

Fixed Transport Network Cost Analysis under 6G Traffic Scenarios and Technologies

Marc Ruiz^{1*}, Miguel Ángel Carnero², Adrián Maeso², and Luis Velasco¹

¹ *Optical Communications Group - Universitat Politècnica de Catalunya, Barcelona, Spain*

² *Telefónica GCTO, Madrid, Spain*

e-mail: marc.ruiz-ramirez@upc.edu

Abstract— This paper addresses the challenges of techno-economic evaluation of transport networks toward 6G and autonomous operation. It presents a methodology used to generate an open reference traffic dataset combining real 4G/5G measurements with forecasted 6G traffic, leveraging both human operator expertise and large language models. The dataset aims to mitigate the lack of realistic data for next-generation network planning studies. Using this dataset, the paper evaluates the techno-economic impact of the DESIRE6G control and orchestration framework in operator-like scenarios. Results show significant improvements in optical bandwidth utilization and sustained cost reduction of 60–85% per user compared to benchmarking network operation approaches.

Keywords—5G/6G traffic data, techno-economic evaluation

I. INTRODUCTION

The ITU-T Focus Group on Network 2030 is guiding network operators towards the definition of the requirements and use cases that future transport networks must support in the short, medium, and long term throughout this decade. To meet the demand of emerging 6G use cases and applications, operators must evolve their networks to satisfy stringent performance criteria across multiple dimensions, including the ability to manage large-scale bandwidth capacity in a flexible and adaptive manner [1]. To achieve such target behaviour, autonomous networking is a cornerstone for enabling seamless, zero-touch operation in next-generation networks. To support this transition, TM Forum offers strategic guidance to help operators adopt frameworks that enable the implementation and assessment of progressively higher levels of autonomy [2]. Leading operators' roadmaps target the realization of fully autonomous systems across multiple network technologies and domains by 2030 [3].

The adoption of novel technologies enabling full network automation to support unprecedented 6G services poses critical challenges during techno-economic evaluation phase. First, emerging technologies need to be evaluated following an incremental approach, on top of current legacy infrastructures and without affecting the assurance of current 4G/5G services. Second, the lack of actual 6G use cases data, e.g., 6G traffic databases, makes difficult to apply typical incremental network planning approaches where available traffic measurements are scaled according to an annual traffic growth rate to generate future term scenarios [4]. And finally, the still not extensive adoption of 5G systems (in both mobile and fixed transport networks) in real deployments hinder operators and network planners in the pursuance of clear answers about the cost and convenience of adopting novel 6G technologies for both data and control planes.

Recent research works have proposed frameworks for techno-economic evaluation in brownfield, gradual scenarios in support of technologies and services beyond 5G [5]. By means of aggregated traffic flows and matrices generated according to the expected characteristics of 6G use cases,

traffic volumes forecasts are used for network planning decisions, i.e., optical network capacity expansion. Other works have proposed the use of large language models (LLM) for traffic analysis and network planning, which complements human operators/engineers view and experience towards the construction of unseen and unprecedented scenarios [6].

In this work, we go beyond previous research works and present a two-fold contribution aiming at alleviating the obstacles towards open-wide research on operator network techno-economic evaluation in the advent of 6G services and technologies. First, a reference traffic dataset aggregating 4G/5G traffic measurements obtained from real network operation, as well as forecasted traffic for expected 6G services is presented. The methodology, that combines human operator expertise with the use of a LLM engine is detailed (Section II). The produced dataset, containing traffic time series for multiple 4G/5G categories and 6G services, has been made openly accessible according to the FAIR (Findable, Accessible, Interoperable, Reusable) principles [7]. Second, the reference traffic data has been used to evaluate the impact in terms of optical bandwidth utilization and cost of deploying a 6G-based control and orchestration system (Section III) on top of a legacy operator network. In particular, we adopt the innovations of DESIRE6G project [8], that proposes multi-scale control loops to enable full flexibility and real-time adaptability of heterogeneous services. Numerical evaluation has been carried out by means of a simulation environment that replicates real operator scenarios (Section IV).

II. REFERENCE TRAFFIC DATASET DESCRIPTION

This section describes the methodology and details of the reference 4G, 5G, and 6G traffic dataset, publicly available for techno-economic and performance evaluation studies. Due to confidentiality reasons, the descriptions have been carefully provided to meet the FAIR principles without compromising the privacy of network operator and customers.

A. Methodology

The methodology to generate the reference traffic dataset consists in two differentiated and consecutive phases. Phase I (Fig. 1) comprises real data collection from current network operation, as well as the procedure followed by operators to learn how the LLM engine analyzes raw traffic and extracts its main characteristics. This knowledge of the interaction between human operators and the LLM engine is required by Phase II (Fig. 2), where the reference traffic dataset for 4G, 5G, and 6G services is generated by the combination of collected data, human expertise, and LLM-based analysis.

The first step of Phase I is the collection of real traffic from the Telefónica mobile network infrastructure. In particular, a large *Raw Traffic* database (DB) containing the downlink (DL) traffic from the 100 domains with the highest traffic load and the uplink (UL) traffic to the top 100 destination domains with the highest traffic load is directly obtained from the




Fig. 1. Phase I scheme

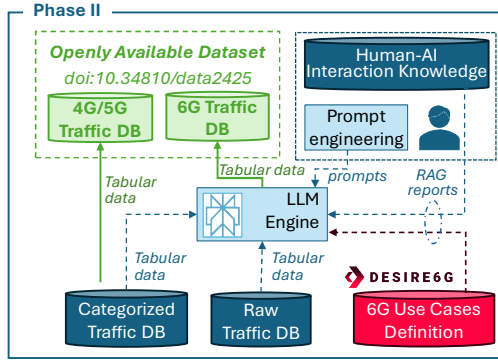


Fig. 2. Phase II scheme

interconnection network management system (NMS). This raw data aggregates the entire Telefónica network traffic measured during one week (from March, 10th 2025 to March 16th 2025), with a granularity of 30 minutes. Then, a *Human-Based Categorization* process, based on operator knowledge and background, is carried out to generate the *Categorized Traffic DB* where traffic is divided into 10 known classes:

- *Content Delivery Network (CDN)* including CDN provider domains, storage services platforms, software updates, sites/platforms managing secure sockets layer (SSL) certificates.
- *Cloud* including cloud provider domains, cloud apps and services, cloud computing
- *Email*
- *Internet-of-Things (IoT)*, including domains related with web analytics, geolocation services, telemetry services.
- *Non-Human Generated*, including traffic from devices without human interaction
- *Social Network*, including messaging and media sharing platforms, and social network apps.
- *Streaming*, including video-on-demand and live TV streaming sites, gaming platforms, music streaming, audio/video conference services and video surveillance service platforms.

- *Technology*, including services/programming interfaces for web developers, online tools for video editing, graphic design, business intelligence analytics and certificate validation services.
- *Web Browsing*, including web stores, search engines, advertisement services and platforms.
- *Other*, including the traffic not categorized as any of the other classes.

Once both raw and categorized DBs are available, they are injected as inputs to the Perplexity *LLM engine* [9] that produces artificial intelligence (AI)-based analysis reports from the traffic DBs and the prompts injected by the operator, e.g., asking the engine for finding unsupervised patterns and using that patterns to identify traffic/service classes. The reports are processed by the human operator with the objective to learn and produce *Human-AI Interaction Knowledge* reports, which consists in a set of rules and instructions to guide AI engine towards operator-like analysis. These reports contain not only corrections and indications to steer analysis, but also context of the traffic that is not specified in the collected data, e.g., the requirements of the different services. The interaction with the LLM engine is iteratively repeated until the knowledge is solid and the provided AI-based reports match with operator view.

Phase II generates the openly available dataset in [7], which contains two differentiated DBs. On the one hand, the *4G/5G Traffic DB*, which contains the Categorized Traffic DB obtained in Phase I, conveniently anonymized and normalized for reusability purposes. On the other hand, the *6G Traffic DB*, obtained by the LLM engine ingested with both Raw and Categorized Traffic DBs. Moreover, a Retrieval-Augmented Generation (RAG) strategy is followed to improve AI-based analysis [10]. Thus, two external sources are added: *i)* the Human-AI Interaction Knowledge reports generated in Phase I; and *ii)* a detailed report with the description of characteristics and requirements of the 6G use cases described in [11]. A summary of the considered 6G use cases along with their main characteristics/requirements is provided next:

- *Augmented Reality/Virtual Reality (AR/VR)*: Relative high bandwidth in DL direction, ultra-low latency, and high resilience.
- *Digital Twins (DT)*: ultra-low latency, moderated bandwidth.
- *Intelligent Image Monitoring*: Low-latency access, variable bandwidth video streaming.
- *Latency Sensitive Robot Control*: ultra-low latency, low bandwidth.
- *Cloud Rendered Gaming*: high throughput, low latency.

By means of prompt engineering, the human operator is able to retrieve the synthetic 6G Traffic DB from the extended, augmented LLM engine.

B. Dataset descriptive analysis

Each available DB in [7] consists in a large tabular file (.tab) with four columns:

- *Timestamp*: in normalized epochs.

- *Category/Use case*: either the category for 4G/5G traffic or the use case for 6G traffic.
- *DL traffic*: traffic in normalized units for the DL direction.
- *UL traffic*: traffic in normalized units for the UL direction

The number of rows (samples) is 3360 and 1680 for 4G/5G and 6G DBs, respectively. For illustrative purposes, Fig. 3 plots the data available for two example cases: category *Streaming* in 4G/5G traffic DB (Fig. 3a) and AR/VR use case in 6G traffic DB (Fig. 3b). Each time series contains 330 points and plots the traffic in the DL direction.




Fig. 3. Streaming category (a) and AR/VR 6G use case (b) examples

III. NETWORK COST ANALYSIS UNDER 6G OPERATION

A. Preliminaries

DESIRE6G contributed to 6G networks research with a novel, flexible, programmable, and multi-site 6G-native system designed to support latency-critical, distributed applications [12]. This DESIRE6G system allows verticals to express service requirements in a simple, intent-based way. Then, DESIRE6G employs several control loops at different timescales to ensure the committed service requirements (Fig. 4): *i*) real-time traffic management at mobile/fixed transport network level, *ii*) cloud-native user plane with local optimisation capabilities, *iii*) AI-driven, distributed service optimisation, and *iv*) intelligent service management and orchestration.

The essential unit to be dynamically managed is the *service graph*, i.e., the concatenation in both UL and DL directions of a set of virtual network functions (VNF) and their virtual links that describes a given service [8]. An instance of a service graph is supported, at a given time instant, by a set of resources (i.e., computing resources at sites and connectivity at network links and devices). Without loss of generality, in the context of this paper, we refer to these set of resources as *service slice*. Fig. 5 sketches an example where two service graph instances (in green and purple) span across several sites and networks. Each service graph is supported by its corresponding slice (in orange). For the sake of simplicity, in this paper we will focus on dynamic management of the allocated bandwidth resources in each fixed network segment/domain. This allocated bandwidth in a given network segment must ensure the required bitrate of the service graph virtual links. Note that the




Fig. 4. DESIRE6G multi-scale control loops (adapted from [12])




Fig. 5. Services graphs and service slices

heterogeneous nature of service graphs composition and NF placement makes service slices to be very variable and dynamic along time.

B. Problem statement

The techno-economic study carried out in this paper, based on a multi-period network planning and reconfiguration approach [4], can be stated as follows:

Given:

- A network topology $G(N, E)$ that connects the set of DESIRE6G sites N by means of a set of links E . Each link $e \in E$ is defined by a maximum number k_e of capacity units (e.g., optical interfaces) of bandwidth w (in Gb/s) available in that link, as well as the cost c_e of using a capacity unit in that link.
- A set of service graphs R , where each $r \in R$ is defined by the UL/DL graph (F_r, L_r) , being F_r the set of VNFs and L_r the virtual links connecting VNFs in F_r .
- A set of time intervals T to evaluate.
- The traffic demand B , containing the required bitrate $b_{r,l,t}$ of service graph r in virtual link l at time t . This depends on the variable traffic requirements in time.
- The set of potential VNFs placements Q , where $q_{r,f,n}$ is equal to 1 if and only if VNF f of service graph r can be potentially placed at site n . This depends on the latency and computing requirements of each NF and virtual link connecting functions.

Output:

- The actual placement of VNFs X , where $x_{rftn} \in X$ is equal to 1 if and only if VNF f of service graph r is placed at site n at time t .
- The allocated bandwidth Y , where $y_{et} \in Y$ is the number of capacity units of size w are needed at network link e to support the bitrate of all the service graphs allocated in that link at time t .

Objective: minimize the total cost of network bandwidth usage, while guaranteeing that bitrate requirements are accomplished along all time periods in T .

C. Methodology

For evaluation purposes, we developed a simulation environment (implemented in Python) that solves the optimization problem stated in Section III.B. The simulator reproduces a realistic operator network infrastructure following the details and characteristics described in [13]. In particular, 5 hierarchical levels (HL) for sites N are defined, namely: HL5 (access level), HL4 (aggregation level), HL3 (aggregation-transit level), HL2 (transit level), and HL1 (interconnection level). A topology $G(N, E)$ providing connectivity between all HLs, including a gateway for interconnectivity with other networks is also configured. In addition, a number of base stations (BS), also included in set N , are connected to HL4 and HL3 nodes. Each link $e \in E$ connecting two different types of sites (BS-to-HL or HL-to-HL) is equipped with a maximum capacity of $k_e \cdot w$ Gb/s.

The simulator allows configuring a number of traffic generators, each injecting the traffic of a single service graph $r \in R$. The reference traffic dataset in Section II is used to generate traffic demand B . In particular, different types of services based on the confidential number of active subscribers that the operator internally manages, are created. Then, a generator is attached to each BS, thus emulating the users that are served from that BS. Each generator is configured with the total number of users for that service, as well as with the normalized per-user traffic profiles (defined in terms of bitrate). In this way, scaling the traffic with time is a straightforward operation.

Besides the normalized per-user traffic profile, each service graph r includes the disaggregated radio access network (RAN) functions (RU/DU/CU), the 5G core (5GC) and the user plane function (UPF) elements, as well as a generic processing function that holds the specific tasks of the service, e.g., video processing tasks for the case of AR/VR use case. The service graphs distinguish between UL and DL, so both directions need to be configured for each service. Moreover, a bitrate scaling factor is specified for each of the virtual links L_r . Note that this is required because traffic volume is not constant along the service graph, e.g., for the same input user traffic, front-haul traffic (between RU and DU) is larger than backhaul one (between CU and 5GC/UPF).

The simulator finds the optimal VNF placement (variable matrix X) according to the traffic demand produced by the generators, as well as the allocated network bandwidth (variable matrix Y), and computes the total cost of allocated bandwidth:

$$\text{cost} = \min \sum_{t \in T} \sum_{e \in E} c_e \cdot y_{et} \quad (1)$$

Several primary constraints are always active to guarantee feasibility, e.g. NF placement and link capacity constraints:

$$x_{rftn} \leq q_{rfn}, \forall r \in R, f \in F_r, n \in N, (t, t') \in T \quad (2)$$

$$y_{et} \leq k_e, \forall e \in E, t \in T \quad (3)$$

A set of constraints force that all the VNFs are placed in a site, and account for the total bitrate to be supported at each network link (Boolean δ parameters define the adjacencies between N and E , and between F_r and L_r):

$$\sum_{n \in N} x_{rftn} = 1, \forall r \in R, f \in F_r, t \in T \quad (4)$$

$$y_{et} = \sum_{r \in R} \sum_{f \in F_r} \sum_{l \in L_r} \delta_{fl} \sum_{n \in N} \sum_{e \in E} \delta_{ne} \cdot b_{rlt} \cdot x_{rftn}, \quad \forall e \in E, t \in T \quad (5)$$

At this point, we can distinguish two different operation modes that can be executed in the simulator, which affects the additional constraints that are considered. The first one is the benchmarking *5G stand-alone (5G SA)* operational mode, where the heterogeneous allocation of slices is allowed, even for service graphs of the same type. However, due to the limitations of real-time optimization, slices remain invariable during the service lifetime, i.e., the allocated resources are fixed for the entire duration of the simulation. More formally, it can be expressed as follows:

$$x_{rftn} = x_{rftn'}, \forall r \in R, f \in F_r, n \in N, (t, t') \in T \quad (6)$$

$$y_{et} = y_{et'}, \forall e \in E, (t, t') \in T \quad (7)$$

A simple example of this operation is sketched in Fig. 6a, showing two different time steps of a single service graph: $t1$ represents a peak hour where demand is high, and $t2$ represents a valley hour where demand is low. As can be seen, both the placement of service functions and the allocated bandwidth resources remain fixed and sufficiently high to guarantee proper service operation during service lifetime.

In contrast to 5G SA operational mode, the *DESIRE6G operational mode* is based on all the DESIRE6G innovations introduced in Section III.A. To allow this 6G operation, constraints in equations (6) and (7) are relaxed (omitted), which allows different placement of VNFs along time for the same service (only constrained by the pre-computed potential placements in $\mathcal{Q}$). The example in Fig. 6b shows that service functions can be now grouped at BS and HL2 sites at time $t2$, thus releasing the capacity resources previously allocated in HL4 and HL3 at time $t1$. Moreover, bandwidth in aggregation and transit network segments can be reduced, which saves costs associated to network bandwidth utilization.

The simulator is configured with one scenario (topology + service demand) and one operational mode and executed for a simulated time duration, e.g., one week. To accelerate the simulation, the optimization problem is solved by means of a randomized heuristic that explore a wide range of feasible solutions and returns the best one after a number of iterations [4]. For the 5G SA approach, the optimal VNF placement is computed first, guaranteeing the service and optimizing the overall cost function. For the sake of fairness, we assume perfect knowledge of future traffic, so that slices are tightly dimensioned to ensure the continuous service performance. In the case of DESIRE6G approach, at every t , the demand is forecasted based on the previous time periods using a polynomial time series predictor [4] and the best placement and resource allocation of each service graph (using the same heuristic algorithm) is computed and evaluated.




Fig. 6. 5G SA operation mode (a) vs 6G operation mode (b)

IV. ILLUSTRATIVE RESULTS

For numerical evaluation purposes, we configured our simulation environment with the dense topology in Fig. 7, which is an anonymized version of a real Telefonica infrastructure. It covers an urban area of 176 km² providing connectivity service to around 1.5 million of potential users/subscribers. The topology consists in 660 BSs connected to 9, 2, and 1 HL4, HL3 sites, and HL2 sites, respectively. BSs consist of macro cells at all available frequencies, as well as 5G small cells in mid-band supporting massive multiple-input and multiple-output (MIMO) technology. Regarding fixed transport network, access (between BS and HL4), aggregation (between HL4 and HL3) and transit (between HL3 and HL2) levels are considered, with bandwidth supported by 100 Gb/s optical interfaces [11]. Each HL4 node is connected to both HL3 sites, and both HL3 connected to the HL2 one. Table I presents the main link parameters for all network levels.

Several simulations have been carried out for an incremental number of users, thus representing an evolutionary scenario based on actual traffic growth figures. Hereafter, the number of users has been normalized in order to facilitate results analysis. Thus, the minimum number of users represents the current starting point, where the actual mix of 4G/5G traffic is assumed and no 6G traffic is injected. As soon as the number of users increases, the proportion of 6G traffic use cases increases. We assume a yearly increment of 25% and 50% for 4G/5G and 6G traffic, respectively. When normalized number of users reaches one, the proportion of 6G traffic is near 50% of total traffic. Moreover, a normalized number of users larger than one produces some service rejection due to the lack of capacity resources.

Fig. 8 and Fig. 9 show the performance results of both 5G SA and DESIRE6G operation modes for both aggregation and transit levels, respectively; the figures show the mean and maximum bandwidth allocated as a function of the incremental number of users. As can be observed, DESIRE6G




Fig. 7. Dense urban topology

TABLE I. NETWORK LINK PARAMETERS

	w	k_e	c_e
Access (BS - HL4)	100 Gb/s	1	1 c.u.
Aggregation (HL4 - HL3)		10	10 c.u.
Transit (HL3 - HL2)		10	100 c.u.




Fig. 8. Mean (a) and maximum (b) used bandwidth in aggregation level




Fig. 9. Mean (a) and maximum (b) used bandwidth in transit level

operation provides a much better average use of resources than that of 5G SA operation. This leads to a reduction of costs derived from resource usage of services in a pay-as-you-go business model. Moreover, it can be observed that maximum network resource utilization, which is directly related with capital expenditures (CAPEX), scales better with the expected 6G traffic volume increase. In fact, the number of users that can be served with the same amount of capacity resources is larger with DESIRE6G operation. In addition, operational expenditure (OPEX) savings, such as energy consumption and maintenance, are reduced by a significant lower amount of active and running capacity resources, i.e. physical and virtual servers/machines running in sites and active transponders/interfaces in network links.

Finally, Fig. 10 focuses on cost figures, showing two different metrics: *i*) the total cost as defined in equation (1) but normalized to the maximum (Fig. 10a), and *ii*) the cost per user (Fig. 10b). In view of the figures, we can conclude that the DESIRE6G innovations allows reducing costs derived from operation with respect to 5G SA operation. The cost per users is remarkably low, and the savings compared with 5G SA increase with the expected users and traffic evolution (savings are embedded as percentages in Fig. 10)




Fig. 10. Cost figures: total cost (a) and cost per user (b)

V. CONCLUSIONS

This paper has presented the methodology used to generate the reference traffic dataset containing time series data of 4G, 5G, and 6G services. Specially, the great ability of the Perplexity LLM engine has resulted essential to validate the proposed methodology by expert human operators. The impact of publishing the reference dataset in open access mode is expected to be very high, since it aims at mitigating the lack of realistic data for next-generation network planning studies combining current and novel services and technologies.

The aforementioned dataset has been demonstrated as useful for performing techno-economic studies. In particular, the impact of the DESIRE6G control and orchestration innovations have been numerically evaluated in operator-like scenarios in terms of bandwidth and cost. The main conclusions and lessons learnt are three-fold. First, DESIRE6G enables efficient and scalable use of resources

along services lifetime leading to OPEX savings and reducing service costs by guaranteeing a feasible and robust pay-as-you-go business model. Second, DESIRE6G solutions allow large exploitation of deployed network capacity resources. Thus, investment rounds to increase capacity, i.e., CAPEX, can be delayed or at least, planned with more anticipation, due to its better efficiency in the use of resources. Finally, cost figures are remarkably low with DESIRE6G. The cost per service users drops when comparing with a 5G SA approach, reaching savings between 60% and 85% sustained in time (and even increasing), which again offers the possibility to operators to potentially increase revenues, leading to large incomes and operator infrastructure value.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the Smart Networks and Services Joint Undertaking under the European Union's Horizon Europe research and innovation programme under G.A. No. 101096466 (DESIRE6G), project PID2024-157824OB-I00 funded by MICIU /AEI /10.13039/501100011033 / FEDER, UE, and from ICREA.

REFERENCES

- [1] B. Da and M. Carugi, "Representative use cases and key network requirements for network 2030," Tech. rep., ITU-T FG-NET2030, 2020.
- [2] TM Forum, "Autonomous Network Levels Evaluation Methodology," v1.2.0, 2023, <https://www.tmforum.org/resources/introductory-guide/fig1252-autonomous-network-levels-evaluation-methodology-v1-2-0/>
- [3] O. González de Dios, P. Armingol, L. Roelens, *et al.*, "Automation of multi-layer multi-domain transport networks and the role of AI [invited]," J. Opt. Commun. Netw. 17, A124–A133, 2025.
- [4] L. Velasco and M. Ruiz, *Provisioning, Recovery and In-operation Planning in Elastic Optical Networks*, ISBN 978-1-119-33856-7, Wiley, 2017.
- [5] M. Ruiz, J. A. Hernandez, M. Quagliotti, *et al.*, "Network Traffic Analysis under Emerging Beyond-5G Scenarios for Multi-Band Optical Technology Adoption," J. Opt. Commun. Netw., 15, F36–F47, 2023.
- [6] C. Hang, P. Yu, R. Morabito, *et al.*, "Large Language Models Meet Next-Generation Networking Technologies: A Review," Future Internet, 16, 365, 2024.
- [7] M. Ruiz and M. A. Carnero Fernandez, "Network traffic measurements for 4G, 5G and 6G services", <https://doi.org/10.34810/data2425>, CORA.Repositori de Dades de Recerca, V1.
- [8] G. Pongracz, A. Mihaly, I. Godor, *et al.*, "Towards extreme network kpis with programmability in 6g," in Proc. 34th International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing, ser. MobiHoc23. New York, NY, USA: ACM, 2023, p. 340–345.
- [9] Perplexity, <https://www.perplexity.ai/> (last access: 01/27/2026)
- [10] W. Fan, Y. Ding, L. Ning, *et al.*, "A Survey on RAG Meeting LLMs: Towards Retrieval-Augmented Large Language Models," in Proc. ACM SIGKDD, 2024.
- [11] R. Lopez Da Silva, M. Groshev, S. Ertas, *et al.*, "D2.1: Definition of Use Cases, Service Requirements and KPIs/KVIs," Zenodo, 2023. doi: 10.5281/zenodo.10446047.
- [12] C. Papagianni, "DESIRE6G - Meeting extreme KPIs through Deep Programmability in 6G AI-Native Systems," SNS JOURNAL 2025, pp. 48–49, 2025.
- [13] M. Ruiz and M. A. Carnero Fernandez, "D2.3: Techno-economic Analysis and Data Set Definition," Zenodo, 2025. doi: 10.5281/zenodo.15799684.