A Value-Driven System Design Framework for Sustainable 6G Networks

Giancarlo Sciddurlo[a,b], Federica de Trizio[a,b], Giuseppe Piro[a,b], Gennaro Boggia[a,b]

[a] *Department of Electrical and Information Engineering, Politecnico di Bari, Bari, Italy*
[b] *CNIT, Consorzio Nazionale Interuniversitario per le Telecomunicazioni*

Abstract

The global challenges of the 21st century, including climate change, the digital divide, social inequalities, and poverty, pose significant obstacles to sustainable development, as highlighted by the United Nations agenda. Its 17 Sustainable Development Goals (SDGs) outline a vision for addressing these issues, necessitating innovative solutions powered by infrastructures such as 6G, a key driver of high-performance networking and societal progress. Achieving this dual objective requires a value-driven approach that extends beyond traditional communication network functionalities. In this context, Key Value Indicators (KVIs) emerge as essential metrics that capture intangible yet critical societal values, complementing traditional Key Performance Indicators (KPIs). However, the integration of ethical principles and social values into networking remains largely unexplored in the scientific literature and research communities. Building on these premises, this work introduces a system design framework that formalizes and evaluates KVIs alongside KPIs, leveraging Intent-Based Networking (IBN) to embed ethical and social dimensions into network services. To enable responsible resource-service provisioning, a bi-objective optimization problem is formulated, balancing network performance and societal value. The problem is solved using an exact ε-constraint method, ensuring optimal trade-offs between KPIs and KVIs. The results validate the effectiveness of the proposed service orchestration framework, demonstrating an improvement of up to 70% in the ethical and social value provided by the network compared to baseline solutions. This highlights the framework's capability to integrate social and ethical considerations into the service allocation process while preserving competitive network performance.

Keywords: 6G, Key Value Indicators, Sustainable Development Goals, Intent-Based Networking

1. Introduction

The global challenges of the 21st century, including climate change, the digital divide, social inequalities, and poverty, present significant barriers to sustainable development worldwide. The scale and complexity of these issues are underscored by the United Nations (UN) agenda, which delineates a strategic plan to address them through 17 Sustainable Development Goals (SDGs) [1]. Realizing the vision of global sustainability, founded on the principles of inclusivity, trust, and resilience, requires innovative solutions enabled by advanced technological infrastructures [2].

Building upon the advancements of previous network generations, 6G aims to deliver broader coverage, higher data rates, denser connectivity, and ultra-low latency to enable next-generation applications such as holographic telepresence, extended reality, and collaborative robotics [3]. Furthermore, it is anticipated to play a primary role in promoting societal progress and sustainability [4] by enhancing access to essential services in rural and remote regions and facilitating the development of educational and healthcare infrastructures [5, 6].

The dual objective of enhancing network performance while simultaneously addressing social and ethical values necessitates a value-driven approach that expands the utility and potential of communication networks. In this context, innovative Key Value Indicators (KVIs) metrics emerge to capture intangible and non-quantifiable aspects, typically expressed qualitatively or mapped to specific ranges and targets. Unlike conventional Key Performance Indicators (KPIs), which primarily focus on operational efficiency without considering the broader societal impact of networks, leveraging KVIs enables a greater emphasis on inclusiveness, trustworthiness, sustainable energy consumption, resource efficiency, and reliability [7].

However, despite the significant opportunities presented by value-driven strategies within advanced technological infrastructures, the existing scientific literature has largely overlooked the integration of ethical principles and social values into networking frameworks and services to address the global challenges of the 21st century. On the one hand, apart from performance evaluation, research in this domain predominantly focuses on energy efficiency as the primary network value parameter [8–17]. This emphasis is evident in both research and standardization communities, where broader challenges outlined by the SDG remain insufficiently addressed [18, 19]. On the other hand, when KVIs extend beyond energy efficiency, they are predominantly defined through qualitative methodologies tailored to specific use cases, as demonstrated in both academic research and European initiatives [20, 21].

To the best of the authors' knowledge, no prior research has explored the potential of guiding next-generation networks toward ethical networking, enabling the simultaneous fulfillment of performance, ethical, and social values, while actively addressing global challenges aligned with the UN's SDGs.

To bridge this gap, this study proposes a methodology to formalize and evaluate KVIs alongside KPIs, aiming to enhance the value and impact of networking services for users. This approach leverages the advantages of the Intent-Based Networking (IBN) paradigm [22], which simplifies network management in complex environments with distributed endpoints and diverse applications by enabling a more autonomous operational model. Within this framework, users define high-level objectives or intents in human-readable terms, which are then translated into ethical and actionable policies, followed by systematic deployment and assurance [23]. Differently from previous works, the proposed strategy enables the dynamic integration of social and ethical dimensions into services, incorporating aspects such as trustworthiness, inclusiveness, and environmental sustainability. Consequently, the impact of communication networks on SDGs can be systematically monitored and assessed through a service orchestration solution that optimizes metrics extending beyond traditional performance indicators.

The main contributions of this work can be summarized as follows:

- A system design is proposed for acquiring service requests expressed in natural language, processing them using a Large Language Model (LLM) to generate trustworthiness, inclusiveness, and environmental sustainability metrics alongside traditional performance indicators. The system can precisely interpret service expectations, orchestrate and optimize them, and subsequently deploy them within a communication network. By leveraging IBN, this approach integrates a social dimension directly into the service delivery process.

- A bi-objective optimization problem is formulated to balance network performance, represented by the KPIs of resources providing requested services, with the network's social and ethical value, represented by the KVIs of resources providing these services. This enables the alignment of service provisioning with the most suitable and responsible network resources.

- To optimize the selection of service and resource combinations that maximize network utility and potential in service delivery, the exact ε-constraint method is applied. This technique reformulates the original multi-objective problem into a single-objective problem, subsequently generating and solving a sequence of constrained optimization problems to identify all optimal solutions.

- The obtained results indicate a substantial enhancement in the social and ethical value provided by the network during the service delivery process, achieving up to a 70% increase compared to baseline approaches that focus solely on network performance parameters. This improvement is consistently observed across all evaluated scenarios, which encompass diverse 6G services with varying requirements in terms of environmental sustainability, trustworthiness, and inclusiveness.

The remainder of this paper is organized as follows: Section 2 reviews related works. Section 3 describes the reference scenario, formulates the system model, and defines the optimization problem. Section 4 presents the exact ε-constraint method, while Section 5 provides performance evaluations, including comparisons with baseline approaches. Finally, conclusions are drawn in Section 6.

2. Related Works

Advancing the potential of next-generation networks requires innovative design and orchestration that jointly meet KPIs and KVIs, ensuring efficient, ethical, and sustainable use of network resources.

Initially, scientific research on service orchestration methods primarily focused on network design to minimize delays and maximize throughput [24, 25]. More recently, ongoing studies have increasingly explored solutions to enhance sustainability by evaluating energy efficiency [8–17], as reducing energy wastage is crucial for ensuring both the affordability and longevity of energy resources. In this context, the authors of [8] formulate an optimization problem to minimize long-term carbon emissions and energy trading costs, decomposing it into three independent subproblems using Lyapunov optimization. The study in [10] presents a model for assessing the feasibility of future green cellular networks, formulating an optimization problem to minimize emission costs and greenhouse gas emissions, which is solved numerically for small networks with varying peak traffic profiles. The authors of [11] explore an in-network computing model designed to reduce data center energy consumption by leveraging virtualization and software-defined networking technologies. Specifically, they model a multi-objective task scheduling optimization problem, addressed using an evolutionary algorithm based on multiple target decomposition. Similarly, the authors of [14] propose an optimization strategy for Virtual Machines (VMs) allocation, formulating an optimization problem to minimize both cost and energy consumption, which is solved heuristically using allocation and migration algorithms. While the growing demand for network transmission capacity and data processing significantly increases computational requirements and energy consumption, it is essential to expand the concept of sustainability through an interdisciplinary perspective that integrates broader societal objectives into network design and operation. In this context, IBN offers a promising paradigm for translating broader societal objectives into actionable mechanisms within network orchestration by automatically associating services with KVIs, thereby enabling systematic tracking of progress toward the SDGs. However, existing scientific literature primarily explores this paradigm with a focus on delivering services based on traditional KPIs [26–40]. For instance, the studies in [26, 27] propose an intent fulfillment framework that translates intents into policy trees using LLMs, following the Metro Ethernet Forum (MEF) Policy Driven Orchestration (PDO) model and executing them through Finite State Machines (FSMs). Complementing these efforts, [35] introduces an intent-based system that leverages natural language

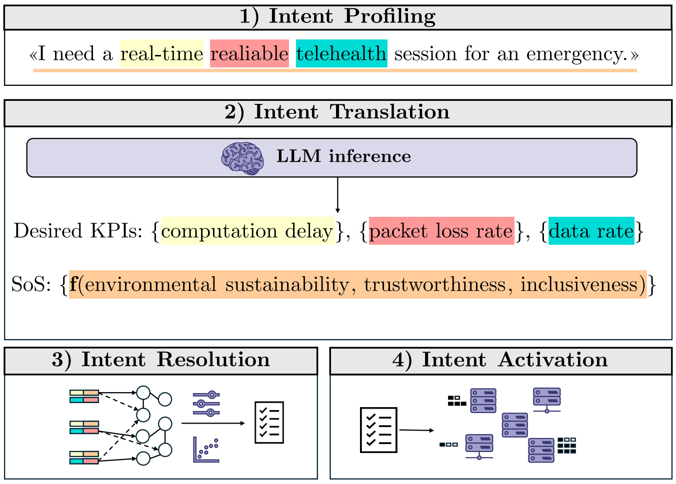

Figure 1: The proposed intent-based framework.

processing to interact with users, interpret their requirements, and provide context-aware provisioning responses, dynamically configuring network paths to meet specific application needs.

Expanding upon the current state of the art, our previous research works [41, 42] utilize IBN to enforce and optimize KVIs, enabling the seamless integration of sustainability, inclusivity, and trustworthiness metrics. Specifically, in [41], we design a malicious intent detection module to identify harmful service requests and prevent their propagation within the network, thereby enhancing trustworthiness, resilience, and availability. Meanwhile, [42] introduces a novel service orchestration framework that aligns network resource management with broader societal objectives, such as economic sustainability and security, while minimizing Service Level Agreement (SLA) violations. This is achieved through a multi-criteria decision-making algorithm and a matching game, enabling dynamic and efficient service orchestration that meets both technical and ethical requirements.

Although some efforts have been made, the integration of societal and ethical dimensions into networking systems through KVIs remains largely unexplored. Previous research has primarily focused on performance, often associating sustainability with KPIs or energy efficiency. Broader values such as inclusiveness, trust, and equity have largely been addressed at a conceptual level or within isolated use cases, due to the inherent challenges in quantifying them compared to technical metrics. To bridge this gap, this work leverages IBN to automatically map natural language requests to explicit parameters expressed through KPIs and KVIs, facilitating the management of network requirements and service provisioning while promoting value-driven and ethical networking.

3. System Model and Problem Formulation

This section introduces the proposed system design based on the IBN paradigm, detailing the system model and problem formulation for orchestrating and optimizing services within a value-driven communication network.

3.1. Proposed Intent-Based System Design

Figure 1 presents the architecture of the proposed intent-based framework, which automatically translates end-user service requests, typically expressed in natural language, into structured network intents enriched with both KPIs and KVIs. To achieve this, end-users interact with the system via high-level interfaces such as applications, templates, or chatbots, without needing to specify technical parameters. To bridge this semantic gap, the framework incorporates a LLM as a core component for interpreting user inputs. The LLM extracts relevant information from natural language requests, inferring both quantitative performance requirements (e.g., latency, data rate, packet loss rate) and value-driven indicators (e.g., trustworthiness, inclusiveness, sustainability), and generates a structured representation that constitutes an intent. Key semantic components are highlighted with distinct colors in Figure 1 to illustrate their roles in identifying both KPIs and KVIs. Following this procedure, intents enable the system to accurately interpret requirements, objectives, and constraints to effectively guide network operations [43]. Their declarative nature abstracts service requesters from the underlying resource and network infrastructure state, delegating request interpretation to a dedicated intent-processing mechanism [44]. Specifically, following the standard intent lifecycle, the proposed framework consists of four key components: Intent Profiling, Intent Translation, Intent Resolution, and Intent Activation.

In detail, the procedure begins with the **Intent Profiling** component, where service requesters define desired outcomes in natural language. These high-level inputs are then forwarded to the **Intent Translation** component, which leverages LLMs to infer the corresponding performance metrics and value indicators, generating structured intents accordingly. LLMs are well-suited for this task, as they are advanced artificial intelligence models based on the Transformer architecture, designed to process and generate human-like language [45, 46]. These models, trained on vast textual corpora and consisting of billions of parameters, exhibit strong capabilities in text generation, factual information retrieval, and complex logical and temporal reasoning. As a result, this module fine-tunes a general-purpose LLM, improving its ability to interpret service requests and generate structured inputs aligned with the intended optimization objectives. For this work, the mapping capability is assumed to function correctly, as the primary focus is not on LLM alignment or its associated challenges, but rather on the system's capacity to detect and interpret human intents expressed in natural language. It is well recognized that achieving robust alignment is highly complex: human preferences and values are diverse, context-dependent, and evolve over time. Additionally, real-world environments involve conflicting objectives and shifting goals, further complicating alignment efforts [47, 48].

Instruction fine-tuning serves as a mechanism to adapt LLMs for more complex tasks. During pre-training, LLMs acquire the capability to comprehend instructions and generate responses; however, this ability remains latent until activated through a form of supervised fine-tuning. By leveraging a dataset comprising instructions and their corresponding outputs, fine-tuning enables the model to learn the specific characteristics of future user inputs, including KPIs and KVIs, which encapsulate human, ethical, and social values. This process mirrors standard model training but is conducted on a significantly smaller dataset compared to pre-training. Within the intent mapping procedure, a service request is transformed into a sequence of tokens, where tokens represent small units of information, such as individual words or word groups. The LLM predicts the next token y_t based on the conditional probability $P(y_t \mid y_{0\to t}, txt)$, where txt denotes the input natural language sequence, and $y_{0\to t}$ represents the previously generated tokens. Through this process, the LLM can comprehend the context of the user-provided natural language request (txt), infer relevant KPIs and KVIs, and subsequently generate performance parameters. Moreover, the model produces an inclusive metric, referred to as the Sustainability of Service (SoS), which encapsulates the three social and ethical dimensions embedded in a given service request [7].

The generated performance parameters correspond to traditional KPIs, such as service delay, data rate, and packet loss rate, which are specified as desired values and acceptable tolerances, respectively. Furthermore, by leveraging LLM inferences, critical social and ethical dimensions, such as trustworthiness, inclusiveness, and environmental sustainability, are encapsulated into KVIs, thereby addressing broader challenges outlined in the UN's SDGs [20, 49]. Aligning with these principles:

- The trustworthiness value aims to ensure network security, robustness, and privacy, thereby safeguarding resilient technological infrastructures.
- The inclusiveness value focuses on equitable access to digital technologies, reducing the digital divide, and empowering diverse communities.
- The sustainability value encompasses the ecological dimension by prioritizing energy efficiency, fostering economic growth, and mitigating environmental impact.

After the translation process, the **Intent Resolution** component orchestrates and optimizes resource-service selection, balancing the trade-off between maximizing network performance and adhering to social and ethical values. The details of the system model and the procedure for optimal service orchestration in this context are described in the following sections.

Finally, the **Intent Activation** component enforces the resolved intent within the physical network infrastructure. A continuous monitoring mechanism ensures that intents remain fulfilled over time, leveraging a closed-loop system to enhance adaptability and learning. This guarantees that the requested service is effectively deployed while maintaining strict adherence to both KPIs and KVIs.

3.2. System Model

Based on the described intent-based system design, this work envisions a communication infrastructure where network resources provide services, defined as $\mathcal{R} = \{r_1, r_2, \ldots, r_n, \ldots, r_N\}$. Each resource $r_n \in \mathcal{R}$ declares its capabilities $(C_{r_n}, y_{1_{r_n}}, y_{2_{r_n}}, \ldots, y_{i_{r_n}}, \ldots, y_{I_{r_n}})$, where C_{r_n} represents the available cores of the r_n-th resource allocated for service execution. Meanwhile, $y_{1_{r_n}}, y_{2_{r_n}}, \ldots, y_{i_{r_n}}, \ldots, y_{I_{r_n}}$ denote the exposed performance attributes used to evaluate the considered KPIs, such as computation delay, data rate, and packet loss rate.

Then, users in the network can request services of various types, represented by the set $\mathcal{S} = \{s_1, s_2, \ldots, s_j, \ldots, s_J\}$.

Each requested service $s_j \in \mathcal{S}$ is defined by a tuple of KPIs, $(D_{s_j}, y_{1_{s_j}}, y_{2_{s_j}}, \ldots, y_{i_{s_j}}, \ldots, y_{I_{s_j}}, y_{1'_{s_j}}, y_{2'_{s_j}}, \ldots, y_{i'_{s_j}}, \ldots, y_{I'_{s_j}})$, where D_{s_j} denotes the required amount of processing resources for executing the s_j-th service, assuming that each resource corresponds to a single task of the service, with each task requiring one processing core for execution; $y_{1_{s_j}}, y_{2_{s_j}}, \ldots, y_{i_{s_j}}, \ldots, y_{I_{s_j}}$ represent the desired performance parameters, such as service delay, data rate, and packet loss rate. Meanwhile, $y_{1'_{s_j}}, y_{2'_{s_j}}, \ldots, y_{i'_{s_j}}, \ldots, y_{I'_{s_j}}$ define the minimum acceptable performance thresholds that the service can tolerate for the corresponding KPIs.

Based on these performance thresholds, the proposed system model defines $Q_{tol_{s_j}}$ as the overall quality tolerance of a service request. This parameter quantifies the permissible deviation from the ideal requested performance within the range $0 \leq Q_{tol_{s_j}} \leq 1$. It can be computed as:

$$Q_{tol_{s_j}} = \sum_{i=1}^{I} w_i \cdot y_{\hat{i_{s_j}}}, \tag{1}$$

where w_i is a weight representing the importance of the i-th KPI, ensuring that $\sum_{i=1}^{I} w_i = 1$ and $0 \leq w_i \leq 1$. Additionally, $y_{\hat{i_{s_j}}}$ represents the normalized value of the i-th KPI for the service request s_j, computed using the min-max normalization model:

$$y_{\hat{i_{s_j}}} = \frac{y'_{i_{s_j}} - \min(y'_i)}{\max(y'_i) - \min(y'_i)}$$

where $\max(y'_i)$ and $\min(y'_i)$ denote the maximum and minimum observed values of the i-th KPI across all service requests.

Given the presence of various standards for quantifying performance attributes, the proposed model defines the normalized performance quality related to the i-th attribute of the r_n-th network resource facing the s_j-th service as $q_{\hat{i_{r_n,s_j}}}$, constrained within the range $[0, 1]$. This normalization is performed using the standard model [50], formulated as follows:

$$q_{\hat{i_{r_n,s_j}}} = \begin{cases} 1 - \frac{\max(y_{i_{r_n}}) - y_{i_{s_j}}}{\max(y_{i_{r_n}}) - y'_{i_{s_j}}}, & \text{if the KPI is a benefit.} \\ 1 - \frac{y_{i_{s_j}} - \min(y_{i_{r_n}})}{y'_{i_{s_j}} - \min(y_{i_{r_n}})}, & \text{if the KPI is a cost.} \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

The overall KPI value, denoted as Q_{r_n,s_j} reflects the level of performance quality by aggregating the normalized values of various attributes and it is computed as follows:

$$Q_{r_n,s_j} = \sum_{i=1}^{I} w_i \cdot q_{\hat{i_{r_n,s_j}}}. \tag{3}$$

Without loss of generality, Table 1 presents a summary of the key symbols used to describe the environment, along with their corresponding definitions.

Table 1: Main Symbols Description.

Symbol	Meaning
$\mathcal{R}$	Set of network resources
$\mathcal{S}$	Set of requested services
N	Number of network resources [#]
J	Number of services [#]
C_{r_n}	Cores of the r_n-th resource
D_{s_j}	Resource demand of the j-th service [#]
$y_{i_{r_n}}$	Exposed performance attribute of the r_n-th resource
$y_{i_{s_j}}$	Desired performance parameter of the s_j-th service
$y_{i'_{s_j}}$	Minimum tolerable value for $y_{i_{s_j}}$ parameter
$\hat{q_{i_{r_n,s_j}}}$	Normalized performance quality of i-th attribute for r_n facing s_j
$Q_{tol_{s_j}}$	Overall quality tolerance of the s_j-th service request
Q_{r_n,s_j}	Overall quality KPI for r_n providing the s_j-th service
V_{r_n,s_j}	Overall social and ethical KVI of r_n providing s_j
w_1, w_2, w_3	Weights for the KPIs relevance
w'_1, w'_2, w'_3	Weights for the KVIs relevance
σ_{r_n,s_j}	Environmental sustainability indicator
θ_{r_n,s_j}	Trustworthiness indicator
ι_{r_n,s_j}	Inclusiveness indicator
$t_{comp}^{r_n,s_j}$	Computing delay of the r_n-th resource facing the s_j-th service [s]
B_{s_j}	Input size of the data to be processed for the s_j-th service request [Mb]
Clk_{r_n}	Processing capability of the r_n-th resource [$Megacycles/s$]
λ_{r_n,s_j}	Average number of service requests that r_n can handle per day
co_{r_n}	Carbon offset associated with the r_n-th resource [$gCOe_2$ per day]
$u_c^{r_n}$	Core usage factor of the r_n-th resource
$P_c^{r_n}$	Power consumption of a single core of the r_n-th resource [kW]
$P_m^{r_n}$	Power consumption of memory of the r_n-th resource [kW]
PUE	Power usage effectiveness coefficient
CI	Carbon intensity factor [$gCOe_2/kWh$]
ρ_{r_n,s_j}	Cyber risk of r_n providing s_j
κ_{r_n,s_j}	Cyber confidence associated with r_n facing s_j-th service
$MTTF_{r_n}$	Mean time to failure of the r_n-th resource [h]
p_{r_n}	Probability of failure of the r_n-th resource within a day
x_{s_j,r_n}	Binary variable indicating if resource r_n provides service s_j
$Q(X)$	Network quality performance function
$V(X)$	Network social and ethical function
δ	Discretization parameter

3.3. *Social and ethical values*

The primary advancement of the proposed approach lies in its ability to deliver services that are aligned with social and ethical KVIs. Specifically, three indicators, such as environmental sustainability, trustworthiness, and inclusiveness, are defined to reflect intangible but essential societal values, measuring the impact of communication networks on SDGs.

The *environmental sustainability indicator* σ_{r_n,s_j} quantifies the climate impact of service provisioning, expressed in grams of carbon dioxide equivalent. This KVI effectively captures the global warming potential of greenhouse gas emissions generated within the provisioning time frame. The indicator accounts for the greenhouse gas emission compensation of the r_n-th resource, achieved through the acquisition of carbon credits or offsets, where each unit corresponds to a tonne of gCO_2e reduced or removed from the atmosphere. This compensation is facilitated by funding initiatives such as reforestation, renewable energy development, or carbon capture projects [51]. This mechanism enables individuals, enterprises, and governmental entities to mitigate their environmental footprint and advance toward carbon neutrality, an urgency widely acknowledged. Environmental sustainability extends well beyond carbon reduction and encompasses metrics at the device, equipment, flow, path, and domain levels, as well as the prevention of various forms of environmental degradation [52]. In this context, the current indicator for environmental sustainability does not preclude the inclusion of complementary metrics addressing other ecological concerns.

Formally, the environmental sustainability indicator quantifies the difference between the carbon offset co_{r_n} of the r_n-th resource, characterized by a carbon intensity factor CI (expressed in gCO_2e/kWh), and its carbon footprint associated with the provisioning of the s_j-th service.

For each $r_n \in \mathcal{R}$, the time required for the r_n-th resource to process the s_j-th service, denoted as computing delay $t_{comp}^{r_n,s_j}$, can be calculated as reported in Eq. 4.

$$t_{comp}^{r_n,s_j} = \frac{\Phi \cdot B_{s_j}}{Clk_{r_n}}, \tag{4}$$

where Φ represents the number of CPU cycles required to process a single bit. According to [53], this value is set to 1000 cycles per bit. Furthermore, B_{s_j} denotes the input size of the data to be processed for the requested service, while Clk_{r_n} represents the processing capability of the r_n-th resource, expressed in [$cycles/s$].

Given λ_{r_n,s_j}, the average number of service requests s_j that the r_n-th resource can handle per day, the environmental sustainability indicator can be computed as follows [54]:

$$\sigma_{r_n,s_j} = co_{r_n} - [t_{comp}^{r_n,s_j} \cdot \lambda_{r_n,s_j} \cdot C_{r_n} \cdot (P_c^{r_n} \cdot u_c^{r_n} + P_m^{r_n}) \cdot PUE \cdot CI], \tag{5}$$

where $P_c^{r_n}$ represents the power consumption of a single processing core, and $P_m^{r_n}$ corresponds to the power consumption of memory, both expressed in kilowatts (kW). The core usage factor ($u_c^{r_n}$), ranging from 0 to 1, represents the fraction of total execution time during which a processor core is actively engaged in computation. A value of 1 indicates full, continuous utilization throughout the entire period, while values below 1 reflect intervals of inactivity or reduced load. This parameter allows adjustment of theoretical power consumption, often based on the processor's thermal design power, to more accurately represent real operational conditions. When detailed usage data is unavailable or impractical to obtain, setting $u_c^{r_n} = 1$ provides a conservative estimate that likely overestimates actual energy consumption. Furthermore, PUE denotes the power usage effectiveness coefficient, set to 1.67 as the globally averaged measured value [54], while CI represents the carbon intensity factor, quantifying the equivalent carbon emissions associated with the consumed energy. The CI varies significantly between countries, reflecting differences in national energy mixes—lower values are typically found in countries relying primarily on renewable sources, while higher values correspond to regions heavily dependent on fossil fuels.

Secondly, the *trustworthiness indicator* θ_{r_n,s_j} quantifies the security level of the service provisioning offered by the r_n-th resource to the s_j-th service. As discussed in our previous work [42], successful cyber-attacks can severely degrade network performance, leading to substantial consequences for both infrastructure and service reliability. Therefore, it is essential to quantify the cyber risk that threatens the confidentiality, integrity, and availability of data and services associated with each r_n-th resource. Following the methodology outlined in [42], the likelihood of a successful cyber attack, denoted as L_{r_n}, can be computed for each resource. Consequently, the associated cyber risk ρ_{r_n,s_j} is defined as the product of L_{r_n} and the impact Δ_{s_j} of an attack on the provisioning of service s_j. The impact parameter Δ_{s_j} quantifies the potential consequences and disruptions resulting from a successful cyber attack. Thus, the cyber confidence κ_{r_n,s_j}, expressing the confidence of the r_n-th resource in providing the s_j-th service is formally calculated as:

$$\kappa_{r_n,s_j} = 1 - \rho_{r_n,s_j}. \tag{6}$$

The indicator θ_{r_n,s_j} is modeled as a generalized logistic function, which increases with cyber confidence. This formulation ensures that more resilient resources, characterized by enhanced security measures and lower cyber risk, attain higher values of θ_{r_n,s_j}. The adoption of a logistic function effectively captures the nonlinear relationship between cyber confidence and trustworthiness, as outlined in [55] and reported as follows:

$$\theta_{r_n,s_j} = L_\theta + \frac{U_\theta - L_\theta}{1 + e^{-B\left(\kappa_{r_n,s_j} - x_0\right)}}. \tag{7}$$

Here, L_θ represents the lower bound of the trustworthiness achievable by a resource lacking strong security guarantees. This value may vary across providers, as even resources without a robust security profile can exhibit a baseline level of trust due to factors such as third-party management, secure default configurations, or adherence to regulatory standards. Conversely, U_θ defines the upper bound of attainable trust, recognizing that even highly secure systems are still subject to residual risks, including zero-day vulnerabilities and insider threats. The parameter B controls the steepness of the curve, determining how rapidly trust increases with an improved security posture. A higher absolute value of B results in a sharper transition near the inflection point, while lower values produce a more gradual increase appropriate for contexts where trust evolves over time due to long-term organizational or cultural developments. The inflection point x_0 denotes the security level at which a resource begins to be perceived as trustworthy. This threshold can vary depending on the complexity of the organization, as more intricate infrastructures typically face greater exposure to systemic threats. Importantly, highly vulnerable resources may exhibit minimal risk reduction from initial countermeasures, whereas mature and secure systems often experience diminishing returns from additional security investments. This reflects real-world cybersecurity dynamics, where the impact of interventions depends heavily on the system's initial resilience, and risk mitigation follows a nonlinear trajectory.

Furthermore, the *inclusiveness indicator*, denoted as ι_{r_n,s_j}, quantifies the network's capability to deliver adequate service to users, ensuring operational continuity and resilience for uninterrupted service functionality. In the context of this work, inclusiveness is approximated using availability as a proxy. This property is evaluated based on the time-to-failure TF_{r_n}, which is modeled using an exponential distribution under the assumption of a constant failure rate. The parameter $MTTF_{r_n}$ represents the mean time to failure (MTTF), which denotes the average time between consecutive failures. This metric is crucial in assessing the reliability of the r_n-th provider. The probability that the r_n-th provider experiences failure or downtime within a given time period t is expressed as:

$$P(TF_{r_n} \leq t) = 1 - e^{-\frac{t}{MTTF_{r_n}}} \tag{8}$$

Let p_{r_n} denote the probability of failure of the r_n-th resource within a day. Assuming that the C_{r_n} available cores of the r_n-th resource fail independently with probability p_{r_n}, the probability of observing k failures within an hour follows a binomial distribution and is given by:

$$P(F_{r_n} = k) = \binom{C_{r_n}}{k}(p_{r_n})^k(1 - p_{r_n})^{C_{r_n}-k}, \quad k = 0, 1, \ldots, C_{r_n} \tag{9}$$

Thus, the probability of no failures occurring in an hour, denoted as $P\left(F_{r_n} = 0\right)$ is given by:

$$P\left(F_{r_n} = 0\right) = \left(1 - p_{r_n}\right)^{C_{r_n}} \tag{10}$$

The inclusiveness indicator can thus be expressed as:

$$\iota_{r_n,s_j} = t_{comp}^{r_n,s_j} \cdot \lambda_{r_n,s_j} \cdot \left(1 - p_{r_n}\right)^{C_{r_n}} \tag{11}$$

Finally, the overall social and ethical KVI associated with the provisioning of the s_j-th service by the r_n-th network resource is defined as follows:

$$V_{r_n,s_j} = w'_1 \cdot \sigma_{r_n,s_j} + w'_2 \cdot \theta_{r_n,s_j} + w'_3 \cdot \iota_{r_n,s_j} \tag{12}$$

where w'_1, w'_2, w'_3 are weighting factors ranging within $[0, 1]$, with their sum equal to 1, representing the relative importance of each specific indicator.

3.4. Problem Formulation

Let x_{r_n,s_j} denote a binary variable that indicates whether service request $s_j \in \mathcal{S}$ is allocated to network resource r_n, and let X be the set of all binary indicator variables:

$$x_{r_n,s_j} = \begin{cases} 1 & \text{if } r_n \text{ provides service for } s_j \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

The network's capability to deliver services can be defined as the overall quality KPI derived from the matching of services with available resources, as reported in the following:

$$Q(X) = \sum_{s_j \in \mathcal{S}} \sum_{r_n \in \mathcal{R}} Q_{r_n,s_j} \cdot x_{r_n,s_j}. \tag{14}$$

Similarly, the network's ability to provide value-oriented and ethical services can be defined as the total KVI derived from the matching of services with available resources:

$$V(X) = \sum_{s_j \in \mathcal{S}} \sum_{r_n \in \mathcal{R}} V_{r_n,s_j} \cdot x_{r_n,s_j}. \tag{15}$$

For users requesting services, the benefit lies in receiving satisfactory services by jointly ensuring performance quality through KPIs and societal and ethical values through KVIs. Accordingly, the following multi-objective optimization problem is formulated and presented in Eq. (16):

$$\max_X (Q(X), V(X)), \tag{16}$$

$$\text{s.t.} \quad \left(Q_{r_n,s_j} - Q_{tol_{s_j}}\right) \cdot x_{r_n,s_j} \geq 0, \forall s_j \in \mathcal{S}, \forall r_n \in \mathcal{R}, \tag{17}$$

$$V_{r_n,s_j} \cdot x_{r_n,s_j} \geq 0, \forall s_j \in \mathcal{S}, \forall r_n \in \mathcal{R}, \tag{18}$$

$$\sum_{s_j \in \mathcal{S}} D_{s_j} \cdot x_{r_n,s_j} \leq A_{r_n}, \forall r_n \in \mathcal{R}, \tag{19}$$

$$\sum_{r_n \in \mathcal{R}} x_{r_n,s_j} = 1, \forall s_j \in \mathcal{S} \tag{20}$$

$$x_{r_n,s_j} \in \{0,1\}, \forall s_j \in \mathcal{S}, \forall r_n \in \mathcal{R} \tag{21}$$

Constraint (17) ensures that the service assigned to network resources meets the overall quality tolerance of a service request. Similarly, constraints (18) ensure that the social and ethical value provided by network resources is greater than zero. Constraint (19), instead, ensures that the total requested services do not exceed the availability of network resources. Constraints (20) ensure that each service is assigned to exactly one network resource, whereas constraints (21) enforce the binary nature of the indicator decision variables.

The defined problem (16) is a generalization of the maximum Generalized Assignment Problem (GAP), which is known to be NP-hard [56]. In practice, efforts to enhance the quality and performance of service provision may lead to increased system complexity or an expanded attack surface, thereby potentially introducing new security concerns. Consequently, since the problem involves two conflicting objective functions, a solution that simultaneously satisfies both objectives may not always exist. In the context of multi-objective optimization problems, the concept of Pareto optimality can be utilized to assess the optimality of solutions. A solution X^* is considered Pareto optimal or dominant if there exists no solution X such that $Q(X) \geq Q(X^*)$ and $V(X) \geq V(X^*)$, where X is any solution distinct from X^* within the feasible region. To obtain all dominant or Pareto-optimal solutions, the exact ε-constrained method can be employed [57, 58].

4. The Proposed Solution

This section presents the solution approach used to address the NP-hard problem described in Section 3, leveraging the exact ε-constraint technique.

4.1. Exact Epsilon-Constraint Method

Multi-objective optimization problems involve two or more conflicting objective functions that must be optimized simultaneously. A commonly used approach for solving such problems is the weighted sum method, which converts the multi-objective formulation into a single-objective problem by assigning weights to each objective function. However, the effectiveness of this method depends on the careful selection of weights, which relies on empirical judgments, introducing subjectivity and increasing both computational and temporal complexity. An alternative and more structured approach is offered by the ε-constraint method [59], which reformulates the problem into a series of single-objective sub-problems, known as ε-constraint problems. This is achieved by selecting one objective as the primary function while transforming the remaining objectives into constraints, each bounded by an ε-value. By systematically varying the bounds, this method enables the generation of the complete set of non-dominated solutions, known as the Pareto front. Compared to the weighted sum method, the ε-constraint approach overcomes several of its limitations. Specifically, it reduces sensitivity to the scaling of objective functions and provides finer control over the distribution of Pareto-optimal solutions by adjusting the granularity of the constraint bounds.

Without loss of generality, the social and ethical KVI objective is treated as the primary objective in this work. Consequently, the original multi-objective problem in Eq.(16) is reformulated as a single-objective optimization problem, as follows:

$$\max_X V(X), \tag{22}$$

$$\text{s.t.} \quad Q(X) \geq \varepsilon, \tag{23}$$

$$(17), (18), (19), (20), (21). \tag{24}$$

In constraint (23), the parameter ε defines the bound for the objective function $Q(X)$. Specifically, ε varies between the lower and upper bounds of $Q(X)$ over the Pareto-optimal set, which are determined by computing the so-called Nadir and ideal points, respectively. Traditionally, the adjustment of ε follows a uniform partitioning approach, where the interval of ε is divided into sub-intervals of equal size, and each interval limit is selected as a candidate value for ε. However, this method does not guarantee the identification of all non-dominated solutions, as some Pareto-optimal points may be missed due to the fixed discretization scheme. To ensure the generation of the complete Pareto front, the exact ε-constraint approach can be employed, which adaptively refines the selection of ε values. The detailed procedure of this approach is presented in Algorithm 1.

The goal of Algorithm 1 is to identify the set of Pareto-optimal solutions X^*, where no solution can be improved in one objective without compromising the other. Accordingly, the algorithm first determines the range of ε, which serves as the constraint bound for one of the objectives. This range is established by solving two independent single-objective optimization problems. The first optimization considers only $V(X)$ as the objec-

Algorithm 1 Exact ε-Constraint Algorithm

Input: $\mathcal{R}, \mathcal{S}$

Output: Pareto-optimal solutions X^*

Compute the ideal points (V_I, Q_I) with **Algorithm** 2.
Compute the Nadir points (V_N, Q_N), by solving problems 25 and 28 with **Algorithm** 2.
Set $\varepsilon_{\min} = Q_N$, $\varepsilon_{\max} = Q_I$, and step size equal to δ.
for $\varepsilon = \varepsilon_{\min}$ to $\varepsilon_{\max}$ with step δ **do**
    Solve each problem 22 with **Algorithm** 2.
    Return the best-known solution X^*.
    Store the optimal solution if non-dominated.
end for
Stop when $\varepsilon >= \varepsilon_{\max}$.

tive function, solving it without any constraints on $Q(X)$. Similarly, the second optimization is performed with $Q(X)$ as the sole objective, without constraints on $V(X)$. These optimizations yield the ideal points. The solution obtained for the first subproblem is denoted as V_I, representing the ideal point of the function $V(X)$. Similarly, the solution to the second problem is denoted as Q_I, corresponding to the ideal point of the function $Q(X)$, defining the best possible outcomes for each objective when considered independently.

Subsequently, two additional optimization subproblems are solved to determine the Nadir values of both functions, denoted as V_N and Q_N, respectively. These values are obtained by optimizing one objective while imposing a constraint that fixes the other objective at its ideal value. Specifically, V_N is computed by optimizing $V(X)$ under the constraint $Q(X) = Q_I$, while Q_N is obtained by optimizing $Q(X)$ with the constraint $V(X) = V_I$, as formulated below:

$$\max_X V(X), \tag{25}$$

$$\text{s.t.} \quad Q(X) = Q_I, \tag{26}$$

$$(17), (18), (19), (20), (21). \tag{27}$$

$$\max_X Q(X), \tag{28}$$

$$\text{s.t.} \quad V(X) = V_I, \tag{29}$$

$$(17), (18), (19), (20), (21). \tag{30}$$

Once these optimization problems are solved, the lower bound of ε is set to the Nadir value of the constrained objective, while the upper bound is set to its ideal value. To systematically explore the Pareto front, a predefined step size δ is introduced to discretize the range of ε. The choice of δ must balance computational efficiency with solution optimality, ensuring that the algorithm captures non-dominated solutions while minimizing computational overhead. The algorithm then proceeds through an iterative process in which ε is incrementally varied from its lower bound to its upper bound. At each iteration, an optimization problem is formulated and solved, maintaining one objective as the primary function while constraining the second objective using the current ε value. The resulting solution is evaluated for dominance, ensuring that no previously identified solution is strictly superior in all objectives. If the solution is non-dominated, it is added to the Pareto-optimal set. The algorithm continues this iterative procedure until ε exceeds its upper bound, indicating that the Pareto front has been fully explored. Upon termination, the algorithm outputs the Pareto-optimal solutions, which represent the trade-offs between the conflicting objectives.

4.2. Branch-and-Cut solutions and Algorithm analysis

Algorithm 2 Branch and Cut Algorithm

Input: Problem 22

Output: Optimal solution.

Solve the relaxed version of the problem, without integrality constraints.
Set the initial solution space as the relaxed problem's feasible region.
while solution space is not empty **do**
    Partition the solution space into subproblems by branching.
    For each subproblem, compute a lower bound of the objective function.
    if a subproblem has a lower bound greater than or equal to the best-known solution **then**
        Prune the subproblem.
    end if
    Apply cutting planes to tighten the feasible region.
    Solve each remaining subproblem and update the best-known solution.
end while
Return the best-known solution.

Each of the aforementioned problems can be solved iteratively using algorithmic frameworks such as Branch and Bound, which can be further enhanced by integrating optimization techniques like pruning and cutting planes [60]. In general, Branch and Bound is a technique used to solve combinatorial optimization problems by systematically exploring the solution space through a tree structure, where each node represents a subproblem derived from the original problem. The algorithm begins by solving a relaxed version of the problem, typically by ignoring integrality constraints to obtain an upper bound for maximization problems. If the solution to this relaxation is already feasible, it is considered optimal. Otherwise, the algorithm branches by dividing the problem into smaller subproblems with additional constraints (e.g., fixing a variable to 0 or 1). To reduce computational effort, pruning is applied to discard subproblems whose bounds indicate they cannot lead to an optimal solution [61]. Moreover, the technique is further enhanced by incorporating cutting planes, which are additional constraints that tighten the feasible region without excluding optimal integer solutions. This results in the Branch and Cut method [62]. Leveraging this method, if the solution is fractional, additional constraints (i.e., cuts) are introduced to eliminate the current fractional solution while preserving all feasible integer solutions.

If the solution remains fractional, additional cutting planes are introduced. If these cuts fail to yield an integer solution, the standard Branch and Bound branching process is applied. This search process reduces the number of branches explored, enhancing computational efficiency. Details on its implementation are provided in Algorithm 2.

In general, the worst-case time complexity of the pure Branch and Bound algorithm is exponential. However, by incorporating integrality relaxations, pruning, and cutting planes, the average-case complexity can be reduced to $O(n^k)$, where k is a small positive integer dependent on the specific problem instance and the effectiveness of the pruning and cutting techniques employed. For commercial solvers such as Gurobi, the value of k can be as low as 1 or 2. In the case of Algorithm 1, the time complexity is given by $O(4 \cdot (J + N)^k + (Q_I - (Q_N - \delta)) \cdot (J + N)^k)$ which simplifies to $O((Q_I - (Q_N - \delta)) \cdot (J + N)^k)$. Here, $O((J + N)^k)$ represents the time complexity associated with computing the ideal and Nadir points, as well as solving each subproblem for a fixed ε. The term $(Q_I - (Q_N - \delta))$ corresponds to the number of iterations of the inner for-loop. The chosen parameter boundaries represent plausible conditions for distributed orchestration scenarios, wherein each orchestrator manages service requests over a constrained subset of network resources, typically limited to a cluster or a defined geographic area. Preliminary scale-up experiments on instances an order of magnitude larger suggested solver convergence within acceptable tolerances. Owing to space limitations and the paper's emphasis on methodological contributions, a comprehensive large-scale benchmarking campaign is deferred to future work.

5. Numerical Results

This section describes the environmental setup and presents the results of the simulation campaigns conducted using a Python script and the Gurobi commercial solver. The evaluation assesses the effectiveness of the proposed service orchestration model in optimizing KVI, extending beyond traditional performance indicators.

5.1. Environmental setup

The investigated environment consists of N network resources, varying within the range $[50, 800]$. Each resource is characterized by different core availability within $[10, 50]$, with processing capabilities ranging from $[40, 150]$ megacycles per second. The MTTF is assumed to be within $[8760, 45000]$ hours, as reported in [63], while the average number of services a resource can handle per day varies between $[150, 250]$.

Power consumption, expressed in kilowatts, falls within the range $[0.01, 0.2]$. Additionally, carbon credit offsets range between $[4109, 6849.31]$ gCO_2e per day per resource. The CI varies widely across countries, ranging from as low as 12 gCO_2e/kWh in regions powered predominantly by renewables (e.g., Switzerland, Norway) to over 800 gCO_2e/kWh in countries heavily reliant on coal or gas (e.g., Australia, South Africa). Although marginal CI provides a more accurate estimate for assessing the environmental impact of relocating computations, it is often unavailable. Therefore, the average CI is commonly used as a practical lower bound and, in the following evaluations, is set to the global average value of 475 gCO_2e/kWh. The likelihood of a successful cyber attack, denoted as L_{r_n}, varies within $[0.25, 1]$. Finally, the lower and upper bounds of the trustworthiness indicator are set to $L_\theta = 1500$ and $U_\theta = 5000$, respectively, representing the realistic minimum and maximum trust levels that a resource can attain. The inflection point x_0 and the growth parameter B are set to 0.5 and 0.6, respectively, consistent with a normalized confidence domain, ensuring that trust increases most rapidly around the midpoint of the security assurance scale. All the aforementioned parameters are summarized in Table 2.

Table 2: Simulation settings.

Parameter	Setting
J	$[80, 1200]$
N	$[50, 800]$
C_{r_n}	$[10, 50]$
D_{s_j}	$[2, 5]$
co_{r_n}	$[4109, 6849.31]$ $gCOe_2$ per day
CI	$[475]$ gCO_2e/kWh
Clk_{r_n}	$[40, 150]$ Megacycles/s
$P_c^{r_n}, P_m^{r_n}$	$[0.01, 0.2]$ kW
$MTTF_{r_n}$	$[8760, 45000]$ hours
λ_{r_n,s_j}	$[150, 250]$ per day
B_{s_j}	$[600, 1200]$ Mb
Δ_{s_j}	$[0.25, 1]$
L_{r_n}	$[0.25, 1]$
U_θ	5000
L_θ	1500
B	0.6
x_0	0.5

The environment also includes an indefinite number of consumers issuing service requests, which are translated into intents. The number of considered service requests varies within $[80, 1200]$.

To define the intent categories, four classes of 6G services have been investigated: Immersive Experience, Collaborative Robots, Physical Awareness, and Trusted Environments, as detailed in [64].

1. **Immersive Experience** encompasses multimedia and extended reality applications such as immersive telepresence, education, and gaming, requiring ultra-low latency, high data rates, and precise positioning. Meanwhile, its SoS focuses on mitigating the digital divide and reducing energy consumption by optimizing computing and transmission efficiency.
2. **Collaborative Robots** involve autonomous robotic systems operating in industrial and public environments, requiring ultra-reliable, low-latency communication while prioritizing sustainability by minimizing waste and reducing environmental impact.
3. **Physical Awareness** services leverage sensing technologies for applications such as network-assisted mobility and

smart environments, demanding high-resolution sensing, accurate positioning, and high reliability while supporting environmental monitoring, energy-efficient connectivity, and enhanced safety in smart cities.

4. **Trusted Environments** focus on secure and privacy-centric applications, including telemedicine and public safety, necessitating high-end security and extreme reliability, alongside resilient infrastructures that ensure societal well-being.

To quantify the system model parameters, relevant metrics have been analyzed, with values determined based on 6G service requirements and the characteristics described above. The service delay, expressed in seconds, varies within $[0.002, 50]$. The data rate, expressed in bits per second, falls within the range $[70, 250]$. The packet loss rate, representing the number of packets lost per second, is within $[20, 50]$. The service demand is defined within $[2, 5]$, while the input size of data to be processed, denoted by B_{s_j} is sampled from the range $[600, 1200]$ Mb for each service. Moreover, the impact Δ_{s_j} of an attack on the provisioning of the service ranges within $[0.25, 1]$. All specified ranges are defined according to the type of deployment and the characteristics of the entities involved, as referenced in [64].

The performance of the proposed solution has been evaluated against three baseline approaches, each characterized by its respective service orchestration method:

- Random Matching (RM): service requests are assigned to network resources randomly, selecting the first available provider without considering the KPI and KVI offered by a given network resource.
- Performance Greedy Matching (PGM): service requests are assigned to network resources using a greedy strategy, iteratively selecting the most favorable available option based solely on the provided KPIs.
- Value Greedy Matching (VGM): service requests are assigned to network resources using a value-oriented greedy strategy, which iteratively selects the most favorable available option based on the provided KVIs.

5.2. *Pareto Front Analysis*

Figure 2 presents the Pareto fronts representing the solution to the proposed optimization problem, obtained for different combinations of J network services and N available resources.

In this representation, the discretization parameter δ is fixed at 0.1 to ensure consistency and favoring computational efficiency in the generation of the Pareto front. In fact, as explained in Section 4, this approach enables the identification of a sufficient number of Pareto-optimal solutions without incurring excessive computational costs, which would otherwise arise with finer discretizations. Considering the presented solution, no single metric was set to dominate the optimization process; therefore, KPIs and KVIs were assigned equal weight. Similarly, the three KVIs were given equal importance. Additional details regarding the behavior of the approach under varying weight configurations and use cases are presented in the following sections.

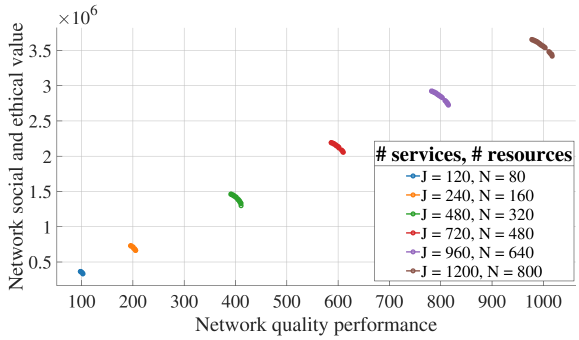

Figure 2: Pareto fronts considering different service requests (J) and network resources (N).

The curves illustrate the trade-off between network quality performance and social and ethical values, highlighting the influence of resource distribution and service requests on the set of optimal solutions. The results indicate that an increase in the number of available resources and service requests leads to an expansion of the Pareto front, thereby allowing for improved trade-offs between the conflicting objectives. For instance, in smaller-scale scenarios ($N = 80, J = 120$), the optimal solutions are constrained within the range $(0.4 \cdot 10^6, 105)$ In contrast, in larger-scale settings ($N = 800, J = 1200$), the Pareto front exhibits a higher density of solutions, reflecting greater flexibility in decision-making. Under these conditions, the network achieves a maximum social and ethical value of $3.6 \cdot 10^6$ while attaining a maximum network quality performance value of 1100. Furthermore, the results underscore the inherent conflict between the two objectives, wherein an improvement in one leads to a corresponding decline in the other.

5.3. *Execution times considerations*

This subsection analyzes the simulation results obtained using the commercial optimization solver Gurobi, focusing on how the construction of the Pareto front varies under different combinations of services and network resources. The analysis incorporates variations in parameters that impact both the ε-constraint method and the previously described system model, including the discretization value δ and the weights assigned to the three KVI, namely environmental sustainability, trustworthiness, and inclusiveness.

Figure 3 examines the effect of varying the number of services J and the discretization parameter δ on the Pareto front, while keeping the number of network resources N fixed at 80. Analyzing the granularity of the solutions, a finer discretization ($\delta = 0.01$) increases the number of optimal trade-offs from 7 to 164 compared to a coarser discretization ($\delta = 1$), offering a more detailed and comprehensive representation of the solution space when J is set to 120. However, this increased solution granularity comes at a significant computational cost, as evidenced by the results in Table 3. The execution time rises from 3.45 s ($\delta = 1, J = 120$) to 20.48 s ($\delta = 0.1, J = 120$) and further to 214.93 s ($\delta = 0.01, J = 120$), demonstrating a substantial increase in computation time when transitioning from $\delta = 1$ to $\delta = 0.01$. The most critical case is observed for $\delta = 0.01, J = 110$, where the execution time reaches 260.43 s, further confirming the substantial increase in computational complexity for finer resolution settings. In contrast, a coarser discretization ($\delta = 1$) significantly reduces computational times, ranging from 3.04 s to 4.08 s. However, this reduction comes at the cost of identifying fewer optimal solutions, potentially overlooking trade-offs between network quality performance and social and ethical values. On the other hand, smaller δ values enhance the optimal solution for the network's social and ethical value by 12%, allowing for a more precise optimization of environmental sustainability, trustworthiness, and inclusiveness.

Table 4 presents the execution times for various scenarios, accounting for different weight assignments to the KVI. This analysis provides insight into the computational impact of prioritizing specific social and ethical dimensions within the optimization process.

Table 3: Simulation run times with varying δ.

N	J	δ	Run time (s)
80	90	1	3.27
80	100	1	3.04
80	110	1	4.08
80	120	1	3.45
80	90	0.1	16.59
80	100	0.1	16.75
80	110	0.1	20.89
80	120	0.1	20.48
80	90	0.01	153.36
80	100	0.01	237.82
80	110	0.01	260.43
80	120	0.01	214.93

Table 4: Simulation run times with varying the number of network resources (N) and the service requests (J).

N	J	Pref. KVI	Run time (s)
80	90	σ_{r_n,s_j}	27.57
80	90	θ_{r_n,s_j}	16.00
80	90	ι_{r_n,s_j}	53.98
80	100	σ_{r_n,s_j}	29.36
80	100	θ_{r_n,s_j}	16.64
80	100	ι_{r_n,s_j}	56.25
80	110	σ_{r_n,s_j}	27.63
80	110	θ_{r_n,s_j}	21.24
80	110	ι_{r_n,s_j}	55.18
80	120	σ_{r_n,s_j}	37.36
80	120	θ_{r_n,s_j}	30.43
80	120	ι_{r_n,s_j}	32.13

In general, the results indicate that execution times remain consistent and within a similar range for all KVIs weight assignments, with only minor fluctuations arising from the different application scenarios considered. Specifically, when $J = 100$ and $N = 80$, execution times range from 16 s to 57 s, with the highest value observed for the inclusiveness indicator. With a slightly higher request load, such as $J = 120$, execution times range from 30.43 s to 37.36 s, with the highest value recorded for the environmental sustainability indicator. Therefore, execution times are primarily influenced by the request load and network resources rather than the weights assigned to the KVIs, as these weights impact the optimization process but do not significantly alter computational complexity across different scenarios.

5.4. KPI and KVI Trade-Off Analysis

Figures 4 and 5 further examine the trade-offs between network quality performance and social and ethical value under varying service-network resource combinations while maintaining a fixed δ parameter, considering a scenario in which neither quality performance nor social and ethical value is explicitly prioritized. The curves provide a comparative analysis of the proposed optimization-based service orchestration framework against baseline approaches, such as the PGM, VGM, and the RM strategy. The results are derived from 500 different seeds, ensuring robustness by accounting for varying distributions of service request types. In detail, Figure 4 provides empirical validation of the theoretical framework introduced in Section 4, demonstrating the system's capability to identify and explore Pareto-optimal solutions as resource availability increases. With a fixed number of service requests ($J = 100$) and a progressively increasing number of available resources (ranging from 20 to 90), the Pareto front shifts, reflecting improvements in both service quality and social and ethical value. This confirms that greater resource diversity facilitates the identification of a broader set of non-dominated configurations and enables more favorable trade-offs between conflicting objectives. This behavior aligns with the core principle of the bi-objective optimization addressed in this work, which jointly considers service demands and network resources to promote the use of ethically valuable resources while satisfying user requirements. In addition, Figures 5a and 5b provide a comparative analysis of network quality performance and the corresponding social and ethical values for scenarios where network resources $N = 80$ and service requests J varies within the range $[80, 120]$. The proposed approach exhibits behavior consistent with the PGM in terms of network quality performance, maintaining coherence in the trend of the curves as the number of considered services increases. A substantial improvement of 10%, instead, is observed in the total network social and ethical value, demonstrating the effectiveness of the proposed framework in integrating social and ethical considerations into service orchestration. Moreover, the approach improves by 13% and 10% over RM in both evaluations, respectively. Furthermore, it maintains a 6% performance gap relative to the VGM in the first case, while demonstrating a reduction of up to 5% in the delivered social and ethical value. Therefore, the proposed strategy outperforms the baseline methods, demonstrating an enhanced capability to allocate resources efficiently while balancing network performance with social and

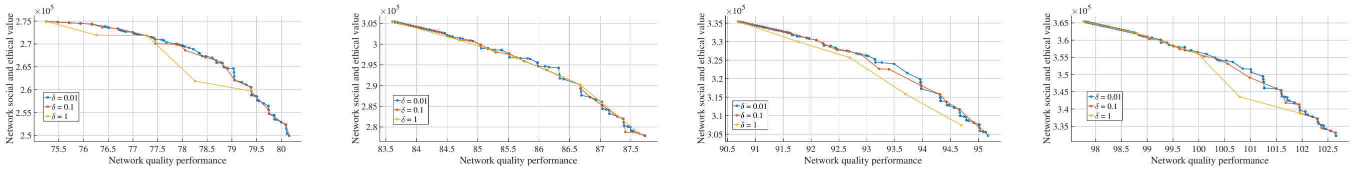

(a) $N = 80, J = 90$ (b) $N = 80, J = 100$ (c) $N = 80, J = 110$ (d) $N = 80, J = 120$

Figure 3: Pareto fronts with varying service requests J and δ.

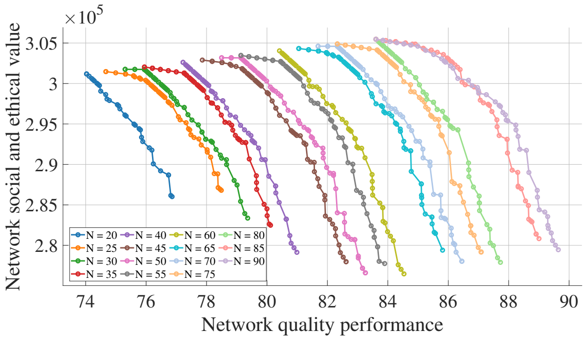

Figure 4: Pareto fronts obtained for varying amounts of network resources, with the number of service requests fixed at $J = 100$.

ethical considerations. This improvement underscores the effectiveness of the approach in optimizing service orchestration beyond traditional performance-centric methods. On the one hand, the greedy approaches achieve optimal results with respect to their respective target criteria but perform suboptimally with regard to the opposing objective. In particular, the PGM employs a purely performance-driven allocation strategy, resulting in higher network quality performance compared to random allocation. However, it fails to integrate social and ethical factors, thereby limiting its improvements to performance-centric metrics while neglecting broader considerations such as environmental sustainability and inclusiveness. In contrast, the VGM attains the highest levels of social and ethical value but demonstrates substantially lower performance when network quality serves as the primary evaluation criterion. On the other hand, the RM demonstrates suboptimal performance due to the absence of an informed resource selection process, leading to inefficient resource-service assignments and reduced overall effectiveness.

Furthermore, Figures 5c and 5d validate the impact of the proposed optimization strategy on both network quality performance and social and ethical value. In this analysis, the number of services is fixed at $J = 100$, while the number of network resources varies within the range $[50, 90]$. Similarly to the previous analysis, the results indicate that the proposed method maintains a performance level comparable to the PGM when evaluating total network quality performance across all resource configurations. However, it achieves up to a 13% improvement over the RM and up to 6% over the VGM. At the same time, the proposed framework surpasses the PGM and RM baseline approaches when considering only the social and ethical value. Specifically, it achieves an increase of up to 20% compared to both the RM and the PGM, resulting in a performance decrease of up to 1% compared to the VGM. Once again, while informed strategies benefit from a larger provider pool by enhancing service–resource matching, the RM approach fails to do so due to its uninformed nature. Conversely, under increasing service demand and fixed resource conditions, RM may occasionally satisfy performance requirements by chance, slightly narrowing the performance gap; however, it consistently remains inferior to informed approaches.

Moreover, Figures 6, 7, and 8 transition the analysis from a general exploration of the trade-off space to specific scenarios, each incorporating realistic KVI and KPI prioritizations that reflect the requirements of their respective 6G service classes. In particular, Figure 6 pertains to the Collaborative Robots and Physical Awareness scenarios, which necessitate prioritization of the environmental sustainability indicator, given the critical importance of minimizing waste and energy consumption, alongside reliability, which demands a minimal packet loss rate. Figure 7, instead, pertains to the Trusted Environments scenario, which requires a stronger emphasis on the trustworthiness indicator due to the paramount importance of security, resilience, and privacy, alongside the need for high data rates and minimal packet loss. Finally, Figure 8 refers to the Immersive Experience scenario, which places greater importance on the inclusiveness indicator, as well as on computation delays and data rates, given the objective of these services to ensure equitable access and bridge the digital divide. In the first scenario, the results indicate that the proposed method achieves a network quality performance that is up to 4% lower than the PGM, and up to 9% higher than the VGM across varying numbers of service requests and network resources. However, it substantially enhances the social and ethical value, surpassing the PGM by up to 13%, and the RM by up to 15% under the same configurations. Notably, the social and ethical value achieved remains approximately 5% lower than that obtained through a purely value-oriented approach. This behavior indicates that resources exhibiting consistently high KPIs or KVIs at the individual level may not constitute a globally optimal allocation when jointly maximizing both categories of indicators, thereby highlighting the existence of locally optimal yet globally suboptimal choices. When prioritizing trustworthiness, the proposed orchestration framework maintains competitive network quality performance, closely aligning with the PGM approach while surpassing the VGM approach by up to 5% and the RM by up to 18%. Moreover, it surpasses the performance-oriented and random baseline strategies in aligning service provisioning with the most suitable and responsible network resources.

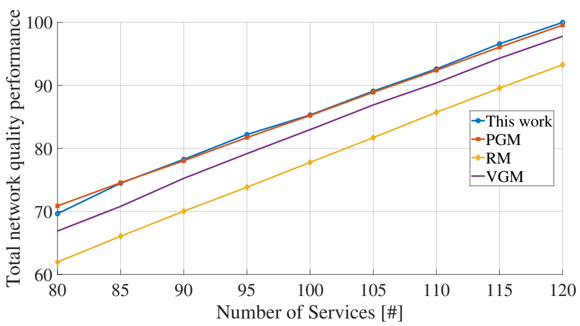

(a) Network quality performance with fixed amount of resources $N = 80$

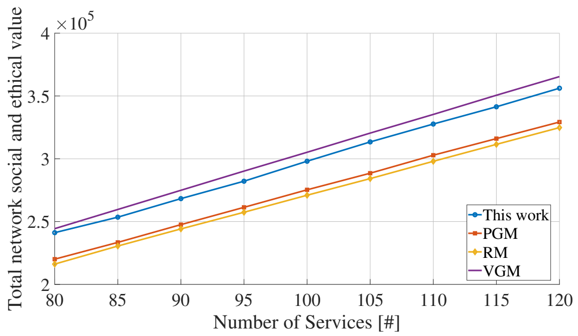

(b) Network social and ethical value with fixed amount of resources $N = 80$

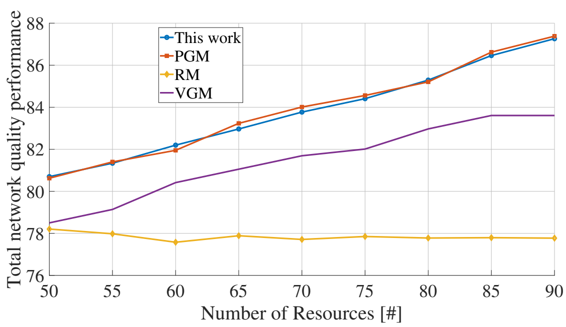

(c) Network quality performance with fixed service requests $J = 100$

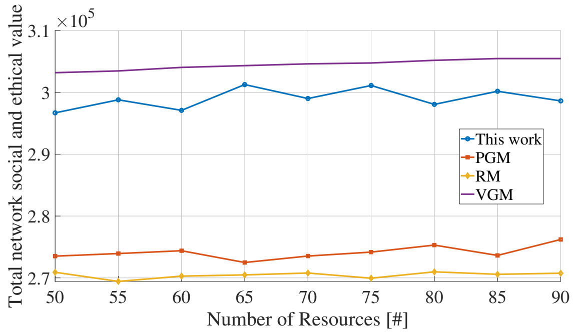

(d) Network social and ethical value with fixed service requests $J = 100$

Figure 5: Network quality performance and social and ethical value considering neither quality nor social and ethical pritiritization.

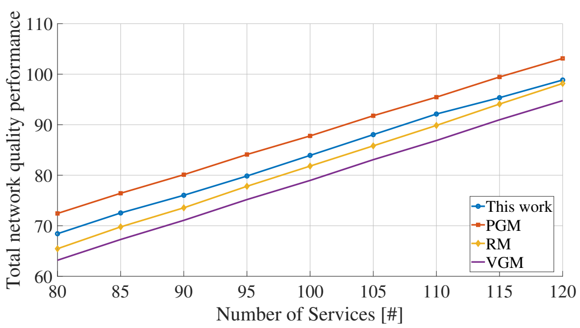

(a) Network quality performance with fixed amount of resources $N = 80$

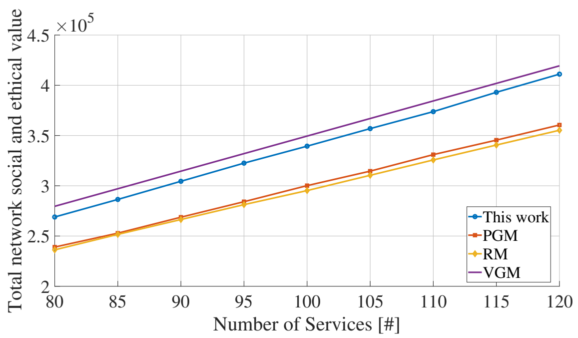

(b) Network social and ethical value with fixed amount of resources $N = 80$

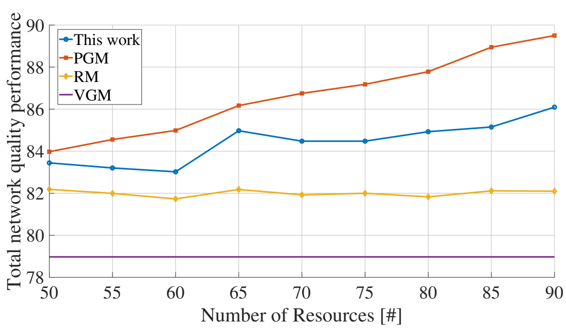

(c) Network quality performance with fixed service requests $J = 100$

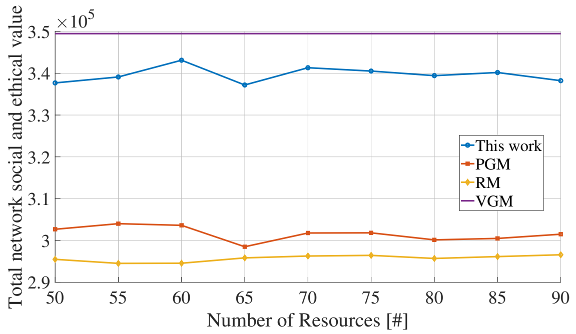

(d) Network social and ethical value with fixed service requests $J = 100$

Figure 6: Network quality performance and social and ethical value considering an environmental sustainability focus.

Specifically, in terms of social and ethical value, the proposed method outperforms both the RM and PGM by approximately up to 4% when $J = 100$, all while consistently matching the VGM, demonstrating its effectiveness in ensuring KVIs adherence within service orchestration, and effectively integrating trustworthiness into resource allocation without compromising network performance, as opposed to the VGM baseline.

Furthermore, while prioritizing inclusiveness leads to only a marginal degradation in network quality performance across varying numbers of service requests and network resources, the

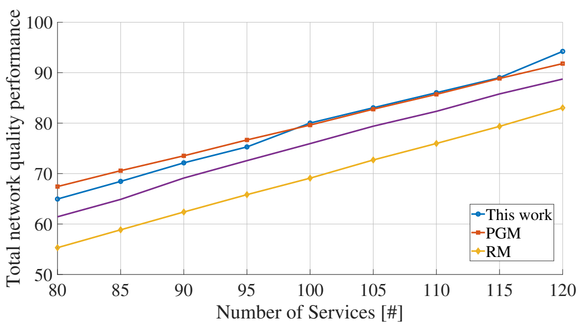

(a) Network quality performance with fixed amount of resources $N = 80$

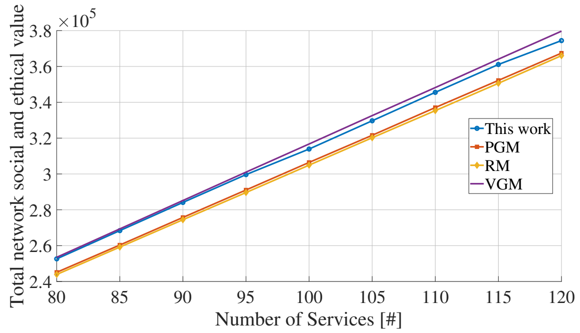

(b) Network social and ethical value with fixed amount of resources $N = 80$

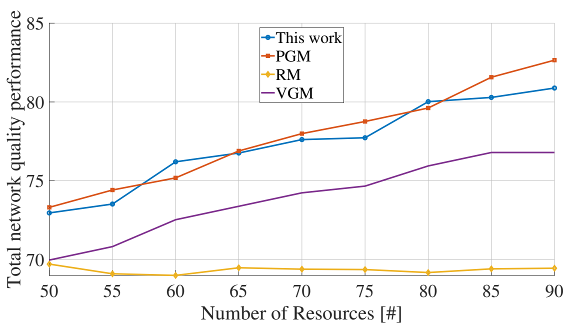

(c) Network quality performance with fixed service requests $J = 100$

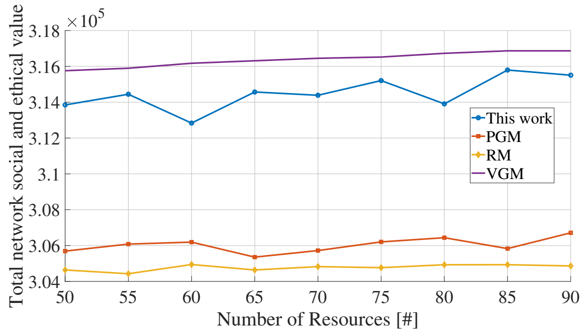

(d) Network social and ethical value with fixed service requests $J = 100$

Figure 7: Network quality performance and social and ethical value considering trustworthiness focus.

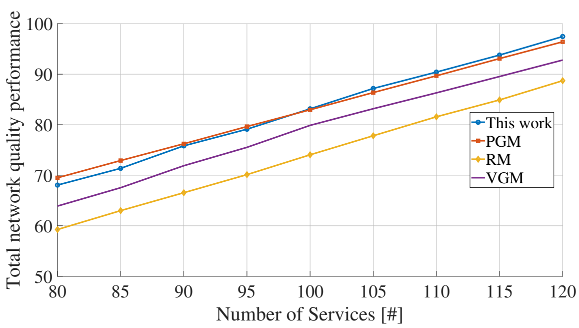

(a) Network quality performance with fixed amount of resources $N = 80$

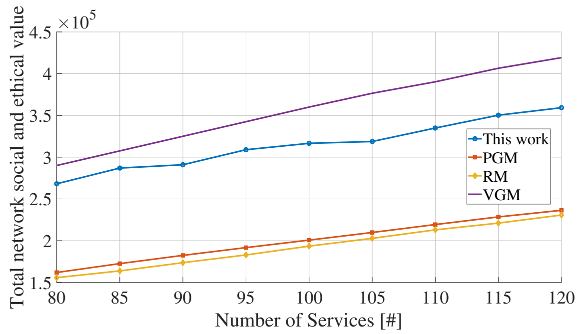

(b) Network social and ethical value with fixed amount of resources $N = 80$

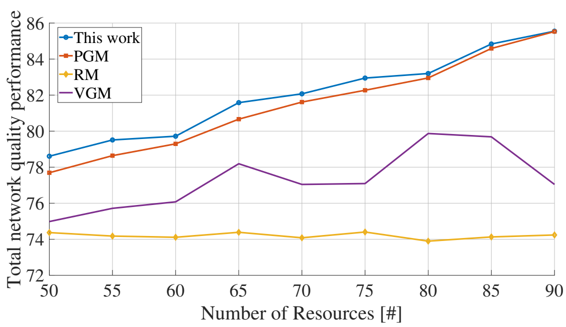

(c) Network quality performance with fixed service requests $J = 100$

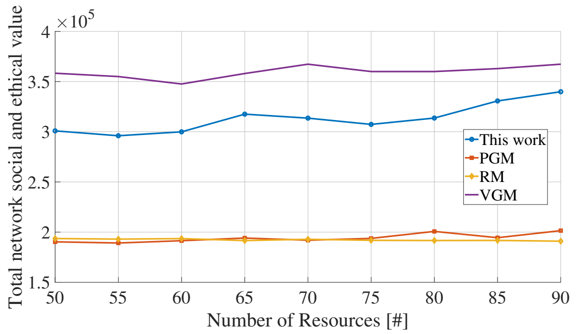

(d) Network social and ethical value with fixed service requests $J = 100$

Figure 8: Network quality performance and social and ethical value considering inclusiveness focus.

proposed approach yields a substantial improvement in social and ethical value. Specifically, for $J = 100$, it achieves an increase exceeding 70% compared to both the PGM and the RM, underscoring its effectiveness in integrating social considerations into network management. Moreover, the proposed approach exhibits a network quality performance gap of up to 12% relative to the VGM. Nonetheless, both approaches demonstrate closely aligned outcomes in terms of social and ethical value under identical conditions, with differences generally within 10% relative to the same baseline. Overall, the proposed approach achieves substantial improvements in the social and ethical value delivered by the network during service orchestration across all evaluated scenarios. This contrasts with value-oriented resource allocation strategies, which prioritize resources with the highest KVIs per request but often fail to achieve globally optimal trade-offs. The performance degradation occasionally observed in the proposed approach relative to the PGM is consistently offset by significant gains in social and ethical value. These results confirm its effectiveness in managing complex service-resource allocations while ensuring a balanced trade-off between network performance and key societal considerations.

6. Conclusions

The proposed work introduced a system design framework for formalizing and evaluating KVIs alongside KPIs, enabling the systematic monitoring and assessment of communication networks' impact on SDGs. This has been achieved through a service orchestration solution that optimizes metrics extending beyond traditional performance indicators. By integrating the IBN paradigm, the framework embeds a social dimension directly into the service delivery process. To balance network performance with social and ethical considerations, a bi-objective optimization problem has been formulated, aligning service provisioning with the most suitable and responsible network resources. The exact ε-constraint method has been employed to solve this problem, transforming the multi-objective formulation into a series of single-objective problems to identify the set of optimal solutions. Computer simulations demonstrate up to a 70% increase in the social and ethical value delivered by the network in service orchestration compared to a baseline approach that focuses solely on network performance parameters. This improvement remains consistent across various scenarios involving 6G services with diverse requirements for sustainability, trustworthiness, and inclusiveness. Future research will investigate the effectiveness of the approach through experimental testbeds and more complex scenarios involving a broader range of KPIs and KVIs. Additionally, the intent translation module, along with a more targeted LLM customization and alignment with human values and preferences, will be implemented to further enhance the proposed orchestration model.

Acknowledgements

This work was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, in the context of partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART", CUP: D93C22000910001), national center on "Sustainable Mobility" (CN00000023 - program "MOST", CUP: D93C22000410001), and partnership on "Cybersecurity" (PE00000007 - program "SERICS", CUP: D33C22001300002, project ISP5G+). It was also supported by the PRIN 2022 projects INSPIRE (grant no. 2022BEXMXN_01) and HORUS (grant no. 2022P44KA8) funded by the Italian MUR, and by the HORIZON MSCA project BRIDGITISE (grant no. 101119554).

References

[1] N. United, Sustainable development goals. transforming our world: the 2030 agenda for sustainable development, Tech. rep., United Nations, New York (2015).
URL https://sustainabledevelopment.un.org/topics

[2] L. Mohjazi, B. Selim, M. Tatipamula, M. A. Imran, The journey toward 6g: A digital and societal revolution in the making, IEEE Internet of Things Magazine 7 (2) (2024) 119–128. doi:10.1109/IOTM.001.2300119.

[3] C.-X. Wang, X. You, X. Gao, X. Zhu, Z. Li, C. Zhang, H. Wang, Y. Huang, Y. Chen, H. Haas, J. S. Thompson, E. G. Larsson, M. D. Renzo, W. Tong, P. Zhu, X. Shen, H. V. Poor, L. Hanzo, On the road to 6g: Visions, requirements, key technologies, and testbeds, IEEE Communications Surveys & Tutorials 25 (2) (2023) 905–974.

[4] K. Ojutkangas, E. Rossi, M. Matinmikko-Blue, A deep dive into the birth process of linking 6g and the un sdgs, Telecommunications Policy 46 (1) (2022) 102283. doi:https://doi.org/10.1016/j.telpol.2021.102283.
URL https://www.sciencedirect.com/science/article/pii/S0308596121001877

[5] C. D. Alwis, A. Kalla, Q.-V. Pham, P. Kumar, K. Dev, W.-J. Hwang, M. Liyanage, Survey on 6g frontiers: Trends, applications, requirements, technologies and future research, IEEE Open Journal of the Communications Society 2 (2021) 836–886. doi:10.1109/OJCOMS.2021.3071496.

[6] L. Pintor, L. Atzori, A. Iera, Sustainability in telecommunication networks and key value indicators: a survey (2025). arXiv:2501.04356.
URL https://arxiv.org/abs/2501.04356

[7] L. Atzori, C. Campolo, A. Iera, G. Morabito, Toward the ethicnet: Challenges and enablers for ethics-aware networks, IEEE Communications Magazine 61 (11) (2023) 192–198. doi:10.1109/MCOM.004.2300110.

[8] K. Chen, Y. Sun, S. Zheng, H. Yang, P. Yu, Online collaborative energy-network resource scheduling for wpt-enabled green edge computing, IEEE Transactions on Green Communications and Networking 8 (2) (2024) 601–618. doi:10.1109/TGCN.2023.3339477.

[9] A. Balakrishnan, S. De, L.-C. Wang, Network operator revenue maximization in dual powered green cellular networks, IEEE Transactions on Green Communications and Networking 5 (4) (2021) 1791–1805. doi:10.1109/TGCN.2021.3094164.

[10] Z. Garroussi, A. W. Badirou, M. D'Amours, A. Girard, B. Sansò, On the economic viability of solar energy when upgrading cellular networks, IEEE Transactions on Green Communications and Networking 7 (2) (2023) 744–758. doi:10.1109/TGCN.2022.3223622.

[11] N. Hu, Z. Tian, X. Du, M. Guizani, An energy-efficient in-network computing paradigm for 6g, IEEE Transactions on Green Communications and Networking 5 (4) (2021) 1722–1733. doi:10.1109/TGCN.2021.3099804.

[12] Z. Zhou, M. Shojafar, M. Alazab, J. Abawajy, F. Li, Afed-ef: An energy-efficient vm allocation algorithm for iot applications in a cloud data center, IEEE Transactions on Green Communications and Networking 5 (2) (2021) 658–669. doi:10.1109/TGCN.2021.3067309.

[13] E. Tong, W. Niu, Y. Tian, J. Liu, T. Baker, S. Verma, Z. Liu, A hierarchical energy-efficient service selection approach with qos constraints for internet of things, IEEE Transactions on Green Communications and Networking 5 (2) (2021) 645–657. doi:10.1109/TGCN.2021.3069121.

[14] Y. Huang, H. Xu, H. Gao, X. Ma, W. Hussain, Ssur: An approach to optimizing virtual machine allocation strategy based on user requirements for cloud data center, IEEE Transactions on Green Communications and Networking 5 (2) (2021) 670–681. doi:10.1109/TGCN.2021.3067374.

[15] T. Ding, L. Liu, Z. Yan, L. Cui, Umix: Sustainable multi–uav coordination for aerial-terrestrial networks, IEEE Transactions on Network Science and Engineering 11 (6) (2024) 5190–5200. doi:10.1109/TNSE.2024.3361731.

[16] A. Awad Abdellatif, A. Abo-Eleneen, A. Mohamed, A. Erbad, N. V. Navkar, M. Guizani, Intelligent-slicing: An ai-assisted network slicing framework for 5g-and-beyond networks, IEEE Transactions on Network

and Service Management 20 (2) (2023) 1024–1039. doi:10.1109/TNSM.2023.3274236.

[17] W. Yang, X. Chi, L. Zhao, Z. Xiong, W. Jiang, Task-driven semantic-aware green cooperative transmission strategy for vehicular networks, IEEE Transactions on Communications 71 (10) (2023) 5783–5798. doi:10.1109/TCOMM.2023.3300869.

[18] N. Alliance, Green Future Networks Sustainability Challenges and Initiatives in Mobile Networks, Tech. rep., NGMN Alliance (07 2021).

[19] G. Wikström, N. Bledow, M. Matinmikko-Blue, H. Breuer, C. Costa, G. Darzanos, A. Gavras, T. Hossfeld, I. Mesogiti, K. Petersen, P. Porambage, R.-A. Stoica, S. Wunderer, Key value indicators: A framework for values-driven next-generation ict solutions, Telecommunications Policy 48 (6) (2024) 102778. doi:https://doi.org/10.1016/j.telpol.2024.102778.
URL https://www.sciencedirect.com/science/article/pii/S0308596124000752

[20] G. I. Association, et al., What societal values will 6g address? 6g infrastructure association vision and societal challenges working group, Tech. rep., 6G SNS IA (2022).

[21] B. Masood Khorsandi, et al., Expanded 6g vision, use cases and societal values – including aspects of sustainability, security and spectrum, Deliverable 1.1, A flagship for B5G6G vision and intelligent fabric of technology enablers connecting human, physical, and digital worlds (2021).
URL https://hexa-x.eu/deliverables/

[22] A. Leivadeas, M. Falkner, A Survey on Intent-Based Networking, IEEE Communications Surveys & Tutorials 25 (1) (2023) 625–655. doi:10.1109/COMST.2022.3215919.

[23] ITU-T, Intent-based network management and orchestration for network slicing in IMT-2020 networks and beyond, Recommendation Y.3161, ITU (12 2023).

[24] J. Navarro-Ortiz, P. Romero-Diaz, S. Sendra, P. Ameigeiras, J. J. Ramos-Munoz, J. M. Lopez-Soler, A survey on 5g usage scenarios and traffic models, IEEE Communications Surveys & Tutorials 22 (2) (2020) 905–929. doi:10.1109/COMST.2020.2971781.

[25] M. Liu, G. Feng, Y. Sun, N. Chen, W. Tan, A network function parallelism-enabled mec framework for supporting low-latency services, IEEE Transactions on Services Computing 16 (1) (2023) 40–52. doi:10.1109/TSC.2021.3130247.

[26] K. Dzeparoska, A. Tizghadam, A. Leon-Garcia, Emergence: An intent fulfillment system, IEEE Communications Magazine 62 (6) (2024) 36–41. doi:10.1109/MCOM.001.2300270.

[27] K. Dzeparoska, J. Lin, A. Tizghadam, A. Leon-Garcia, Llm-based policy generation for intent-based management of applications, in: 2023 19th International Conference on Network and Service Management (CNSM), Vol. 33, IEEE, 2023, p. 1–7. doi:10.23919/cnsm59352.2023.10327837.
URL http://dx.doi.org/10.23919/CNSM59352.2023.10327837

[28] C. Wang, M. Scazzariello, A. Farshin, S. Ferlin, D. Kostić, M. Chiesa, Netconfeval: Can llms facilitate network configuration?, Proc. ACM Netw. 2 (CoNEXT2) (Jun. 2024). doi:10.1145/3656296.
URL https://doi.org/10.1145/3656296

[29] J. Lin, K. Dzeparoska, A. Tizghadam, A. Leon-Garcia, AppleSeed: Intent-Based Multi-Domain Infrastructure Management via Few-Shot Learning, in: 2023 IEEE 9th International Conference on Network Softwarization (NetSoft), 2023, pp. 539–544. doi:10.1109/NetSoft57336.2023.10175410.

[30] A. S. Jacobs, R. J. Pfitscher, R. H. Ribeiro, R. A. Ferreira, L. Z. Granville, W. Willinger, S. G. Rao, Hey, Lumi! Using Natural Language for Intent-Based Network Management, in: 2021 USENIX Annual Technical Conference (USENIX ATC 21), USENIX Association, 2021, pp. 625–639.

[31] E. El-Rif, A. Leivadeas, M. Falkner, Intent Expression Through Natural Language Processing in an Enterprise Network, in: 2023 IEEE 24th International Conference on High Performance Switching and Routing (HPSR), 2023, pp. 1–6. doi:10.1109/HPSR57248.2023.10148046.

[32] R. Mondal, A. Tang, R. Beckett, T. Millstein, G. Varghese, What do llms need to synthesize correct router configurations?, in: Proceedings of the 22nd ACM Workshop on Hot Topics in Networks, HotNets '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 189–195. doi:10.1145/3626111.3628194.
URL https://doi.org/10.1145/3626111.3628194

[33] C. Wang, M. Scazzariello, A. Farshin, D. Kostic, M. Chiesa, Making network configuration human friendly (2023). arXiv:2309.06342.
URL https://arxiv.org/abs/2309.06342

[34] F. Settanni, A. Zamponi, C. Basile, Dynamic security provisioning for cloud-native networks: An intent-based approach, in: 2024 IEEE International Conference on Cyber Security and Resilience (CSR), 2024, pp. 321–328. doi:10.1109/CSR61664.2024.10679397.

[35] M. Kiran, E. Pouyoul, A. Mercian, B. Tierney, C. Guok, I. Monga, Enabling intent to configure scientific networks for high performance demands, Future Generation Computer Systems 79 (2018) 205–214. doi:https://doi.org/10.1016/j.future.2017.04.020.
URL https://www.sciencedirect.com/science/article/pii/S0167739X1730626X

[36] C. V. Nahum, V. H. L. Lopes, R. M. Dreifuerst, P. Batista, I. Correa, K. V. Cardoso, A. Klautau, R. W. Heath, Intent-aware radio resource scheduling in a ran slicing scenario using reinforcement learning, IEEE Transactions on Wireless Communications 23 (3) (2024) 2253–2267. doi:10.1109/TWC.2023.3297014.

[37] E. J. Scheid, C. C. Machado, M. F. Franco, R. L. dos Santos, R. P. Pfitscher, A. E. Schaeffer-Filho, L. Z. Granville, INSpIRE: Integrated NFV-based Intent Refinement Environment, in: 2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM), 2017, pp. 186–194. doi:10.23919/INM.2017.7987279.

[38] G. Landeau, M. Avgeris, A. Leivadeas, I. Lambadaris, Security-Oriented Network Intent Placement using Particle Swarm Optimization, in: 2023 7th Cyber Security in Networking Conference (CSNet), 2023, pp. 19–22. doi:10.1109/CSNet59123.2023.10339783.

[39] I. Chatzistefanidis, A. Leone, N. Nikaein, Maestro: Llm-driven collaborative automation of intent-based 6g networks, IEEE Networking Letters (2024) 1–1doi:10.1109/LNET.2024.3503292.

[40] K. Abbas, A. Nauman, M. Bilal, J.-H. Yoo, J. W.-K. Hong, W.-C. Song, Ai-driven data analytics and intent-based networking for orchestration and control of b5g consumer electronics services, IEEE Transactions on Consumer Electronics 70 (1) (2024) 2155–2169. doi:10.1109/TCE.2023.3324010.

[41] F. de Trizio, G. Sciddurlo, D. Rutigliano, G. Piro, G. Boggia, A novel malicious intent detection approach in Intent-Based enterprise networks, in: 1st International wokshop on Network Security Operations (NeSecOr) 2024 (NeSecOr 2024), Prague, Czech Republic, 2024, p. 7.

[42] F. de Trizio, G. Sciddurlo, I. Cianci, G. Piro, G. Boggia, Optimizing key value indicators in intent-based networks through digital twins aided service orchestration mechanisms, Computer Communications 228 (2024) 107977. doi:https://doi.org/10.1016/j.comcom.2024.107977.
URL https://www.sciencedirect.com/science/article/pii/S0140366424003244

[43] A. Clemm, L. Ciavaglia, L. Z. Granville, J. Tantsura, Intent-Based Networking - Concepts and Definitions, RFC 9315 (oct 2022). doi:10.17487/RFC9315.
URL https://www.rfc-editor.org/info/rfc9315

[44] K. Mehmood, K. Kralevska, D. Palma, Intent-driven autonomous network and service management in future cellular networks: A structured literature review, Computer Networks 220 (2023) 109477. doi:https://doi.org/10.1016/j.comnet.2022.109477.
URL https://www.sciencedirect.com/science/article/pii/S1389128622005114

[45] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, Y. Du, C. Yang, Y. Chen, Z. Chen, J. Jiang, R. Ren, Y. Li, X. Tang, Z. Liu, P. Liu, J.-Y. Nie, J.-R. Wen, A survey of large language models (2024). arXiv:2303.18223.
URL https://arxiv.org/abs/2303.18223

[46] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need (2023). arXiv:1706.03762.
URL https://arxiv.org/abs/1706.03762

[47] Y. Bengio, G. Hinton, A. Yao, D. Song, P. Abbeel, T. Darrell, Y. N. Harari, Y.-Q. Zhang, L. Xue, S. Shalev-Shwartz, G. Hadfield, J. Clune, T. Maharaj, F. Hutter, A. G. Baydin, S. McIlraith, Q. Gao, A. Acharya, D. Krueger, A. Dragan, P. Torr, S. Russell, D. Kahneman, J. Brauner, S. Mindermann, Managing extreme ai risks amid rapid progress, Science 384 (6698) (2024) 842–845. arXiv:https://www.science.org/doi/pdf/10.1126/science.adn0117,

doi:10.1126/science.adn0117.
URL https://www.science.org/doi/abs/10.1126/science.adn0117

[48] T. Xiao, J. Zhu, Foundations of large language models (2025). arXiv:2501.09223.
URL https://arxiv.org/abs/2501.09223

[49] H. Wymeersch, H. Chen, H. Guo, M. Furkan Keskin, B. M. Khorsandi, M. H. Moghaddam, A. Ramirez, K. Schindhelm, A. Stavridis, T. Svensson, V. Yajnanarayana, 6G Positioning and Sensing Through the Lens of Sustainability, Inclusiveness, and Trustworthiness, arXiv e-prints (2023). doi:10.48550/arXiv.2309.13602.

[50] H. Wang, C. Yu, L. Wang, Q. Yu, Effective bigdata-space service selection over trust and heterogeneous qos preferences, IEEE Transactions on Services Computing 11 (4) (2018) 644–657. doi:10.1109/TSC.2015.2480393.

[51] I. O. for Standardization, Greenhouse Gases: Carbon Footprint of Products: Requirements and Guidelines for Quantification, International Organization for Standardization, 2018.

[52] A. Clemm, C. Pignataro, E. Schooler, L. Ciavaglia, A. Rezaki, G. Mirsky, J. Tantsura, Green Networking Metrics for Environmentally Sustainable Networking, Internet-Draft draft-cx-green-green-metrics-00, Internet Engineering Task Force, work in Progress (Oct. 2024).
URL https://datatracker.ietf.org/doc/draft-cx-green-green-metrics/00/

[53] G. Zhang, F. Shen, Y. Zhang, R. Yang, Y. Yang, E. A. Jorswieck, Delay minimized task scheduling in fog-enabled iot networks, in: 2018 10th International Conference on Wireless Communications and Signal Processing (WCSP), 2018, pp. 1–6. doi:10.1109/WCSP.2018.8555532.

[54] L. Lannelongue, J. Grealey, M. Inouye, Green algorithms: Quantifying the carbon footprint of computation, Advanced Science 8 (12) (2021) 2100707. doi:https://doi.org/10.1002/advs.202100707.

[55] M. Battaglioni, G. Rafaiani, F. Chiaraluce, M. Baldi, Magic: A method for assessing cyber incidents occurrence, IEEE Access 10 (2022) 73458–73473. doi:10.1109/ACCESS.2022.3189777.

[56] S. Sahni, T. Gonzalez, P-complete approximation problems, J. ACM 23 (3) (1976) 555–565. doi:10.1145/321958.321975.
URL https://doi.org/10.1145/321958.321975

[57] Y. H. Yv, L. S. Lasdon, D. Da, On a bicriterion formulation of the problems of integrated system identification and system optimization, IEEE Transactions on Systems, Man, and Cybernetics SMC-1 (3) (1971) 296–297. doi:10.1109/TSMC.1971.4308298.

[58] G. Mavrotas, Effective implementation of the ε-constraint method in multi-objective mathematical programming problems, Applied Mathematics and Computation 213 (2) (2009) 455–465. doi:https://doi.org/10.1016/j.amc.2009.03.037.
URL https://www.sciencedirect.com/science/article/pii/S0096300309002574

[59] J.-F. Bérubé, M. Gendreau, J.-Y. Potvin, An exact ε-constraint method for bi-objective combinatorial optimization problems: Application to the traveling salesman problem with profits, European Journal of Operational Research 194 (1) (2009) 39–50. doi:https://doi.org/10.1016/j.ejor.2007.12.014.
URL https://www.sciencedirect.com/science/article/pii/S0377221707012106

[60] T. Denat, A. Harutyunyan, N. Melissinos, V. T. Paschos, Average-case complexity of a branch-and-bound algorithm for min dominating set, Discrete Applied Mathematics 345 (2024) 4–8. doi:https://doi.org/10.1016/j.dam.2023.11.021.
URL https://www.sciencedirect.com/science/article/pii/S0166218X23004353

[61] A. Basu, M. Conforti, M. Di Summa, H. Jiang, Complexity of branch-and-bound and cutting planes in mixed-integer optimization, Mathematical Programming 198 (1) (2023) 787–810. doi:10.1007/s10107-022-01789-5.
URL https://doi.org/10.1007/s10107-022-01789-5

[62] M.-F. Balcan, S. Prasad, T. Sandholm, E. Vitercik, Improved sample complexity bounds for branch-and-cut (2022). arXiv:2111.11207.
URL https://arxiv.org/abs/2111.11207

[63] N. Milanovic, B. Milic, Automatic generation of service availability models, IEEE Transactions on Services Computing 4 (1) (2011) 56–69. doi:10.1109/TSC.2010.11.

[64] S. Roessel, C. Drewes, et al., 6g use cases and requirements, Deliverable 1.2, SNS-JU, Horizon, EU (2023).
URL https://hexa-x-ii.eu/results/