telecom

MDPI

Review

Resource Allocation and Optimization for 6G Wireless Networks: Requirements, Design Principles, and Open Challenges

Gilang Raka Rayuda Dewa [1,*] and Illsoo Sohn [2]

[1] Department of Computer Science, Sampoerna University, Jakarta 12780, Indonesia
[2] Department of Computer Science and Engineering, Seoul National University of Science and Technology, Seoul 01811, Republic of Korea; isohn@seoultech.ac.kr
* Correspondence: gilang.dewa@sampoernauniversity.ac.id

Abstract

Resource allocation plays a crucial role in determining the overall performance of 6G wireless networks; however, the growing number of users and the integration of heterogeneous infrastructures make this task increasingly challenging. Even with higher capacity and wider coverage, spectrum and energy resources remain limited. Furthermore, different network infrastructures operate under rapidly changing wireless conditions. Therefore, achieving the expected 6G network performance also depends on allocating these limited resources based on current network conditions. This requirement has led to the development of various optimization methods for wireless-resource allocation. However, the performance improvements reported in existing studies do not necessarily establish whether the corresponding methods can operate under 6G networks. Most evaluations emphasize the resulting allocation quality while providing limited evidence on whether the complete decision process remains efficient under changing network conditions. This paper therefore evaluates the operational evidence reported for centralized, distributed, and hybrid resource allocation against criteria derived from 6G requirements. The assessment of several state-of-the-art solutions shows that solution quality is consistently demonstrated; however, evidence concerning decision timeliness, signaling overhead, and resource-management energy remains limited. Based on these findings, the paper derives four operational design principles: validity-aware control, efficient coordination, feasibility-preserving adaptation, and net-benefit management. Accordingly, it identifies open challenges for practical 6G resource allocation.

Keywords: 6G networks; wireless; resource allocation; optimization; requirements

Check for updates

Academic Editor: Barbara M. Masini

Received: 31 August 2026
Revised: 29 October 2026
Accepted: 5 November 2026
Published: 9 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Wireless communication networks have evolved from systems mainly designed for voice and broadband access into an established infrastructure for transportation, healthcare, manufacturing, and public services. Fifth-generation (5G) networks have supported this development via enhanced mobile broadband (eMBB), massive machine-type communications (mMTC), and ultra-reliable low-latency communications (URLLC) [1]. However, the increasing number of users has introduced higher performance demands [2,3]. These demands are expected to increase as wireless networks become more closely integrated with computing, sensing, and autonomous control [4,5].

The International Telecommunication Union (ITU) introduced the IMT-2030 framework to describe the expected usage scenarios and capabilities of future mobile communication systems [4]. Unlike previous generations, the sixth generation (6G) is expected

to support services in numerous environments, including terrestrial, aerial, maritime, and remote environments. This development requires integrating several heterogeneous nodes, e.g., terrestrial base stations, edge servers, satellites, aerial platforms, and others [6,7]. Although this integration potentially extends network coverage, it also increases the number of involved nodes. For example, a LEO-based network may involve satellites, ground stations, gateways, and terrestrial network elements, adding complexity to user association and handover decisions [8,9].

Based on this framework, network performance depends heavily on how resources are managed. Resource allocation determines how channels, transmission power, computing capacity, storage, energy, and other wireless parameters are assigned among network users [2,3,10]. The main decisions include channel allocation, power control, beam selection, user association, scheduling, routing, handover management, computation offloading, and network slicing. Each decision must consider the available resources and the quality-of-service (QoS) requirements of different applications.

Resource allocation decisions are typically interdependent with other network control decisions; changes in one decision can alter the network's requirements. For example, a handover decision changes the serving node and may require a new allocation of channels, power, and computational resources. These relationships become more difficult to manage as the number of users, network nodes, and decision variables increases. The resulting allocation problem can be nonconvex, and finding a globally optimal decision may require substantial computation; practical methods therefore often seek a feasible, high-quality approximation within the available decision time [5,11]. With high computational complexity, network topology may change before a calculated allocation can be implemented.

Therefore, the suitability of a resource allocation method cannot be decided only from the objective value. The time required to calculate the solution must be considered concurrently with the amount of information exchange required during the decision process. A method that provides a highly accurate solution may still be unsuitable when its computational time exceeds the available decision time. These implementation requirements become important in 6G networks because resource decisions may involve large numbers of users and nodes.

Various optimization methods have been developed to address wireless-resource allocation problems. Mathematical optimization explicitly formulates the objective functions, decision variables, and system constraints; however, its computational requirements may increase with the number of involved nodes [11]. Distributed optimization splits the problem among network nodes and reduces dependence on a central coordinator. However, it may require iterative information exchange and sufficient time to converge [9,12,13]. Learning-assisted optimization uses previous observations and network interactions to estimate suitable resource decisions [14,15]. Model-driven and graph-based methods further combine network models or relationships among network nodes into the learning process [16,17]. These methods provide different advantages, but their suitability depends on the resource constraints, network conditions, and implementation requirements.

Several surveys have discussed resource allocation and optimization for future wireless networks. Alhashimi et al. examined resource allocation strategies for heterogeneous 6G networks [2]. Sefati et al. focused on resource allocation for IoT-based 6G networks [3]. Other studies have discussed machine learning for large-scale wireless optimization and AI-enabled resource allocation methods [11,15]. More recently, de Faria et al. presented a broader survey of radio-resource allocation for 6G networks [10]. These studies provide important discussions of their respective applications, resources, and optimization methods.

However, existing survey studies are mainly organized by network architecture, application type, resource type, and optimization technique. As a result, methods developed for

different network scenarios are difficult to compare using the same basis. In addition, many relevant methods were originally developed for current or emerging wireless systems and were not specifically identified as 6G network methods. Their relevance to 6G networks should therefore be determined from their operating characteristics. This requires an analysis that starts from the requirements expected for 6G networks and then examines how existing methods can support them.

Table 1 presents the main difference between existing surveys and this paper. This paper synthesizes existing findings to connect 6G requirements with resource allocation problems, decision architectures, optimization mechanisms, and implementation conditions. This structure aims to identify the properties required of resource allocation methods.

Figure 1 summarizes the analytical framework of this paper. The framework first translates the expected 6G capability indicators into resource allocation requirements. It then relates these requirements to the corresponding resource allocation techniques. Here, this paper also examines decision-making architectures and optimization methods to address different resource allocation scenarios. Finally, the methods are assessed using evaluation criteria related to solution quality, decision timeliness, information exchange, constraint handling, and scalability. This structure allows methods developed for different wireless scenarios to be examined from the same 6G perspective.

Table 1. Comparison of existing surveys and this paper.

Paper	Main Scope	Multi-Resource	Methods and Architectures	6G Requirement Mapping	Design Principles
Alhashimi et al. [2]	Resource allocation in heterogeneous 6G networks	△	△	△	–
Sefati et al. [3]	Resource allocation in IoT-based 6G networks	✓	△	△	–
Shi et al. [11]	Learning methods for large-scale wireless optimization	△	△	△	–
Alhashimi et al. [15]	AI-enabled resource allocation for 6G networks	✓	△	△	–
de Faria et al. [10]	Radio-resource allocation for 6G networks	–	✓	△	–
This paper	Resource allocation and optimization evaluated from 6G requirements	✓	✓	✓	✓

Note: ✓ indicates explicit coverage, △ indicates partial coverage, and – indicates that the topic is not a primary focus.

This paper is a requirements-based review. It addresses three questions: (RQ1) How can IMT-2030 capabilities be translated into operational requirements for resource allocation? (RQ2) Which allocation and implementation properties are supported by the reported evidence in representative primary studies? (RQ3) What design guidance follows from properties that remain unsupported or incompletely reported? The reviewed applications include radio-resource allocation, vehicular communication, industrial control, edge inference, and computation offloading. Their numerical objectives are not pooled because the

underlying system models, constraints, and baselines differ. Instead, all selected studies are examined using the same evidence criteria.

Accordingly, the main contributions of this paper are summarized as follows.

- **Requirements-based evaluation framework:** The expected capabilities of 6G networks are translated into eight criteria for evaluating both the resulting allocation and the process used to produce it. This framework provides a common basis for examining methods developed for different wireless environments.
- **Implementation-oriented method synthesis:** Resource allocation methods are organized into centralized model-based, distributed, and learning-assisted or hybrid approaches. The discussion focuses on how their information requirements, decision processes, and constraint-handling mechanisms affect practical implementation.
- **Direct evidence mapping of existing studies:** Representative primary studies are mapped to the same evaluation criteria using direct, partial, or insufficient evidence. The resulting comparison identifies which operational properties are supported by the reported evaluations and which remain insufficiently validated within the selected literature.
- **Operational design principles and open challenges:** The evidence gaps are translated into design principles for future wireless communication systems. Subsequently, the paper further identifies several open challenges to potentially explore in further research.

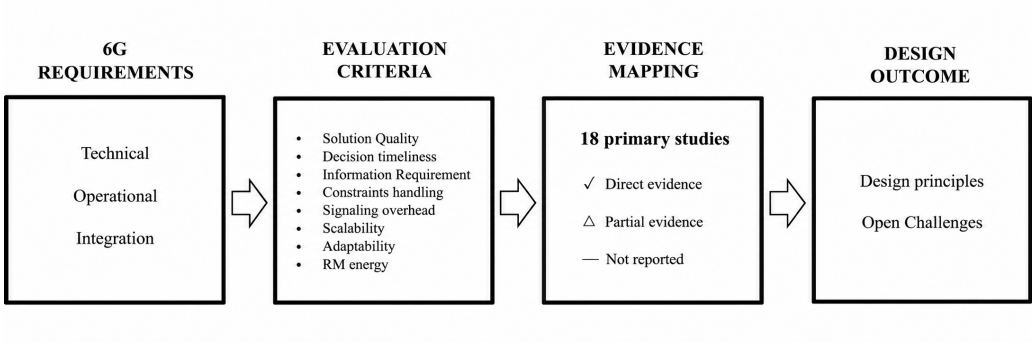


Figure 1. Requirements-to-evidence framework for evaluating resource allocation methods and deriving design guidance for 6G networks.

The remainder of this paper is organized as follows. Section 2 introduces the 6G requirements, evaluation criteria, and decision timescales considered in this study. Section 3 discusses the three resource allocation technique groups and their implementation characteristics. Section 4 maps the reviewed studies directly to the evaluation criteria and compares the completeness of their reported evidence. Section 5 presents the resulting design principles and open research challenges. Finally, Section 6 concludes the paper.

2. 6G Network Requirements and Resource Allocation Evaluation Criteria

This section reviews 6G network requirements based on ITU recommendations. Subsequently, the recommendations are translated into allocation evaluation criteria.

2.1. IMT-2030 Capability Indicators

The IMT-2030 framework formulates several expected 6G quantitative performance metrics, providing a starting point for evaluating resource allocation methods. Recommendation ITU-R M.2160 formulates a capability indicator. In addition, Report ITU-R M.2516 characterizes the supporting technology trends [4,18]. Table 2 summarizes the capability indicators in the IMT-2030 framework. Here, the capability indicators depend heavily on the usage scenario [4,5,19]. For example, a scenario requiring very high reliability, e.g., a healthcare service, does not necessarily require the maximum peak data rate. Resource

allocation should therefore determine which indicator is dominant for each scenario and allocate resources accordingly.

Table 2. IMT-2030 capability indicators and their resource allocation requirements.

Capability	Value Range	Main Resource Allocation Requirement
Peak data rate	50–200 Gbit/s	Bandwidth and power allocation with interference control.
User-experienced data rate	300–500 Mbit/s or higher	User assignment and fair resource allocation.
Spectrum efficiency	1.5–3 times IMT-2020	Efficient channel reuse and interference management.
Area traffic capacity	30–50 Mbit/s/m^2 or higher	Load balancing across central networks.
Connection density	10^6–10^8 devices/km^2	Scalable access and scheduling with low signaling.
Mobility	500–1000 km/h	Mobility-aware association and handover management.
Radio latency	0.1–1 ms	Low-latency resource allocation.
Reliability	$1-10^{-5}$ to $1-10^{-7}$	Reliability-aware allocation with limited violations.
Positioning accuracy	1–10 cm	Joint communication and sensing allocation.

The capability indicators show that the network size and decision time require careful consideration. Consider an assignment problem with N users and K available resources. If each user selects one resource, the direct search space contains K^N possible assignments. For $N = 100$ users and $K = 10$ channels or serving nodes, this produces 10^{100} combinations. An optimization method does not necessarily investigate every combination; however, this example shows that the exhaustive search cannot support dense networks. The problem becomes more complicated when assignment is combined with continuous power, bandwidth, or computational variables [2,11].

A similar issue appears in the time domain. A target radio latency of 0.1–1 ms leaves only a very short interval for measurement, signaling, optimization, and decision implementation. Moreover, the allocation must remain valid long enough to be used. This is particularly important at mobility levels of 500–1000 km/h, where the channel model can change rapidly. Therefore, an algorithm with a small optimality gap may still be unsuitable if its decision time exceeds the validity period of its input information.

The reliability range also affects how optimization results should be reported. If the failure probability is 10^{-5}, then 10^5 statistically comparable transmissions have an expected failure count of one; similarly, 10^{-7} gives an expected count of one in 10^7 transmissions. These expectations do not bound the realized number of failures. Average throughput and a zero observed failure count are therefore insufficient unless the failure event, number of evaluated cases, sampling assumptions, and uncertainty are also reported. This requirement is highly important for learning-assisted methods because a high average reward does not guarantee that every generated decision is feasible [14,15].

IMT-2030 also includes coverage, sensing, AI integration, sustainability, security, and interoperability as new capability indicators [4,20]. These indicators do not yet have a fixed target that applies to every scenario. They should therefore be evaluated via application-dependent performance metrics. Examples include coverage probability, sensing error, inference latency, energy per decision, recovery time, and the success rate of cross-domain operations.

2.2. From 6G Capabilities to Resource Allocation Evaluation

The capability indicators in Table 3 describe the expected network performance; however, achieving a resource allocation objective does not necessarily mean that the allocation can be applied under the required operating conditions. The evaluation should therefore consider not only the quality of the resulting allocation, but also the decision time, coordination requirements, constraint satisfaction, and energy required by the allocation process. Peak data rate, user-experienced data rate, spectrum efficiency, area traffic capacity, and positioning accuracy mainly relate to solution quality when they are included in the allocation objective. Reliability and radio latency require explicit consideration of system constraints, while radio latency and mobility also require timely allocation decisions. Mobility further requires the allocation to adapt to changing network conditions. Connection density affects the amount of information that must be collected and exchanged, and therefore influences signaling overhead and scalability. Area traffic capacity also affects scalability because higher traffic demand increases the size and complexity of the allocation problem. Finally, resource-management energy is considered as an operational criterion because obtaining network information, exchanging control messages, and computing allocation decisions also consume energy.

Table 3. Relationship between IMT-2030 capability indicators and resource allocation evaluation criteria.

Evaluation Criterion	Related Capability Indicator	Evaluation Requirement
Solution quality	Peak data rate, user-experienced data rate, spectrum efficiency, area traffic capacity, positioning accuracy	Compare the achieved objective with a stated baseline or reference solution when the capability is represented in the allocation problem.
Decision timeliness	Radio latency, mobility	Measure the complete time from the defined observation start to allocation activation and compare it with the relevant decision deadline and state-validity horizon.
Information requirement	Connection density, mobility	State which network information is required, where it is obtained, and how frequently it must be updated.
Constraint handling	Reliability, radio latency	Report whether resource and service constraints are satisfied, including violation probability when the requirement is probabilistic.
Signaling overhead	Connection density	Measure the information exchanged for state collection and coordination, including message volume and coordination rounds.
Scalability	Connection density, area traffic capacity	Evaluate how decision time, memory, and signaling change as network size or traffic demand increases.
Adaptability	Mobility	Evaluate performance and update time after the operating condition changes from the nominal or training condition.
Resource-management energy	Connection density, area traffic capacity	Measure energy used by state acquisition, signaling, computation, model updating, and allocation application when applicable.

It should be noted that these relationships are not one-to-one. A single capability may relate to several operational criteria, while one criterion may also support several capabilities. In addition, the eight criteria do not directly cover security or interoperability. For example, satisfying resource allocation constraints does not automatically demonstrate

security, and low signaling overhead does not guarantee interoperability across different technologies. Positioning accuracy represents only one sensing-related capability and does not cover sensing and perception more broadly. Therefore, security, interoperability, and broader sensing performance are outside the scope of this paper.

2.3. Resource Allocation Evaluation Criteria

The capability indicators describe expected network performance; however, they do not show the additional metrics required to find a final decision. An optimization method may provide high throughput while requiring outdated network information, excessive computation, or repeated signaling. Therefore, the allocation result and the process used to obtain it should be evaluated. Table 4 summarizes the measures used in this paper. There are 8 evaluation dimensions, i.e., solution quality, decision timeliness, information requirement, constraint handling, signaling overhead, scalability, adaptability, and resource-management energy.

Table 4. Criteria for evaluating resource allocation methods.

Evaluation Dimension	Measure	Minimum Reporting
Solution quality	$G = 100(U_{\text{ref}} - U)/\lvert U_{\text{ref}}\rvert$	Reference solution and worst-case gaps.
Decision timeliness	$\rho_t = T_{\text{RM}}/T_{\text{valid}}$	Complete delay and percentage of decisions with $\rho_t < 1$.
Information requirement	$\eta_{\text{info}} = B^{\text{used}}_{\text{state}}/B^{\text{all}}_{\text{state}}$	Required network state, information collecting, and update interval.
Constraint handling	$P_{\text{vio}} = N_{\text{vio}}/N_{\text{test}}$	Individual and joint constraint-violation rates.
Signaling overhead	$\eta_{\text{sig}} = B_{\text{sig}}/(B_{\text{sig}} + B_{\text{data}})$	Bytes per decision, messages per node, and coordination rounds.
Scalability	Growth of $T_{\text{RM}}(N)$, $M(N)$, and $B_{\text{sig}}(N)$	Measurements across several network sizes.
Adaptability	$D_{\text{shift}} = 100(U_{\text{nom}} - U_{\text{shift}})/\lvert U_{\text{nom}}\rvert$	Performance degradation and required update time.
Resource-management energy	$E_{\text{RM}} = E_{\text{meas}} + E_{\text{sig}} + E_{\text{comp}}$	Energy per decision and net system-energy change.

The evaluation begins with the quality of the resulting resource allocation. The reference gap G compares the utility achieved by the evaluated method, denoted by U, with the utility of an exact or reference solution, denoted by U_{ref}. For a maximization problem, a smaller value of G indicates that the resulting allocation is closer to the reference solution. An exact solution can be used for small network problems, while the best-known solution can be used when the problem cannot be solved exactly. In either case, the reference must be stated clearly because an improvement over a weak baseline does not necessarily indicate an effective resource allocation. Mean results should also be accompanied by worst-case or percentile values when the allocation is intended for services with strict performance requirements.

However, a high-quality resource allocation is useful only if it can be applied before network conditions change. For one resource-management cycle, let t_0 denote the start of the defined measurement procedure and t_{act} indicate the time at which the resulting allocation becomes active at the affected nodes. The complete elapsed resource-management delay is

$$T_{\text{RM}} = t_{\text{act}} - t_0. \tag{1}$$

For a sequential implementation with nonoverlapping measurement, signaling, solution, and application stages, Equation (1) reduces to $T_{\text{meas}} + T_{\text{sig}} + T_{\text{solve}} + T_{\text{apply}}$. When

an implementation is represented by an acyclic dependency graph with stage and transfer delays, its scheduled completion is determined by the longest executed path from the measurement start to allocation activation.

The validity period therefore also accounts for the different rates at which the required state components evolve. Suppose $\mathcal{K}$ contains the state components required by the decision, t_k^{obs} indicates the observation time of component k, and τ_k^{valid} is its justified validity duration. At cycle origin t_0, a conservative remaining validity period is

$$T_{\text{valid}} = \min_{k \in \mathcal{K}} \left[\tau_k^{\text{valid}} - (t_0 - t_k^{\text{obs}}) \right]. \tag{2}$$

The validity duration of each component should be determined from measurement or a justified state-evolution model using a decision-relevant tolerance. For example, channel-state validity may be determined by the time over which the channel-prediction error remains within an acceptable bound, queue-state validity by the time over which the queue-estimation error remains acceptable, and topology validity by the predicted remaining availability of the corresponding link. The timeliness ratio is then expressed as

$$\rho_t = \frac{T_{\text{RM}}}{T_{\text{valid}}}, \tag{3}$$

when $T_{\text{valid}} > 0$. A decision is timely under this definition when $\rho_t < 1$ and the application-specific deadline is also satisfied.

For an offloaded task whose allocation, upload, queueing, execution, and result return are sequential, the end-to-end task delay can be expressed as

$$T_{\text{task}} = T_{\text{RM}} + T_{\text{UL}} + T_{\text{queue}} + T_{\text{exec}} + T_{\text{DL}}, \tag{4}$$

where T_{RM} includes only the resource-management time that the task actually waits for. When tasks or sub-tasks execute concurrently, the dependency graph and release times determine the completion delay; overlapping durations should not be added twice. Communication and computation resources must therefore be considered jointly when they determine the same task deadline [7,21].

A similar condition applies to mobility management. If handover execution is not already included in allocation application, the decision should satisfy

$$T_{\text{RM}} + T_{\text{HO}} < T_{\text{remain}}, \tag{5}$$

where T_{remain} is the predicted remaining service time. If T_{apply} already includes the handover, T_{HO} must not be counted again. This distinction is important in vehicular and LEO networks, where serving-node availability can change before the decision is installed [8,9,12].

Achieving a timely decision also depends on the amount of network information required by the method. The information ratio η_{info} compares the state information used to calculate the allocation, denoted by $B_{\text{state}}^{\text{used}}$, with the information required to represent the complete network state, denoted by $B_{\text{state}}^{\text{all}}$. A value close to one indicates that the method requires a nearly complete network view, while a smaller value indicates that the decision can be calculated using more limited information. This ratio should be reported together with the types of information collected, their update intervals, and whether they are obtained locally or through a central controller. Otherwise, a method may appear to provide fast decision-making while excluding the time and communication required to collect its input information.

In addition, the decision must satisfy the system constraints. Let $g_j(\mathbf{x}, \mathbf{s})$ represent the jth constraint for resource decision $\mathbf{x}$ under network state $\mathbf{s}$. A probabilistic requirement can be written as

$$\Pr\left[g_j(\mathbf{x}, \mathbf{s}) > 0\right] \leq \epsilon_j, \tag{6}$$

where ϵ_j is the maximum permitted violation probability. If $N_{\text{vio},j}$ violations are observed in $N_{\text{test},j}$ evaluated cases, the empirical estimate is

$$\widehat{P}_{\text{vio},j} = \frac{N_{\text{vio},j}}{N_{\text{test},j}}. \tag{7}$$

The violation count and sample size must be reported together with a confidence level and interval procedure. Under independent Bernoulli trials, a one-sided exact upper confidence bound can be expressed as

$$p_j^{\text{U}} = \text{Beta}^{-1}(1 - \alpha; N_{\text{vio},j} + 1, N_{\text{test},j} - N_{\text{vio},j}), \tag{8}$$

where $1 - \alpha$ is the confidence level. If no violations are observed, the upper bound becomes $p_j^{\text{U}} = 1 - \alpha^{1/N_{\text{test},j}}$. At 95% confidence, at least 299,572 independent trials with zero violations are needed to obtain an upper bound of 10^{-5}. In addition, 29,957,322 trials are needed for 10^{-7}. If wireless events are correlated or the operating conditions change over time, the sampling model should account for this dependence. The evaluation should also distinguish between violations of individual constraints and the joint event in which at least one required constraint is violated.

When resource decisions are coordinated among different network nodes, the required information exchange should also be included in the evaluation. The signaling ratio η_{sig} compares the signaling bits B_{sig} with the combined signaling and data traffic $B_{\text{sig}} + B_{\text{data}}$. Here, B_{sig} includes the information exchanged to obtain and coordinate the resource decision, while B_{data} represents the transmitted user data during the same evaluation period. For iterative distributed methods, total signaling depends on message size, the number of neighboring nodes, and the number of coordination rounds. Therefore, reducing the size of one message does not necessarily reduce the total overhead when the same message is exchanged repeatedly [9,12,13].

The computational and signaling measurements should then be examined as the network size increases. Scalability is represented by the growth of decision time $T_{\text{RM}}(N)$, memory usage $M(N)$, and signaling $B_{\text{sig}}(N)$ with the number of network nodes N. These values should be measured at several network sizes because a single large simulation does not show the method's actual growth. A method cannot be considered scalable only because its execution time remains low if its memory or signaling grows excessively [11,16,17].

In addition to network size, operating conditions may differ from those used during method design or training. The adaptation loss D_{shift} compares the utility under the original condition, denoted by U_{nom}, with the utility after the network condition changes, denoted by U_{shift}. A smaller value indicates that the method maintains performance after network changes. This evaluation should state which condition changed and report the time required to recalculate, update, or retrain the method. Without this information, a method may appear adaptable only because it is evaluated under conditions similar to those used during its development.

Finally, the resource-management procedure consumes energy during network measurement, information exchange, and solution calculation. This energy is expressed as

$$E_{\text{RM}} = E_{\text{meas}} + E_{\text{sig}} + E_{\text{comp}}, \tag{9}$$

where E_{meas} represents the energy used to obtain the network information, E_{sig} indicates the energy required for signaling, and E_{comp} represents the processing energy. These components should be included in the energy-efficiency evaluation because lower transmission

power may be offset by greater signaling or computational energy. Therefore, the benefit of a resource allocation method should be determined by the net change in total system energy [7,20,22].

2.4. Resource Allocation Problems and Decision Timescales

The capability indicators and evaluation criteria are realized through decisions over several wireless resources, e.g., radio, connectivity, mobility, computing, storage, and energy consumption. Although these decisions are often studied separately, they share a similar structure. Let $\mathbf{s}$ denote the available network state, e.g., channel conditions, traffic demand, node locations, queue states, and resource availability. A resource allocation method determines a decision vector $\mathbf{x}$ by solving

$$\begin{aligned} \max_{\mathbf{x}} \quad & U(\mathbf{x};\mathbf{s}) \\ \text{s.t.} \quad & g_j(\mathbf{x};\mathbf{s}) \le 0, \qquad j = 1,\dots,J, \\ & \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{10}$$

where $U(\mathbf{x};\mathbf{s})$ represents a system objective and $g_j(\mathbf{x};\mathbf{s})$ represents an operational constraint. The feasible set $\mathcal{X}$ may contain continuous or discrete wireless variables. The coexistence of these variables leads to nonconvex resource allocation problems [5,11]. Table 5 summarizes the main decision groups considered in this paper. It should be noted that the decision timescales are relative. Their actual values depend on channel variation, mobility, and traffic conditions.

Table 5. Representative resource allocation problems and their decision timescales.

Resource Allocation Problem	Main Decision	Decision Timescale
Channel and power allocation	Assign channels, bandwidth, beams, and transmission power	Very fast to fast
User association and scheduling	Select serving nodes, scheduled users, access opportunities, and traffic priorities	Fast to medium
Routing and mobility management	Select communication paths, relay nodes, handover targets, and serving layers	Medium
Computation offloading and processing	Select offloading destinations, processing order, and computing-resource allocation	Fast to medium
Storage and content placement	Determine caching, model placement, and storage-resource allocation	Medium to slow

3. Resource Allocation Methods for 6G Networks

Existing resource allocation methods are organized here into three groups: centralized model-based optimization, distributed optimization, and learning-assisted or hybrid optimization. This grouping is based on their implementation costs. This section presents these methods in detail, as illustrated in Figure 2 and as summarized in Table 6.

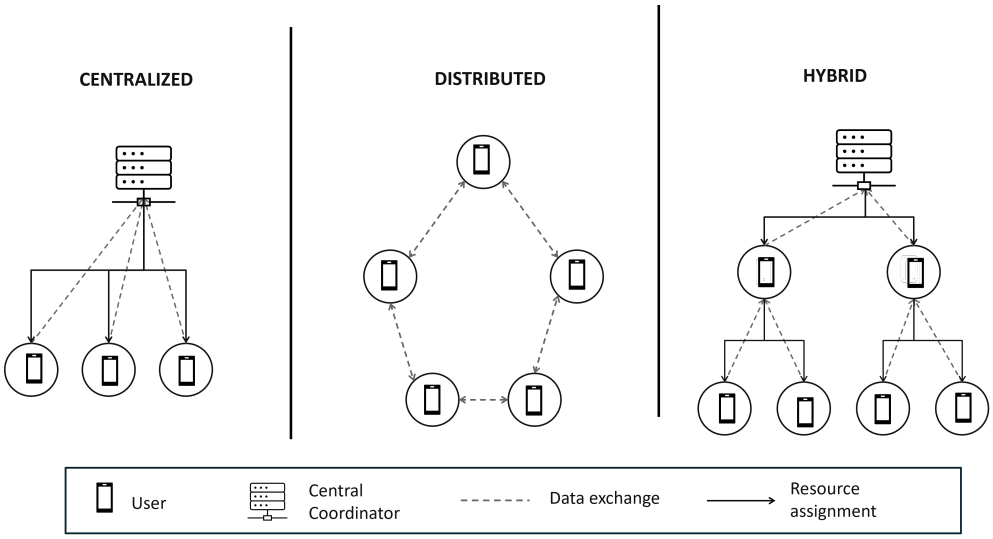


Figure 2. Information exchange and resource assignment in centralized, distributed, and hybrid architectures.

Table 6. Implementation characteristics of the three resource allocation method groups.

Method Group	Information	Constraints	Main 6G Risk
Centralized model-based	Global state	Explicit	Invalid decisions
Distributed optimization	Local messages	Coordinated	Signaling delay
Learning-assisted and hybrid	Features or models	Structured	Changes in multiple network conditions

3.1. Centralized Model-Based Optimization

Centralized model-based methods use a global view of the network to jointly allocate resources to involved users. Their main advantage is the ability to optimize system-wide objectives while explicitly enforcing operational constraints. However, this advantage diminishes when collecting global information takes longer than the network topology remains valid.

Recent studies illustrate both properties. Wu et al. jointly considered channel assignment, beamforming, and power allocation in Simultaneously Transmitting and Reflecting Reconfigurable Intelligent Surface (STAR-RIS) networks using matching and alternating optimization [23]. Zhai et al. coordinated Unmanned Aerial Vehicle (UAV) trajectory, RIS configuration, offloading, and computing allocation through successive convex approximation and fractional programming [24]. Both studies show that explicit models can coordinate multiple resource types.

Similarly, in edge computing, Nguyen et al. combined metaheuristic task assignment with convex communication-resource allocation for dependent tasks in collaborative UAV networks [25]. Liu et al. used queue-aware dynamic optimization for computation offloading under static and dynamic subchannels in energy-harvesting Multi-access Edge Computing (MEC) networks [26]. For multiuser edge inference, Liu et al. jointly considered batching, early exiting, scheduling, and bandwidth allocation to improve throughput [27]. These methods maintain task, capacity, and deadline constraints, but their suitability for low-latency 6G networks also depends on decision time, global information collection, and dynamic network changes.

Centralized control is suitable when the network can collect the required information, solve the joint resource allocation problem, and apply the resulting decision within the required control time. Its practicality therefore depends strongly on the problem size, includ-

ing the number of users, discrete allocation variables, antennas, RIS elements, constraints, and jointly optimized resource domains. As the problem size increases, the computational burden can also increase significantly. For example, [23] evaluates exhaustive-search benchmarks only for four- and six-user cases, providing a useful reference for small-scale optimality comparison. Centralized control is therefore more practical for planning and supervisory allocation, where longer decision times are acceptable. It can also support faster control, provided that the complete process, from information collection to allocation, can still be completed within the application deadline.

3.2. Distributed Optimization

Distributed methods split the allocation burden among network nodes using small amounts of information. They eliminate dependence on a central coordinator and exploit parallel computation. However, locality does not automatically provide low overhead. Total cost highly depends on neighborhood density, message size, and the number of coordination rounds.

Bai and Wang proposed distributed optimization followed by centralized improvement, allowing devices to produce local allocations before a controller resolves remaining conflicts [28]. Wu et al. decomposed communication and computing decisions for IIoT (Industrial Internet of Things) networks, showing the need to coordinate several processing locations without solving the complete mixed-integer problem at one location [29]. Ji and Qin used federated reinforcement learning for distributed channel and power allocation, reducing the transfer of local channel information while improving robustness to channel variation [30]. These studies demonstrate different forms of partial coordination.

Message-based methods provide further evidence. The hierarchical affinity-propagation method in [12] distributed routing decisions across a multilayer vehicular network and reported execution and route-request signaling as the number of vehicles increased. The priority-aware message-passing method in [13] enabled healthcare IoT nodes to coordinate transmission schedules without a central controller while accounting for both medical urgency and network throughput. Although the method remained effective in dense deployments, its suitability for 6G also depends on whether coordination can be completed within the service requirement under imperfect message exchange.

A simple example is bandwidth allocation across cells that use orthogonal bandwidth resources and have independent local constraints. If the objective set can be separated, $U(\mathbf{x}) = \sum_i U_i(x_i)$ and $\mathcal{X} = \prod_i \mathcal{X}_i$, each cell can solve its own local problem independently without causing resource conflicts. However, this does not directly apply when the cells share common spectrum resources. If the cells share a total bandwidth budget $\sum_i x_i \leq B$, their local decisions must be coordinated, for example, using a common nonnegative price λ. Under the usual convexity conditions for dual decomposition, each cell solves its local problem based on this price, while λ updates based on the difference between the total requested bandwidth and the available bandwidth. Before convergence, the combined local allocations may still exceed the shared budget; if an intermediate solution must be used, a correction rule is needed. Distributed optimization is therefore suitable when the interaction among local problems can be coordinated properly.

3.3. Learning-Assisted and Hybrid Optimization

Learning-assisted methods replace repeated optimization with a trained allocation policy. Hybrid methods also maintain model-based components to guide learned decisions.

Guo and Yang used a heterogeneous graph neural network (GNN) to learn power allocation while maintaining the permutation structure of multi-cell multi-user systems [31]. Wang et al. introduced a general-edge update of the GNN to update both node and edge

representations for beamforming and evaluated generalization across several network conditions [32]. Peng et al. showed that the suitability of vertex- or edge-based GNNs depends on the location and dimension of the decision variables, linking model architecture to the requirements of scheduling, power control, and precoding [33].

A hybrid method can reduce the gap between model structure and learned inference. Yang et al. combined weighted minimum mean-square error (WMMSE) operations into a GNN for wireless-resource allocation, maintaining algorithmic knowledge while reducing repeated online computation [34]. For RAN slicing, Akyildiz et al. used hierarchical reinforcement learning to coordinate inter-slice and intra-slice allocation for eMBB and URLLC traffic [35]. These results support learning-assisted optimization for repeated and structured decisions.

Huang et al. jointly considered collaborative sensing-task offloading and communication-bandwidth allocation in an integrated sensing, communication, and computation vehicular system [36]. Uslu et al. developed state-augmented graph learning with dual-variable regression for constrained wireless power allocation [37]. Yu et al. combined graph features with dueling double deep Q-learning for V2X spectrum and power allocation [38]. These studies broaden the reviewed applications and learning mechanisms, but their recent publication does not by itself establish complete operational validation.

4. Comparative 6G Evaluation Criteria

The reported improvement in a resource allocation objective does not establish that a method can operate under 6G networks. The allocation must also be produced using information that can be collected in practice, completed before network conditions change, and applied without violating resource constraints. This section therefore evaluates the reviewed studies according to the criteria defined in Table 4. The objective is to determine which aspects of their 6G suitability have been demonstrated and which remain unsupported.

4.1. Evaluation Procedure and Evidence Overview

The quantitative mapping includes 18 representative primary studies. A study is included when it defines a wireless-resource allocation decision, describes the optimization or learning mechanism, and reports analytical or numerical evidence sufficient to assess at least the allocation objective and relevant operating assumptions. Surveys are used as background and are not counted as primary evidence. Duplicate preprint and published versions are treated as one study, and papers without an identifiable allocation mechanism or sufficient technical information are excluded from the quantitative mapping. The selection is intended to support a requirements-based comparison.

Each study is assigned one of three evidence levels for every criterion. Direct support is assigned when the reported measurement corresponds directly to the criterion. Partial support indicates that the study reports a related but incomplete measure. An unmarked entry indicates that the study does not provide sufficient information for an assessment.

The distinction is particularly important for decision timeliness. For example, iteration count indicates computational load, but it excludes the time required to collect network information, transfer it to the decision, establish an output, and implement the resulting allocation. These measurements are therefore considered partial evidence.

A similar distinction is applied to energy. Transmission, computation, or UAV propulsion energy describes the service being optimized, whereas resource-management energy refers to the energy consumed to obtain and maintain the allocation. The latter includes measurement, signaling, optimization, model inference, and model updating. A study that minimizes service energy without measuring these control activities is therefore classified as providing only partial evidence of energy. Table 7 applies these criteria to 18 primary studies.

Table 7. Evidence reported by 18 studies against the evaluation criteria.

Study	Quality	Decision Time	Info.	Constr.	Signaling	Scale	Adapt.	RM Energy
Centralized model-based optimization								
Wu et al. [23]	✓	△	△	✓	–	△	–	–
Zhai et al. [24]	✓	△	△	✓	–	–	–	△
Nguyen et al. [25]	✓	△	△	✓	–	△	–	–
Liu et al. [26]	✓	△	△	✓	–	△	△	△
Liu et al. [27]	✓	△	△	✓	–	△	–	–
Distributed optimization								
Bai and Wang [28]	✓	△	✓	✓	△	✓	–	–
Wu et al. [29]	✓	△	✓	✓	△	△	△	△
Ji and Qin [30]	✓	△	✓	△	△	△	✓	△
Dewa et al. [12]	✓	△	✓	✓	✓	✓	△	–
Dewa et al. [13]	✓	△	✓	✓	△	✓	–	–
Learning-assisted and hybrid optimization								
Guo and Yang [31]	✓	△	△	△	–	✓	△	–
Yang et al. [34]	✓	△	△	△	–	✓	△	–
Wang et al. [32]	✓	△	△	△	–	✓	✓	–
Peng et al. [33]	✓	△	△	△	–	✓	△	–
Akyildiz et al. [35]	✓	△	△	△	–	△	△	–
Huang et al. [36]	✓	△	△	✓	–	△	–	–
Uslu et al. [37]	✓	△	△	✓	–	✓	△	–
Yu et al. [38]	✓	△	△	△	–	△	△	–

Note: ✓ denotes direct evidence, △ denotes partial or indirect evidence, and – denotes insufficient evidence.

The mapping shows an initial distinction between method groups. Centralized model-based studies provide the clearest evidence of constraint handling because they explicitly represent it accordingly. Distributed methods describe their local information and exchanged variables more clearly, although the amount of signaling is rarely measured completely. Learning-assisted studies provide evidence of short online inference and scaling across network sizes; however, their feasibility and adaptability often remain dependent on the conditions represented during training.

4.2. Centralized Model-Based Optimization

Centralized model-based methods calculate an allocation from global information by aggregating involved users. The main strength of centralized methods is that objectives and constraints are visible. This structure enables determining whether the calculated solution satisfies certain performance requirements. Their principal limitation in 6G networks is that a mathematically valid solution may no longer be valid when the information required becomes very expensive to collect.

Wu et al. [23] investigated sum-rate-maximizing resource allocation in a STAR-RIS-aided downlink network. For the orthogonal multiple access (OMA) case, the proposed method determines channel assignment using matching theory and then optimizes time allocation, transmission power, and reflection coefficients. For non-orthogonal multiple access (NOMA), a location-based matching algorithm pairs users from the transmission and reflection regions on each subchannel. The decoding order, STAR-RIS coefficients, and power allocation are subsequently optimized using semidefinite programming, convex upper-bound approximation, and geometric programming.

Comparisons with exhaustive search for four- and six-user settings show that the proposed channel-assignment schemes achieve near-optimal sum rates, providing direct evidence of solution quality. The formulations also enforce the minimum user rate, total transmission-power budget, per-element energy conservation, phase-shift range, channel

assignment, user multiplexing, time allocation, and successive-interference-cancellation constraints. These results support direct evidence of constraint handling.

The evidence for the operational criteria is less complete. The paper proves the convergence of the proposed algorithms and provides their computational complexity. The OMA algorithm converges mostly within three iterations, while the iterations required by the NOMA coefficient-optimization stage increase with the number of STAR-RIS elements. These results describe the computational behavior; however, they do not capture the full delay from channel configuration to applying the resulting configuration. The method assumes perfect channel-state information at the access point and that the STAR-RIS controller exchanges information with the access point via a separate wireless link. However, the method does not evaluate the amount, update interval, and delay of this information exchange. Complexity analysis and experiments with different numbers of users and STAR-RIS elements provide partial evidence of scalability. Variations in transmission power and user location show sensitivity within the assumed model, but they do not show how the method updates an allocation after the operating condition changes. Accordingly, adaptability is marked as insufficient evidence in Table 7. Finally, the transmission-power budget and per-element energy-conservation constraint do not represent the energy consumed to collect network information, execute the optimization, and configure the STAR-RIS. Thus, they do not provide direct evidence of resource-management energy.

Zhai et al. [24] considered a UAV-mounted RIS that supports task offloading to an MEC server. The method jointly optimizes the UAV trajectory, passive beamforming, and computing-resource allocation to improve energy efficiency. The formulation explicitly represents flight constraints, phase-shift limits, and computing-resource limits. The reported energy-efficiency improvement therefore demonstrates solution quality. The formulation also provides direct evidence that the stated constraints are handled.

The reported energy result should not be interpreted as complete resource-management energy. It mainly represents the energy associated with communication, computing, and UAV operation. The additional energy required to estimate the channels, calculate the trajectory, exchange control information, and configure the RIS is not measured separately. Decision time is also supported only indirectly because convergence of the proposed iterative procedure does not show whether the trajectory and beamforming decisions can be updated before the channels and user locations change. The study therefore shows that joint allocation is effective under the evaluated model, but not that it can be repeatedly recalculated within a practical UAV control cycle.

Nguyen et al. [25] investigated dependent-task offloading in collaborative MEC-enabled UAV networks. A central controller determines where dependent sub-tasks execute, while communication resources are allocated among mobile users and their associated UAVs. The method combines a discrete whale-optimization procedure for offloading with a convex communication-allocation subproblem. Task dependencies determine when subtasks can start and whether independent branches can execute concurrently. The reported latency and data-rate comparisons support solution quality within the evaluated model, while the modeled resource and service restrictions support constraint handling.

The controller requires task and network information, but the study does not separately measure the complete state-acquisition delay, allocation-installation delay, control-message volume, or optimizer energy. Algorithmic search behavior and tests at different problem sizes provide partial operational evidence, but service latency cannot substitute for a measurement of the complete allocation cycle. The study should therefore not be described as a joint satellite-path, remote-cloud, UAV-placement, and computing-capacity optimization.

The dynamic MEC method of Liu et al. [26] addresses time-varying computation requests, energy-harvesting devices, and static–dynamic assigned subchannels. The re-

source allocation combines queueing and long-term resource constraints, thereby ensuring that current decisions also consider future queue stability. This treatment reflects repeated network operation. The reported delay, energy, and computation performance support solution quality. In addition, the modeled queue, battery, subchannel, and computing constraints provide direct evidence of these constraints.

The use of stochastic arrivals and changing subchannels provides some evidence of adaptability; however, the adaptation remains within the process assumed by the model. It does not show robustness when the arrival distribution, channel behavior, and energy-harvesting process differ from the evaluated assumptions. Similarly, the energy result considers only task execution and device operation; accordingly, it does not include the full resource-management energy required. Decision time remains partial because the reported algorithmic runtime is not compared with the duration for which the queue, battery, and channel state remain representative.

Liu et al. [27] investigated joint communication and computation allocation for multiuser edge inference. Batching allows several inference tasks to be processed together. In addition, early exiting reduces computation for tasks that can meet their accuracy requirements without using the full neural network. The allocation selects users, bandwidth, batch size, and exit points to improve completed-task throughput under latency and accuracy requirements. The resulting throughput and delay performance provide direct evidence of solution quality, and the scheduling and computing constraints are included explicitly.

The reported latency describes the time needed to transmit and execute inference tasks after an allocation is available. It does not include the time required to collect task demands, channel conditions, and processor availability. The algorithm is evaluated under different workloads, providing partial evidence of scalability. However, this study does not evaluate adaptation to changes in model architectures. Signaling and the energy consumed by the allocation process are also not reported.

Across the five centralized studies, allocation quality is supported by comparisons of each scenario-specific objective with stated baselines or reference solutions. In [23], small-user channel-assignment experiments compare sum-rate performance with exhaustive search. The offloading study in [25] compares latency and data-rate results with alternative schemes, while [24,26,27] report improvements in their respective energy, delay, or inference-throughput objectives. These findings demonstrate benefit within the tested models; they do not provide a common optimality gap or establish global optimality across studies. Explicit constraints are assessed separately under constraint handling and do not by themselves establish objective quality. None of the studies measures the complete validity of the input state at activation, the full communication required to construct the global view, and the complete resource-management energy.

4.3. Distributed Optimization

Distributed methods divide the decision process among devices, access points, controllers, or network domains. This reduces the need to transmit every local observation to one optimizer. However, it does not remove communication from the decision process. Local decisions must exchange enough information to remain mutually consistent, and this exchange may continue for several rounds. For this reason, the evaluation distinguishes between describing the exchanged variables and quantifying the complete signaling required for one allocation.

Bai and Wang [28] proposed a distributed optimization with a centralized-refining mechanism. Devices first optimize local resource reservations and report the resulting information. A central optimizer then constructs a resource-trading graph and refines the feasible allocation within a reduced search region. The method is classified as distributed

because the devices perform the initial optimization workload, although the final refinement remains centralized. Its near-optimal utility and reduced central processing time provide direct evidence of quality and scalability. The local reservations and reports also make the information flow explicit.

The reported reduction in central processing time is not a complete decision-time measurement because it does not include local computation, reporting, graph construction, refinement, and allocation delivery. Signaling is classified as partial because the reservation and reporting procedure is specified, but the total exchanged bytes and the effect of unreliable links are not fully measured. Constraint handling is direct because the refinement is performed within the feasible region. The study does not establish adaptability when user behavior, resource demand, or the trading graph differs from the evaluated cases.

Wu et al. [29] considered multidomain resource allocation for an Industrial IoT-based wireless control system by jointly considering the control, communication, and computational complexity. The method minimizes the energy required to complete control tasks via resource-sharing decisions, computation allocation, and transmission-power control. The control-convergence requirement is translated into communication-reliability constraints. The proposed distributed algorithm then separates task-processing decisions from computation and power allocation while coordinating the two levels through their energy results.

The benchmark comparisons support solution quality, the stated information requirements, and constraint handling under the evaluated model. Scalability is marked as partial evidence because the reported results do not jointly show how the complete decision delay, memory, and signaling increase with network size. Decision timeliness remains partial because the study measures algorithmic processing time without the complete coordination and application delays. Signaling is also partial because the exchanged variables are defined; however, their total volume, delivery delay, and reliability are not evaluated. Variations in network size and system parameters provide partial evidence of adaptability. The reported energy considers task communication and computation; however, it excludes the energy required for measurement, coordination, and optimization.

Ji and Qin [30] used federated learning to support distributed energy-efficient resource allocation. Local agents learn from their own wireless conditions, while model aggregation shares knowledge without transferring the underlying observations directly. This structure provides direct evidence of the information architecture because the local inputs and model exchange are part of the method. Its resource allocation performance establishes solution quality, and evaluation under heterogeneous or changed local conditions supports the direct adaptability rating used in Table 7.

Nevertheless, federated aggregation introduces a second communication process in addition to resource allocation. Model parameters must be uploaded, aggregated, and redistributed. The study reports this process but does not provide complete signaling and energy considerations for every training round. The energy-efficiency objective mainly concerns wireless operation; thus, resource-management energy is only partially considered. Constraint handling is also partial because learned decisions and reward-based control do not establish the frequency or magnitude of constraint violations. The reported training or inference time does not represent the complete delay required to observe the network, update the model when necessary, and apply the allocation.

Dewa et al. [12] developed a hierarchical affinity-propagation method for distributed routing in an integrated multilayer V2X network. The method exchanges local information to select routes under changing vehicular conditions. Evaluations across different network sizes show improved routing performance and lower signaling than the baseline, providing direct evidence of solution quality, signaling, and scalability. However, decision timeliness remains partial because the reported execution time does not include the complete process

from topology observation to route installation. Variations in vehicle density and mobility provide partial evidence of adaptability, while route validity during execution and resource-management energy are not evaluated.

Dewa et al. [13] proposed a distributed priority-aware scheduling method for healthcare IoT networks. Local messages allocate shared wireless resources while considering device priority, throughput, and concurrent-transmission constraints. The reported results support solution quality and constraint handling, while the complexity analysis provides evidence of computational scalability. However, signaling is classified as partial because the exchanged packet counts are reported only for selected configurations without considering message size, retransmissions, or communication delay. Decision timeliness remains partial because the end-to-end allocation delay and deadline violations are not evaluated. Adaptability is marked as insufficient because density and scheduling comparisons do not show adaptation to a changed operating condition or to asynchronous message delivery. Resource-management energy is not reported.

The distributed studies provide stronger evidence than the centralized group because the variables they exchange form part of the algorithm design. However, knowing the content of one message is not equivalent to knowing the communication cost of the complete decision. The total number of messages depends on neighborhood size, iteration count, retransmission, and synchronization. Only the V2X study reports a direct signaling rate, while the other distributed studies provide indirect or configuration-specific evidence. Distributed optimization is therefore most convincing when local interaction remains sparse and the coordination process produces a usable allocation before the network changes.

4.4. Learning-Assisted and Hybrid Optimization

Learning-assisted methods replace repeated numerical optimization with a trained mapping from network observations to resource decisions. Hybrid methods retain part of an established optimization structure while learning the components that are expensive to calculate repeatedly. Their main advantage is predictable online inference. Their main limitation is that successful inference under evaluated conditions does not guarantee feasibility or performance under an unfamiliar network condition.

Guo and Yang [31] applied a heterogeneous graph neural network to power allocation in multicell multiuser systems. The graph distinguishes different node and link types, allowing one set of learned parameters to represent interference relationships across cells and users. The reported sum-rate performance provides direct evidence of solution quality. Evaluation across different network sizes provides direct scalability evidence because the shared graph parameters can be applied without defining a separate fully connected model for every system size.

The method still requires channel and interference features to construct the graph. These inputs are described, but their acquisition and graph-construction time are not included in the reported inference time. Information evidence is therefore partial. Constraint support is also partial because a valid output parameterization or training penalty does not provide the probability and magnitude of violations under unseen network conditions. Evaluation at another node count provides partial evidence of adaptability, but it does not establish robustness to different propagation, interference, mobility, traffic, or service constraints. Signaling and lifecycle energy are not reported.

Yang et al. [34] embedded the structure of WMMSE optimization into a graph neural network. Instead of learning the complete allocation as an unconstrained black-box mapping, the network follows repeated WMMSE-style updates using a bounded number of trainable layers. This structure improves interpretability and reduces the online iterations required to approximate the model-based solution. The reported objective performance provides

direct evidence of solution quality. In addition, evaluation across different problem sizes provides direct evidence of scalability.

The reduced inference time is classified as partial decision evidence. A practical implementation must also construct the interference graph, obtain its channel features, execute the network, correct infeasible outputs if necessary, and deliver the power allocation. The WMMSE structure provides more support for constraint handling than an unconstrained neural network, but it does not automatically guarantee feasibility under every input. Adaptability is partial because size generalization does not establish robustness when the relationship between the channel state and optimal allocation changes. Training, updating, signaling, and resource-management energy remain unreported.

Wang et al. [32] proposed an edge-update-empowered graph neural network for radio-resource management. In addition to updating node representations, the architecture updates edge features that represent wireless links and interference relationships. This makes the model applicable to decisions whose outputs are naturally associated with links. The study evaluates the architecture on several resource allocation settings and reports competitive allocation performance.

The evaluation provides direct scalability evidence because the architecture is tested across different network sizes. It also provides direct adaptability evidence in Table 7 because it evaluates performance beyond the original training configuration. However, the evidence remains tied to the tested range of wireless conditions. The graph features and their acquisition are not included in the reported inference cost; thus, information and decision-time support remain partial. Constraint handling is also partial unless the output is followed by a feasibility test or the violation probability is reported. The study does not quantify signaling for distributed graph execution or the energy consumed by training and inference.

Peng et al. [33] compared vertex-based and edge-based GNN representations for learning wireless-resource policies. The study examines how the location of decision variables in a graph affects the model's expressive capability. Link scheduling, power control, and beamforming require different graph outputs; consequently, placing every output at a node is not equally suitable for all three problems. The reported experiments provide direct evidence on solution quality and scalability.

The study provides partial evidence of the information required because it does not measure the cost of exchanging required node and edge features. Theoretical expressiveness and empirical feasibility do not replace direct constraint-violation measurements; accordingly, constraint handling remains only partially supported. Generalization across graph sizes provides partial adaptability evidence, but it does not show operation under different channel distributions, mobility, traffic, or service requirements. Decision time also excludes graph preparation and any output correction. Signaling and resource-management energy are not quantified.

Akyildiz et al. [35] developed a hierarchical reinforcement-learning architecture for RAN slicing with URLLC and eMBB traffic. An upper-level agent allocates resources between slices, while lower-level agents allocate the assigned resources to users within each slice. Queue occupancy and channel quality form the observed state, and the objective balances eMBB throughput against the latency requirement of URLLC traffic. The reported throughput and delay results provide direct evidence of solution quality.

The hierarchy reduces the action space and separates inter-slice from intra-slice decisions, but the evaluation provides only partial scalability evidence because the tested hierarchy covers a limited range of users and slice configurations. Constraint handling remains partial because achieving a low average URLLC delay does not establish the tail probability of exceeding the latency limit. Adaptability is also partial: random traffic arrivals

are included, but operation under a new traffic distribution, changed service constraints, or an unfamiliar radio environment is not fully evaluated. Policy inference time excludes state collection, agent interaction, and decision installation. Training energy, model updating, and coordination signaling are also not reported.

The three recent studies add complementary evidence. Huang et al. [36] report processing-delay improvements for collaborative sensing-aware offloading and bandwidth allocation; this supports solution quality and the modeled constraints, while the complete measurement-to-activation delay and control signaling remain unmeasured. Uslu et al. [37] provide constrained power-control results, convergence analysis, and tests that transfer the graph-based policy across network sizes; the long-term constraints provide direct constraint evidence but should not be interpreted as per-packet reliability guarantees. Yu et al. [38] report V2I capacity and V2V transmission-success improvements while varying the number of vehicles and payload; these results support the evaluated performance and provide partial scaling and adaptation evidence, but they do not measure the complete control cycle, signaling volume, or resource-management energy.

The learning-assisted group demonstrates that online calculation can be shortened and models can often be reused across different network sizes. Its remaining weakness is the absence of an operational boundary. Most studies do not specify when an input should be considered outside the conditions represented during training, how to detect an unsafe output, or whether correction remains faster than model-based optimization. For 6G operation, inference speed must therefore be evaluated together with graph preparation, uncertainty detection, feasibility correction, and model updating.

4.5. Cross-Method Evidence Synthesis

The preceding evaluation shows that no method group dominates all eight criteria. Centralized model-based methods provide the strongest structural support for feasibility, but their operation depends on timely global information. Distributed methods make information exchange more explicit and reduce central processing, but repeated coordination can become the dominant delay and signaling cost. Learning-assisted methods provide bounded online inference and stronger scaling across network sizes, but their feasibility and reliability under unseen network conditions remain less established.

To summarize reporting coverage, let N denote the number of mapped primary studies and $w \in [0, 1]$ denote the declared weight assigned to partial evidence. The reporting-coverage score for criterion q is

$$C_q(w) = \frac{N_q^{\checkmark} + wN_q^{\triangle}}{N} \times 100\%. \tag{11}$$

where $N_q^{\checkmark}$ and $N_q^{\triangle}$ denote the numbers of studies providing direct and partial evidence, respectively.

Direct evidence has weight one, insufficient evidence has weight zero, and this paper uses $w = 0.5$ as an explicit midpoint convention between these two evidence levels, which is reported in Table 8. The score summarizes reporting completeness within the selected studies and does not rank scientific importance.

Figure 3 evaluates the sensitivity of the reporting-coverage scores for all eight criteria over the full range $0 \leq w \leq 1$. Solution quality is unaffected by w because all 18 studies provide direct evidence. The principal ranking change identified in the sensitivity analysis occurs between scalability and information requirement. Scalability has eight direct and nine partial entries, while information requirement has five direct and thirteen partial entries, giving an unnormalized difference of $(8 + 9w) - (5 + 13w) = 3 - 4w$. Thus, scalability has higher reporting coverage for $w < 0.75$, the two criteria tie at $w = 0.75$, and information requirement becomes higher for $w > 0.75$. Decision timeliness, which consists entirely of

partial evidence, also exceeds scalability when $w > 8/9$. Other criteria retain their relative ordering over the interior of the examined range, although several scores converge when w approaches one. These results show that the reported ordering depends on the declared treatment of partial evidence and should therefore be interpreted as a reporting summary, not as an invariant ranking of scientific importance.

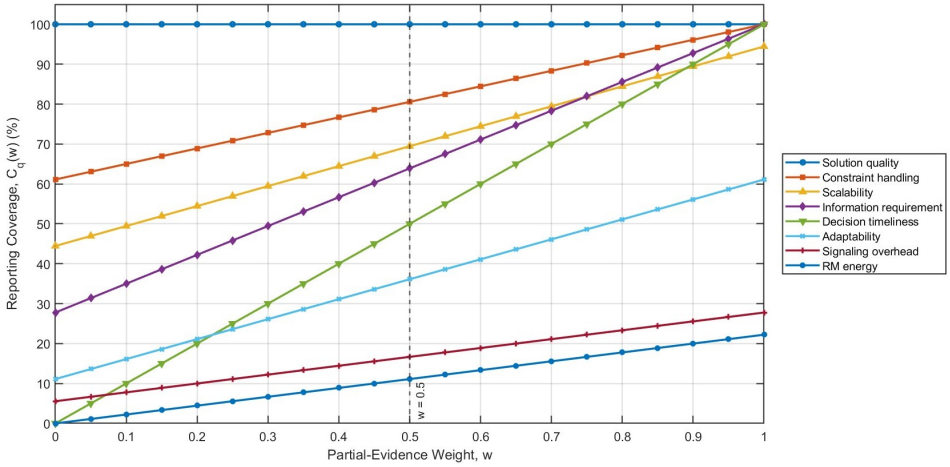


Figure 3. Sensitivity of reporting coverage to the partial-evidence weight w for all eight evaluation criteria.

Table 8. Evidence counts across the 18 reviewed primary studies.

Rank	Criterion	(✓)	(△)	–	$C_q(0.5)$
1	Solution quality	18	0	0	100.0%
2	Constraint handling	11	7	0	80.6%
3	Scalability	8	9	1	69.4%
4	Information requirement	5	13	0	63.9%
5	Decision timeliness	0	18	0	50.0%
6	Adaptability	2	9	7	36.1%
7	Signaling overhead	1	4	13	16.7%
8	RM energy	0	4	14	11.1%

Solution quality obtains complete coverage because every study reports the performance objective used to justify its proposed method. This consistency reflects the usual structure of resource allocation research, in which a new method is compared with several baselines using throughput, energy efficiency, latency, or another scenario-specific objective. However, the result does not show whether the allocation can be obtained and maintained under the timing and control limitations of a practical network.

Constraint handling ranks second because the centralized and distributed model-based methods combine feasibility directly. The six partial entries are dominated by learning-assisted methods. In these studies, constrained parameterization, reward penalties, and knowledge-driven architectures reduce the likelihood of an invalid output, but the evaluations generally do not report per-constraint violation probabilities. This distinction becomes important for URLLC, healthcare traffic, and safety-related control, where a small average violation may still hide an unacceptable event.

Scalability and information requirements rank in the middle. Several studies evaluate more users, links, or graph nodes, but most vary only one dimension while retaining the original wireless model. The resulting evidence establishes computational or architectural scaling, not robustness to another deployment. All studies state information requirements, but centralized and learning-assisted methods rarely measure the cost of acquiring the

required state or features. Distributed methods are clearer about their local messages, although message descriptions are rarely converted into total bytes or delay.

Decision timeliness has 50% coverage under the midpoint convention. The mapped studies provide partial timing evidence, but none measures the full resource allocation cycle relative to state validity. A method may reduce solver or inference time and still produce an invalid allocation when it includes channel measurement, information transfer, queueing, coordination, feasibility correction, and installation.

Adaptability receives substantially weaker coverage. Only two studies provide direct evaluation beyond their original operating configuration. The remaining partial evidence typically comes from varying node count, traffic within a predefined model, or another parameter already represented in training or optimization. These tests should not be interpreted as evidence that a method can respond to a different propagation environment, mobility pattern, traffic distribution, service constraint, or network architecture.

Signaling and resource-management energy rank last. Only one mapped study reports direct signaling evidence, while four additional studies provide packet counts, coordination procedures, or model-exchange descriptions that support an indirect assessment. No mapped study directly measures the complete energy consumed by resource management. The four partial entries report transmission, computation, UAV, or service energy but omit at least part of the control process.

The ranking therefore identifies a difference between producing a high-quality allocation and demonstrating an implementable decision process. Existing studies provide strong evidence that the proposed methods improve their intended objectives. They provide much less evidence that network information can be acquired, processed, exchanged, and converted into a feasible allocation before it becomes outdated and without eliminating the resulting performance or energy benefit.

5. Operational Design Principles and Open Challenges

The evidence mapping in Section 4 shows strong reporting of allocation objectives but substantially less reporting of the complete operational process. The four design principles below are recommendations synthesized from the IMT-2030-derived criteria in Section 2 and the evidence gaps in Section 4. Recommendation ITU-R M.2160 motivates the underlying capabilities [4], while the reviewed studies illustrate mechanisms and limitations that lead to the proposed guidance. The design principles should therefore be validated only for the relevant deployment.

5.1. Design Principles

The first design principle is validity-aware control. It follows from decision timeliness and adaptability: an allocation should be associated with the observation time and validity of the state used to calculate it. The network should activate the decision only when the measured complete cycle satisfies the state-validity and application deadlines. Snapshot-based optimization such as [23] motivates this requirement because high allocation quality under a fixed state does not show that the state remains current at activation.

The second design principle is efficient coordination. It follows from information requirement, signaling overhead, and scalability. Centralized, distributed, and federated methods can reduce one component of computation while increasing state collection or coordination. The relevant quantity is therefore the total information exchange required to obtain a usable allocation. Local message exchange in [12,13] and model exchange in [30] illustrate why message content, rounds, delivery delay, and retransmissions should be considered together.

The third design principle is feasibility-preserving adaptation. It follows from constraint handling and adaptability. A method should state the operating range for which its guarantees or empirical evidence apply, detect when the current condition leaves that range, and provide correction or fallback when needed. Explicit model constraints in [23], structured learning in [34], and long-term constrained learning in [37] illustrate different guarantee types; a long-term average constraint should not be interpreted as an instantaneous packet-deadline or outage guarantee.

The fourth design principle is net-benefit management. It follows from solution quality and resource-management energy. The performance gain of an allocation should be compared with the additional measurement, signaling, processing, model-update, and installation cost required to obtain it. Service-energy objectives such as [24,29] illustrate useful resource objectives, but they do not by themselves measure the energy consumed by the allocation process. Table 9 summarizes the proposed operational design principles.

Table 9. Operational design principles derived from the requirements-based evaluation.

Design Principle	Derivation from the Framework	Supporting Evidence and Required Verification
Validity-aware control	Decision timeliness and adaptability; Equations (1)–(3)	Snapshot-based optimization in [23] motivates checking state age before activation. Measure the complete cycle, component validity, and constraint satisfaction when the allocation is applied.
Efficient coordination	Information requirement, signaling overhead, scalability	Local message exchange in [12,13] and federated exchange in [30] motivate full coordination accounting. Report messages, encoded bytes, rounds, retransmissions, and delivery delay.
Feasibility-preserving adaptation	Constraint handling and adaptability; Equations (6) and (7)	Explicit constraints in [23], structured learning in [34], and long-term constraints in [37] require distinct guarantee interpretations. Verify transient violations, recovery time, and fallback behavior.
Net-benefit management	Solution quality and resource-management energy; Equation (9)	Service-energy objectives in [24,29] motivate separating service benefit from control cost. Measure the gain remaining after state acquisition, coordination, computation, updating, and installation.

5.2. Future Directions and Open Challenges

The first open challenge is maintaining feasibility when channel information becomes outdated. Existing studies generally calculate and verify a resource allocation using the same channel-state information (CSI). In practice, the channel may change before the allocation is applied, causing an initially feasible decision to violate its constraints. Future methods should measure how CSI validity affects feasibility and determine when a decision must be updated. Developing CSI-validity-aware constraint margins without making the allocation overly conservative remains an open problem.

The second open challenge is delayed control signaling. Distributed methods often assume synchronized iterations and reliable message delivery. In practical wireless networks, control messages may arrive late or be lost, causing nodes to make decisions using different network information. Waiting for every information arrival increases delay; however, processing with incomplete information can reduce suitability. Future methods should

therefore support delay-tolerant coordination and provide a feasible intermediate allocation when the coordination process is interrupted.

The third open challenge is adaptation to channel and traffic variation. Scalability evaluations typically increase the number of users while keeping the original operating conditions. Such evaluations demonstrate computational scalability but do not show whether the method can adapt to changes in wireless conditions. A method may support a larger network while losing performance when the channel changes. Future studies should evaluate these variations explicitly and determine whether the method requires correction, model updating, or complete retraining.

The fourth open challenge is control-plane energy. Most studies evaluate transmission energy without measuring the energy required to produce the resource allocation. This omitted cost can include information collection, signaling, computation, and model updating. Future research should establish a consistent measurement boundary and report the control-plane energy required per decision. The operational energy saved by an allocation should exceed this additional cost before the method is considered energy-efficient. Table 10 summarizes the open challenges for resource allocation techniques in next-generation wireless.

Table 10. Open challenges identified from the requirements-based evaluation.

Open Challenge	Current Limitation	Research Target
Outdated-CSI feasibility	Allocation uses outdated channel information	CSI-validity-aware constraints
Delayed control signaling	Ideal message delivery is assumed	Delay-tolerant coordination
Channel and traffic variation	Fixed operating conditions are assumed	Flexible adaptation
Control-plane energy	Only transmission energy is measured	Signaling and processing energy need to be considered

6. Conclusions

This paper presents a requirements-based review of resource allocation and optimization for 6G wireless networks. Eighteen representative primary studies were examined across centralized model-based, distributed, and learning-assisted or hybrid approaches using eight operational criteria derived from IMT-2030 capabilities. The comparison separates improvement in the allocation objective from evidence concerning the process required to obtain and maintain the decision. Under the declared midpoint convention for partial evidence, solution quality has 100.0% reporting coverage, followed by constraint handling at 80.6% and scalability at 69.4%. Decision timeliness remains only partially supported at 50.0%, while adaptability, signaling overhead, and resource-management energy have lower coverage of 36.1%, 16.7%, and 11.1%, respectively. These percentages summarize the selected evidence mapping and should not be interpreted as a ranking of method performance. Future evaluations should report the age of the network information at activation, the complete measurement-to-activation delay, relevant constraint guarantees or violation rates, total coordination signaling, adaptation after operating-condition changes, and the energy used by resource management. The appropriate architecture remains application-dependent; no reviewed method group is best for every 6G scenario.

Author Contributions: Conceptualization, G.R.R.D. and I.S.; Methodology, I.S.; Formal Analysis, G.R.R.D. and I.S.; Writing—Original Draft Preparation, G.R.R.D. and I.S.; Writing—Review and Editing, I.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. ITU-R. *IMT Vision—Framework and Overall Objectives of the Future Development of IMT for 2020 and Beyond*; Technical Report Recommendation ITU-R M.2083-0; International Telecommunication Union: Geneva, Switzerland, 2015.
2. Alhashimi, H.F.; Hindia, M.N.; Dimyati, K.; Hanafi, E.B.; Safie, N.; Qamar, F.; Azrin, K.; Nguyen, Q.N. A Survey on Resource Management for 6G Heterogeneous Networks: Current Research, Future Trends, and Challenges. *Electronics* **2023**, *12*, 647. [CrossRef]
3. Sefati, S.S.; Haq, A.U.; Nidhi; Craciunescu, R.; Halunga, S.; Mihovska, A.; Fratu, O. A Comprehensive Survey on Resource Management in 6G Network Based on Internet of Things. *IEEE Access* **2024**, *12*, 113741–113784. [CrossRef]
4. ITU-R. *Framework and Overall Objectives of the Future Development of IMT for 2030 and Beyond*; Technical Report Recommendation ITU-R M.2160-0; International Telecommunication Union: Geneva, Switzerland, 2023.
5. Islam, M.M.; Hasan, K.; Jeong, S.H. Performance Evaluation and Optimization for 6G Networks: A Survey of KPIs, Tools, and AI Models. *ICT Express* **2026**, *12*, 390–416. [CrossRef]
6. Mahboob, S.; Liu, L. Revolutionizing Future Connectivity: A Contemporary Survey on AI-Empowered Satellite-Based Non-Terrestrial Networks in 6G. *IEEE Commun. Surv. Tutor.* **2024**, *26*, 1279–1321. [CrossRef]
7. Ergen, M.; Saoud, B.; Shayea, I.; El-Saleh, A.A.; Ergen, O.; Inan, F.; Tuysuz, M.F. Edge Computing in Future Wireless Networks: A Comprehensive Evaluation and Vision for 6G and Beyond. *ICT Express* **2024**, *10*, 1151–1173. [CrossRef]
8. Alsaeedy, A.A.R.; Chong, E.K.P. A Survey of Mobility Management in Non-Terrestrial 5G Networks: Power Constraints and Signaling Cost. *IEEE Access* **2024**, *12*, 107529–107551. [CrossRef]
9. Dewa, G.R.R.; Sohn, I.; Djamari, D.W. Message Passing-Based Assignment for Efficient Handover Management in LEO Networks. *Telecom* **2025**, *6*, 76. [CrossRef]
10. de Faria, A.J.; Brito, J.M.C.; Spadoti, D.H.; Borges, R.M. A Broader Survey on 6G Radio Resource Management. *Sensors* **2026**, *26*, 2497. [CrossRef] [PubMed]
11. Shi, Y.; Lian, L.; Shi, Y.; Wang, Z.; Zhou, Y.; Fu, L.; Bai, L.; Zhang, J.; Zhang, W. Machine Learning for Large-Scale Optimization in 6G Wireless Networks. *IEEE Commun. Surv. Tutor.* **2023**, *25*, 2088–2132. [CrossRef]
12. Dewa, G.R.R.; Park, C.; Sohn, I. Hierarchical Affinity Propagation for Reliable Integrated Multilayer V2X Networks. *IEEE Internet Things J.* **2023**, *10*, 16316–16331. [CrossRef]
13. Dewa, G.R.R.; Park, C.; Sohn, I. Priority-Aware Scheduling for High-Dense Healthcare IoT (H-IoT) Networks Using Message-Passing Algorithm. *IEEE Internet Things J.* **2024**, *11*, 21604–21619. [CrossRef]
14. Alwarafy, A.; Abdallah, M.; Ciftler, B.S.; Al-Fuqaha, A.; Hamdi, M. The Frontiers of Deep Reinforcement Learning for Resource Management in Future Wireless HetNets: Techniques, Challenges, and Research Directions. *IEEE Open J. Commun. Soc.* **2022**, *3*, 322–365. [CrossRef]
15. Alhashimi, H.F.; Hindia, M.N.; Dimyati, K.; Hanafi, E.B.; Alden, F.Z.; Qamar, F.; Nguyen, Q.N. Survey on AI-Enabled Resource Management for 6G Heterogeneous Networks: Recent Research, Challenges, and Future Trends. *Comput. Mater. Contin.* **2025**, *83*, 3585–3622. [CrossRef]
16. Shen, Y.; Zhang, J.; Song, S.H.; Letaief, K.B. Graph Neural Networks for Wireless Communications: From Theory to Practice. *IEEE Trans. Wirel. Commun.* **2023**, *22*, 3554–3569. [CrossRef]
17. Dewa, G.R.R.; Kim, J.; Sohn, I. Message-Passing Neural Network (MPNN) for Large-Scale Healthcare IoT Networks in Smart City. *IEEE Internet Things J.* **2026**, *13*, 4765–4781. [CrossRef]
18. ITU-R. *Future Technology Trends of Terrestrial International Mobile Telecommunications Systems Towards 2030 and Beyond*; Technical Report Report ITU-R M.2516-0; International Telecommunication Union: Geneva, Switzerland, 2022.
19. Hossain, E.; Vera-Rivera, A. 6G Cellular Networks: Mapping the Landscape for the IMT-2030 Framework. *IEEE Trans. Technol. Soc.* **2025**, *6*, 377–392. [CrossRef]
20. Zhao, Q.; Lister, D.; Saxena, N.; Erfanian, J. *6G Requirements and Design Considerations*; Technical Report; NGMN Alliance: Frankfurt, Germany, 2023.
21. Zhang, X.; Debroy, S. Resource Management in Mobile Edge Computing: A Comprehensive Survey. *ACM Comput. Surv.* **2023**, *55*, 1–37. [CrossRef]

22. NGMN Alliance. *6G Position Statement: An Operator View*; Technical Report, Next Generation Mobile Networks Alliance; NGMN Alliance: Frankfurt, Germany, 2023. Available online: https://www.ngmn.org/wp-content/uploads/NGMN_6G_Position_Statement.pdf (accessed on 20 June 2026).
23. Wu, C.; Mu, X.; Liu, Y.; Gu, X.; Wang, X. Resource Allocation in STAR-RIS-Aided Networks: OMA and NOMA. *IEEE Trans. Wirel. Commun.* **2022**, *21*, 7653–7667. [CrossRef]
24. Zhai, Z.; Dai, X.; Duo, B.; Wang, X.; Yuan, X. Energy-Efficient UAV-Mounted RIS Assisted Mobile Edge Computing. *IEEE Wirel. Commun. Lett.* **2022**, *11*, 2507–2511. [CrossRef]
25. Nguyen, L.X.; Tun, Y.K.; Dang, T.N.; Park, Y.M.; Han, Z.; Hong, C.S. Dependency Tasks Offloading and Communication Resource Allocation in Collaborative UAVs Networks: A Meta-Heuristic Approach. *IEEE Internet Things J.* **2023**, *10*, 9062–9076. [CrossRef]
26. Liu, L.; Yuan, X.; Chen, D.; Zhang, N.; Sun, H.; Taherkordi, A. Multi-User Dynamic Computation Offloading and Resource Allocation in 5G MEC Heterogeneous Networks With Static and Dynamic Subchannels. *IEEE Trans. Veh. Technol.* **2023**, *72*, 14924–14938. [CrossRef]
27. Liu, Z.; Lan, Q.; Huang, K. Resource Allocation for Multiuser Edge Inference With Batching and Early Exiting. *IEEE J. Sel. Areas Commun.* **2023**, *41*, 1186–1200. [CrossRef]
28. Bai, J.; Wang, X. Distributed-Optimization With Centralized-Refining for Efficient Resource Allocation in Future Wireless Networks. *IEEE Trans. Commun.* **2024**, *72*, 4829–4843. [CrossRef]
29. Wu, W.; Yu, W.; Li, H.; Zhu, S.; Ma, Y.; Guan, X. Distributed Multidomain Resource Allocation for IIoT-Based Control Systems. *IEEE Trans. Ind. Inform.* **2024**, *20*, 14006–14016. [CrossRef]
30. Ji, Z.; Qin, Z. Federated Learning for Distributed Energy-Efficient Resource Allocation. In *Proceedings of the IEEE International Conference on Communications (ICC)*; IEEE: New York, NY, USA, 2022; pp. 1–6. [CrossRef]
31. Guo, J.; Yang, C. Learning Power Allocation for Multi-Cell Multi-User Systems With Heterogeneous Graph Neural Networks. *IEEE Trans. Wirel. Commun.* **2022**, *21*, 884–897. [CrossRef]
32. Wang, Y.; Li, Y.; Shi, Q.; Wu, Y.C. ENGNN: A General Edge-Update Empowered GNN Architecture for Radio Resource Management in Wireless Networks. *IEEE Trans. Wirel. Commun.* **2024**, *23*, 5330–5344. [CrossRef]
33. Peng, Y.; Guo, J.; Yang, C. Learning Resource Allocation Policy: Vertex-GNN or Edge-GNN? *IEEE Trans. Mach. Learn. Commun. Netw.* **2024**, *2*, 190–209. [CrossRef]
34. Yang, H.; Cheng, N.; Sun, R.; Quan, W.; Chai, R.; Aldubaikhy, K.; Alqasir, A.; Shen, X. Knowledge-Driven Resource Allocation for Wireless Networks: A WMMSE Unrolled Graph Neural Network Approach. *IEEE Internet Things J.* **2024**, *11*, 18902–18916. [CrossRef]
35. Akyildiz, H.A.; Gemici, O.F.; Hokelek, I.; Cirpan, H.A. Hierarchical Reinforcement Learning Based Resource Allocation for RAN Slicing. *IEEE Access* **2024**, *12*, 75818–75831. [CrossRef]
36. Huang, B.; Fan, X.; Zheng, S.; Chen, N.; Zhao, Y.; Huang, L.; Gao, Z.; Chao, H.C. Collaborative Sensing-Aware Task Offloading and Resource Allocation for Integrated Sensing-Communication- and Computation-Enabled Internet of Vehicles (IoV). *Sensors* **2025**, *25*, 723. [CrossRef] [PubMed]
37. Uslu, Y.B.; NaderiAlizadeh, N.; Eisen, M.; Ribeiro, A. Fast State-Augmented Learning for Wireless Resource Allocation with Dual Variable Regression. *IEEE Trans. Signal Process.* **2026**, *74*, 3553–3568. [CrossRef]
38. Yu, W.; Yang, X.; Yu, S. A V2X Communication Resource Allocation Method Based on Graph Neural Networks and Deep Reinforcement Learning. *Sci. Rep.* **2026**, *16*, 23131. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.