink network deployment cost. Accessed Nov. 2025.

[31] McKinsey & Company, "The economics of satellite broadband: How constellations compete." https://www.mckinsey.com/industries/aerospace-and-defense/our-insights/the-economics-of-satellite-broadband, 2023. Accessed Nov. 2025.

[32] Euroconsult, "Prospects for the small satellite market 2023–2032." https://digital-platform.euroconsult-ec.com/reports/prospects-for-the-small-satellite-market-2023/, 2023. Industry report on CAPEX/OPEX trends for LEO constellations. Accessed Nov. 2025.

[33] European Space Agency (ESA), "Iris^2: Secure satellite constellation for europe." https://www.esa.int/Applications/Connectivity_and_Secure_Communications/IRIS2, 2024. ESA programme overview. Accessed Nov. 2025.

[34] TechCrunch, "Oneweb raises $1.25 billion for global satellite broadband network." https://techcrunch.com/2019/03/18/oneweb-raises-1-25-billion-for-satellite-broadband/, 2019. Accessed Nov. 2025.

[35] European Space Agency (ESA), "Oneweb satellites: A joint esa and airbus development for leo broadband." https://artes.esa.int/projects/oneweb-satellites, 2020. Accessed Nov. 2025.

[36] L. Weaver, "Between air and space – zephyr and the future of high altitude pseudo-satellites within defence," *Air & Space Power Review*, vol. 22, no. 2, pp. 58–74, 2019.

[37] SoftBank Corp., "High altitude platform station (haps) – stable telecommunication for all people and things." https://www.softbank.jp/en/corp/philosophy/technology/special/ntn-solution/haps/. Accessed Nov. 2025.

[38] M. Aslam, S. Y. Shin, and M. Bennis, "Google loon: A review on stratospheric communications, challenges, and future directions," *IEEE Access*, vol. 8, pp. 118818–118830, 2020.

[39] D. Burrus, "Inside project loon: The future of balloon-powered internet." https://www.burrus.com/2017/07/project-loon-the-future-of-balloon-powered-internet/, 2017. Accessed Nov. 2025.

[40] Google X, "What we've learned from launching internet balloons." https://x.company/blog/posts/loon-lessons-learned/, 2020. Accessed Nov. 2025.

[41] "Study on channel model for frequencies from 0.5 to 100 GHz," Tech. Rep. TR 38.901, 3rd Generation Partnership Project (3GPP), 2017. Version 14.0.0.

[42] M. Takahashi, Y. Kawamoto, N. Kato, A. Miura, and M. Toyoshima, "Adaptive power resource allocation with multi-beam directivity control in high-throughput satellite communication systems," *IEEE Wireless Communications Letters*, vol. 8, no. 4, pp. 1248–1251, 2019.

[43] M. Alzenad, M. Z. Shakir, H. Yanikomeroglu, and M.-S. Alouini, "Fso-based vertical backhaul/fronthaul framework for 5g+ wireless networks," in *Proc. IEEE ICC*, pp. 1–6, 2018. Discusses HAPS as aerial backhaul nodes.

[44] T. Qiu, L. Zhao, F. Xia, M. Chen, and S. Yu, "Haps-based emergency communication systems: Capacity, coverage and qos analysis," *IEEE Network*, vol. 34, no. 6, pp. 212–219, 2020.

[45] BBC News, "Spanish blackout: Millions without power across iberia," 2025. Accessed May 2025.

[46] Thales Group, "Space q&a: All about stratobus," 2021. Accessed May 2025.

[47] Thales Alenia Space, "Stratobus factsheet." via Lynceans.org, 2021. Accessed May 2025.

[48] Futura Sciences, "Stratobus: bientôt les premiers vols de démonstration," 2017. Accessed May 2025.

[49] Unmanned Systems Technology, "Solar array validated for stratobus," 2018. Accessed May 2025.

[50] European Space Agency (ESA), "High-altitude aerial platforms," 2017. Accessed May 2025.

[51] SatNews, "Thales alenia space contract received from european commission for eurohaps project," 2023. Accessed May 2025.

[52] Wikipedia contributors, "Stratobus — Wikipedia, the free encyclopedia," 2025. Accessed May 2025.

[53] "Non-terrestrial networks (ntn) overview." `https://www.3gpp.org/technologies/ntn-overview`, 2024. Accessed Oct. 4, 2025.

[54] S. O. Soler and contributors, "Itu-rpy: Python implementation of itu-r propagation models." `https://itu-rpy.readthedocs.io/_/downloads/en/latest/pdf/`, 2023. Accessed Oct. 4, 2025.

[55] ITU, "Attenuation factors for thz earth-space links (incl. p.838-3, p.840, p.676 references)." `https://www.itu.int/dms_pub/itu-s/opb/jnl/S-JNL-VOL5.ISSUE2-2024-A19-PDF-E.pdf`, 2024. International Telecommunication Union Journal, 2024. Accessed Oct. 4, 2025.

[56] International Telecommunication Union (ITU), “Preparing for wrc-19: Background paper,” tech. rep., ITU, 2019.

[57] International Telecommunication Union (ITU), “Wrc-23 closing press release.” ITU Press Release, 2023.

[58] H. Alliance, “Advancing global connectivity: The expanded horizon of the haps spectrum.” HAPS Alliance blog, —.

[59] I. T. U. (ITU), “Resolution 122 (rev. wrc-19) – use of the frequency bands 47.2–47.5 ghz and 47.9–48.2 ghz by haps,” tech. rep., ITU, 2019.

[60] “Telecom security incidents 2024,” tech. rep., European Union Agency for Cybersecurity (ENISA), 7 2025. Annual summary of incidents reported by national authorities (covering year 2024).

[61] “Statement from gen. glen vanherck, commander, norad/usnorthcom,” 2023.

[62] “Statement by secretary of defense on high-altitude object (4 feb 2023),” 2023.

[63] “U.s. downs 3 unidentified flying objects nearly a week after balloon,” 2023.

[64] Singapore MINDEF, “Aster 30 missile system — technical specifications,” 2013.

[65] MBDA, “Aster family / samp/t ng,” 2025.

[66] Missile Defense Advocacy Alliance, “Samp/t air defense system (france & italy),” 2024.

[67] Reuters, “Germany’s diehl to further expand production of iris-t air defence systems,” 2024.

[68] Diehl Defence, “Denmark procures iris-t slm (essi interim solution),” 2025.

[69] Diehl Defence, “armasuisse switzerland procures iris-t slm,” 2025.

[70] Raytheon, “Raytheon completes first amraam-er flight tests from nasams,” 2016.

[71] KONGSBERG, “Nasams air defence system (customer base),” 2024.

[72] Raytheon, “Patriot guidance enhanced missile (gem-t),” 2025.

[73] DOT&E, U.S. DoD, “Patriot system assessment (fy2012 annual report),” 2012.

[74] CSIS Missile Threat, “Patriot — operators and overview,” 2023.

[75] CSIS Missile Threat, “S-400 triumf,” 2021.

[76] “S-400 missile system,” 2025. Dati altimetrici delle varianti, consultati 2025-11-01.

[77] “Hq-9 surface-to-air missile system,” 2025. Operatori/ceiling variante HQ-9B, consultati 2025-11-01.

[78] Israel Aerospace Industries, “Barak er interceptor — 150 km range / 30 km altitude (test 2021),” 2021.

[79] “Barak 8,” 2025. Dati prova ER 2021, consultati 2025-11-01.

[80] “Aster (missile family) — operators,” 2025. Elenco operatori incl. Singapore e Ucraina, consultati 2025-11-01.

[81] 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.

[82] M. A. Mbarek, “HAPS-MIMO: MIMO simulation tools for high-altitude platform stations.” `https://github.com/aminmbare/HAPS-MIMO`. Accessed: Nov. 2025.

Acknowledgments

I would like to thank my supervisor, Prof. Michela Meo, for her guidance, trust, and feedback throughout this work. I am also very grateful to my co-supervisor, Eng. Greta Vallero, Ph.D., for her daily support, attention to detail, and constant availability.

I also thank the researchers at CIRA for their collaboration and for the resources shared during the development of this thesis, as well as Mohamed Amine Mbarek and Maoquan Ni for their contributions to this field.

Finally, I wish to thank my family and friends for their support and patience. I am especially grateful to my parents, to my sister Veronica and to my brother Daniele, who have been role models to me and have always given me comfort, even in the most difficult moments of my studies. I warmly thank my grandparents, Mario and Carla, and all my uncles and aunts.

A special thought goes to Martina, who, like my family, has always supported and endured me, and to Mauro, Cristina, and Sofia. To all my friends, especially Simone, Luca, and Fabio: without you, I would not have been able to reach this goal.