

Goal-Oriented Resource Allocation for 6G Radio Access Networks with Spatial Reuse

Maria Cecilia Fernández Montefiore*[‡], Federico Mason[†], Federico Chiariotti[†],
F. Javier López-Martínez*, and Andrea Zanella[†]

*Dept. Signal Theory, Networking and Communications, CITIC-UGR, University of Granada, Spain

[‡]Instituto de Investigaciones en Ingeniería Eléctrica (IIIE), Dpto. Ingeniería Eléctrica y de Computadoras,
Universidad Nacional del Sur (UNS) - CONICET, Bahía Blanca, Argentina

[†]Dept. Information Engineering, University of Padova, Italy

Emails: mcfm@ugr.es, federico.mason@unipd.it, federico.chiariotti@unipd.it, fjlm@ugr.es, andrea.zanella@unipd.it

Abstract—Compared to previous cellular generations, 6G pushes toward an unprecedented coupling of application semantics and network operation, blurring the traditional boundary between the generation, processing, and transport of information. In this context, the Goal-Oriented Communication (GoC) paradigm introduces the Value of Information (VoI) as a measure of packet importance from the perspective of the end application, going beyond classical network-layer metrics such as throughput and delay. In this work, we propose a resource allocation framework that integrates VoI into the optimization of MIMO-OFDMA in 6G radio access networks. Two optimization problems are formulated: a *sum value problem*, whose goal is to maximize the aggregate VoI of successfully delivered packets, and a *safety problem*, which aims at minimizing the VoI of undelivered packets. Our simulation results demonstrate that: (i) VoI-based allocation outperforms conventional throughput-maximization strategies in both scenarios; (ii) using spatial-layer based precoding strategies further improves performance over single-stream OFDMA. To our knowledge, this is the first work that integrates VoI into the joint optimization of spatial-layer dimensions with time-frequency resources, providing a foundation for the GoC-based resource management in 6G networks.

Index Terms—Value of information, goal-oriented communication, multi-user communications, delay-sensitive communications, non-orthogonal multiple access.

I. INTRODUCTION

UNLIKE previous cellular technologies, the sixth generation (6G) of mobile networks will be based on a tight native coupling of applications and network layers, breaking the long-standing classical boundary between the generation and transport of information [1]. Under this vision, 6G networks will not only improve aggregate network capacity but also fundamentally rethink the mechanisms through which single packet transmissions are scheduled and prioritized [2]. This paradigm change is expected to enable heterogeneous services

with unprecedented time-varying requirements in terms of latency, reliability, and throughput [3].

To achieve its ambitious objectives, 6G will embrace the goal-oriented communication (GoC) paradigm, which discriminates packets based on their importance to the end user [4]. In GoC scenarios, maximizing data throughput is no longer sufficient to ensure high application performance, as different packets within the same data flow may have different *values* [5]. In recent years, numerous works have proposed GoC strategies to optimize packet scheduling based on the value of information (VoI) of each packet [6]. However, most of these prior approaches rely on simplified communication models, and neglect lower-level resource-allocation aspects, such as number of resource blocks, channel gains, inter-user interference, and propagation delay.

The above limitation becomes particularly critical in 6G radio access networks with multiple-input multiple-output (MIMO) capabilities [7]. In particular, the spatial dimension of MIMO-enabled orthogonal frequency-division multiple access (OFDMA) systems [8] allows multiple co-channel users to share the same resource blocks, increasing the aggregate data rate. This leads to a significantly more complex resource allocation process, as the cross-layer optimization of precoding schemes, power allocation, and antenna configurations produces non-convex problems [9]. Therefore, on the one hand, most existing GoC frameworks fail to capture the interaction between VoI and the physical-layer resources that ultimately determine packet delivery. On the other hand, the literature on joint OFDMA and MIMO optimization focuses predominantly on classical, bit-centric performance metrics such as throughput and spectral efficiency [10], [11].

In this work, we introduce a novel GoC resource allocation framework for 6G radio access networks. Specifically, we consider a downlink (DL) transmission scenario in which a multi-antenna base station (BS) equipped with a digital twin (DT) of an environment needs to deliver updates to a set of single-antenna User Equipments (UEs). Each update is tailored to allow the receiver to perform a control task such as robotic manipulation, and is associated with a certain VoI.

As the BS has a full knowledge on the environment via its

This work was partly supported by grant PID2023-149975OB-I00 (COSTUME) funded by MICIU/AEI/10.13039/501100011033 and FEDER/UE, by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation program under grant agreement no. 101292933 (MAGIC-6G), by the European Commission as part of the 6G-LEADER Project under grant no. 101192080, and by grant DGP_PRED_2024_0212 funded by Consejería de Universidad, Investigación e Innovación (Junta de Andalucía) and the European Union.

DT, it can compute the necessary updates and their values, but it is limited by the channel capacity. The BS can then leverage the spatial resources provided by the multiple antennas to either optimize a single-user's transmission or accommodate multiple users on the same resource blocks, balancing inter-user interference against per-user rate. Differently from previous works using similar communication models, we do not evaluate performance solely in terms of throughput, but instead consider the VoI associated with each update.

More specifically, we target two different optimization scenarios: (i), a *sum value problem*, where the goal is to maximize the sum of the VoI of scheduled (and, hence, successfully-delivered) frames; (ii), a *safety problem*, where the goal is to minimize the VoI of data frames that remain unscheduled. The first problem is relevant for high-capacity interactive services, such as massive online gaming or holographic communication, where the aim is to maximize the total VoI delivered to the receiver. In contrast, the second problem addresses critical applications, such as remote surgery or autonomous driving systems, where it is fundamental to set an overall bound on the uncertainty of the system, even at the cost of efficiency.

To the best of our knowledge, this work represents the first attempt to integrate the concept of VoI into the joint optimization of communication resources over space, time, and frequency dimensions. For this reason, we consider a simple system model, which eases the interpretation of the results and sets a baseline for more advanced goal-oriented resource allocation strategies in 6G networks. Our experimental results show that, in both cases, VoI-based allocation outperforms conventional strategies that focus on maximizing communication throughput. Moreover, the adoption of spatial-layer precoding can further improve performance with respect to baseline OFDMA, although at the cost of higher complexity of resource allocation schemes.

II. SYSTEM MODEL

In this section, we define our communication model, explaining how OFDMA resources are allocated and, consequently, how to compute the throughput associated with each user. We use the following *mathematical notation*: boldface lowercase (**a**) and uppercase (**A**) letters denote vectors and matrices, respectively; transpose and conjugate transpose are denoted as $(\cdot)^T$ and $(\cdot)^H$; calligraphic letters, e.g., $\mathcal{S}$, denote sets, while $|\mathcal{S}|$ indicates the set's cardinality; finally, $\mathbb{E}\{\cdot\}$ is the expectation operator and $\|\cdot\|$ is the Euclidean norm.

A. Application Scenario

We consider an application in which a BS equipped with a DT must distribute updates about the state of the system to a set $\mathcal{U} = \{1, \dots, U\}$ of single-antenna users over the DL. Each update contains the information necessary to the user to perform its task. The length of the update to user u in bits is l_u , and the DT can evaluate the positive impact of its reception as the VoI v_u .

The BS has N antennas, which it uses to operate multi-user multiple-input single-output (MU-MISO) transmission. The

system operates according to an OFDMA framework, where the radio resources are organized into a 2D time-frequency grid with N_B frequency bins, with an individual subcarrier width of b and time frames of N_T time slots each, with a time slot duration t . Hence, at the beginning of each frame, the BS is required to allocate the resources to the users, considering the potential value of updates and optimizing application performance.

B. Communication System Model

We denote a specific resource element (or block) of the grid by the pair (k, n) , where $k \in \{1, \dots, N_B\}$ is the discrete-frequency index and $n \in \{1, \dots, N_T\}$ is the discrete-time index. The channel vector between the BS and user $u \in \mathcal{U}$ on resource block (k, n) is a N -dimensional complex column vector $\mathbf{h}_{u,k,n} \in \mathbb{C}^{N \times 1}$, modeled as

$$\mathbf{h}_{u,k,n} = \sqrt{\beta_u} \tilde{\mathbf{h}}_{u,k,n}, \quad (1)$$

where the vector $\tilde{\mathbf{h}}_{u,k,n} = [\tilde{h}_{u,k,n,1}, \dots, \tilde{h}_{u,k,n,N}]^T \in \mathbb{C}^{N \times 1}$ represents the small-scale fading, while $\beta_u \in \mathbb{R}$ represents the large-scale fading. In particular, β_u includes the effects of path loss and shadowing and is assumed to be constant over the $N_B \times N_T$ resource grid due to the stationarity of the users' positions relative to the frame duration.

We consider a 2D block-fading model, according to which the channel $\mathbf{h}_{u,k,n}$ is assumed to be constant over one resource block and independent and identically distributed (i.i.d.) across different resource blocks. Hence, $\tilde{\mathbf{h}}_{u,k,n}$ is a complex Gaussian random vector whose elements $\tilde{h}_{u,k,n,i} \sim \mathcal{CN}(0, 1)$ are i.i.d. standard complex normal random variables for $i = 1, \dots, N$. At the beginning of a frame, the BS selects a subset of users $\mathcal{K}_{k,n} \subseteq \mathcal{U}$ to be served for each resource element (k, n) . Notably, if $|\mathcal{K}_{k,n}| \leq 1$, there is at most one user per resource element, whereas if $|\mathcal{K}_{k,n}| > 1$, the BS exploits spatial multiplexing to serve multiple users. We define the *global set of scheduled users* $\mathcal{S} = \bigcup_{k=1}^{N_B} \bigcup_{n=1}^{N_T} \mathcal{K}_{k,n}$ as the union of all local allocation sets across the entire time-frequency grid. It is important to note that a specific user $u \in \mathcal{S}$ may be included in multiple local sets $\mathcal{K}_{k,n}$ across different resource blocks, depending on the resource allocation policy and the user's channel conditions. Therefore, the cardinality of the global set $|\mathcal{S}| \leq U$ represents the total number of unique users served within the frame.

C. Signal and Precoding Model

The BS applies linear precoding to the data symbols $x_{u,k,n}$ intended for each user $u \in \mathcal{K}_{k,n}$, where $\mathbb{E}[|x_{u,k,n}|^2] = 1$. The precoding vector $\mathbf{w}_{u,k,n} \in \mathbb{C}^{N \times 1}$ is expressed as:

$$\mathbf{w}_{u,k,n} = \sqrt{p_{u,k,n}} \bar{\mathbf{w}}_{u,k,n}, \quad (2)$$

where $p_{u,k,n} \geq 0$ denotes the power allocated to user u , and $\bar{\mathbf{w}}_{u,k,n}$ is the unit-norm beamforming vector ($\|\bar{\mathbf{w}}_{u,k,n}\| = 1$) determining the spatial direction of the transmission.

The transmitted signal vector $\mathbf{s}_{k,n}$ is:

$$\mathbf{s}_{k,n} = \sum_{u \in \mathcal{K}_{k,n}} \sqrt{p_{u,k,n}} \bar{\mathbf{w}}_{u,k,n} x_{u,k,n}, \quad \text{s. t.} \quad \sum_{u \in \mathcal{K}_{k,n}} p_{u,k,n} \leq P. \quad (3)$$

The signal received by user $u \in \mathcal{K}_{k,n}$ at resource (k, n) is:

$$y_{u,k,n} = \sqrt{p_{u,k,n}} \mathbf{h}_{u,k,n}^H \bar{\mathbf{w}}_{u,k,n} x_{u,k,n} + \sum_{j \in \mathcal{K}_{k,n}, j \neq u} \sqrt{p_{j,k,n}} \mathbf{h}_{u,k,n}^H \bar{\mathbf{w}}_{j,k,n} x_{j,k,n} + z_{u,k,n}, \quad (4)$$

where $z_{u,k,n} \sim \mathcal{CN}(0, \sigma^2)$. Therefore, the resulting signal-to-interference-plus-noise ratio (SINR) is:

$$\gamma_{u,k,n} = \frac{p_{u,k,n} |\mathbf{h}_{u,k,n}^H \bar{\mathbf{w}}_{u,k,n}|^2}{\sum_{j \in \mathcal{K}_{k,n}, j \neq u} p_{j,k,n} |\mathbf{h}_{u,k,n}^H \bar{\mathbf{w}}_{j,k,n}|^2 + \sigma^2}. \quad (5)$$

D. Transmission Scenarios

In this paper, we consider two transmission scenarios in which each resource block can be assigned to at most one or two users, respectively.

1) *Baseline OFDMA Allocation (BOA)*: In this case, we set $|\mathcal{K}_{k,n}| = 1, \forall (k, n)$. Therefore, there is no inter-user interference, and the BS allocates the total power of each block to the single scheduled user u^* , i.e., $p_{u^*,k,n} = P$. The beamformer strategy is designed using maximum ratio transmission (MRT), so that the precoding vector is

$$\bar{\mathbf{w}}_{u^*,k,n}^{\text{MRT}} = \frac{\mathbf{h}_{u^*,k,n}}{\|\mathbf{h}_{u^*,k,n}\|}, \quad (6)$$

while the signal-to-noise ratio (SNR) is

$$\gamma_{u^*,k,n} = \frac{P \|\mathbf{h}_{u^*,k,n}\|^2}{\sigma^2}. \quad (7)$$

2) *Spatial OFDMA Allocation (SOA)*: In this case, each user can either be paired with another user or remain unpaired, operating as in the BOA scenario. Consequently, at most two users can be allocated to the same resource block so that $|\mathcal{K}_{k,n}| \in \{1, 2\}, \forall (k, n)$. Furthermore, we assume that each user can be associated with at most one other user across the entire OFDMA grid. In other words, if $u_1 \in \mathcal{U}$ is paired with $u_2 \in \mathcal{U}$, then it cannot also be paired with a third user $u_3 \in \mathcal{U}$. Although this assumption restricts the solution space, it significantly simplifies the optimization problem, eliminating the need to consider configurations in which u_1 is paired with u_2 on some resource blocks and with u_3 on others.

In our system, when two users $\{u_1, u_2\}$ share the same resource block, the beamforming vectors $\bar{\mathbf{w}}_{u_1,k,n}$ and $\bar{\mathbf{w}}_{u_2,k,n}$ are designed based on a zero-forcing (ZF) precoding scheme, so that $\mathbf{h}_{u_1,k,n}^H \bar{\mathbf{w}}_{u_2,k,n}^{\text{ZF}} = 0$. Instead, the power coefficients $p_{u_i,k,n}$ can be optimized with any allocation policy satisfying $p_{u_1,k,n} + p_{u_2,k,n} \leq P$. As in the BOA scheme, there is no inter-user interference and the SNR for user $i \in \{1, 2\}$ is

$$\gamma_{u_i,k,n} = \frac{p_{u_i,k,n} |\mathbf{h}_{u_i,k,n}^H \bar{\mathbf{w}}_{u_{3-i},k,n}^{\text{ZF}}|^2}{\sigma^2}. \quad (8)$$

We remark that, with this transmission scheme, some resource blocks are assigned to pairs of users, while others are assigned to single users, whose SNR is as in (7).

III. GOAL-ORIENTED RESOURCE ALLOCATION

In this section, we present the model used to determine the user requirements in terms of VoI and data rate. We then define two optimization problems, representing high-throughput and mission critical 6G applications, respectively, and propose heuristic approaches to solve both problems, taking advantage of spatial multiplexing of MIMO system.

A. Resource Requirements

We denote by b the bandwidth and by t the time duration of a resource block. The average resource capacity (in bits) for user u is

$$\bar{d}_u = \frac{1}{N_B N_T} \sum_{k=1}^{N_B} \sum_{n=1}^{N_T} b t \log_2(1 + \gamma_{u,k,n}). \quad (9)$$

For each user u , we indicate with l_u the total number of bits to be transmitted to that user in the current frame and by v_u the associated VoI. The average number of resource blocks required to successfully deliver the payload to the user u is $l_u / \bar{d}_u$. We assume that successful message decoding requires that the BS transmits all the l_u bits. Therefore, if the available resources are not sufficient to carry the whole message to a certain user, that user is excluded from the allocation.

Given the pair (l_u, v_u) for each user $u \in \mathcal{U}$, the goal is then to allocate the resources in the OFDMA time-frequency grid to maximize performance during the current frame. We introduce a decision variable $x_u \in \{0, 1\}$, where $x_u = 1$ if user u is scheduled for transmission ($\mathcal{S}$), and $x_u = 0$ otherwise, and denote by $M = N_T N_B$ the total number of resources.

For simplicity, in the allocation problem we assume that a message of l_u bits is delivered successfully if the BS allocates $\bar{r}_u$ resource blocks to the user u . Using the average capacity given by (9) in place of the instantaneous capacity, we can consider all resource blocks equivalent, strongly decreasing computational complexity. However, the actual capacity allocated to users may turn out to be lower than expected and possibly not sufficient to carry l_u bits. In this case, the user u will eventually not be scheduled.

To limit this risk, we consider the resource requested by the user u as equivalent to

$$\bar{r}_u = \lceil \alpha l_u \bar{d}_u^{-1} \rceil, \quad (10)$$

where $\alpha \geq 1$ is a *margin factor*. Given the above definitions, we consider two optimization scenarios, characterized by different performance metrics, which are described below.

B. The Sum Value Problem (SVP)

The objective is to maximize the sum-VoI of the packets delivered to the users:

$$\arg \max_{x_u \in \{0,1\}^K} \sum_{u \in \mathcal{U}} x_u v_u \quad \text{s.t.} \quad \sum_{u \in \mathcal{U}} x_u \bar{r}_u \leq M. \quad (11)$$

This problem can be cast into a *fractional knapsack* problem [12], where M is the system capacity.

As the knapsack problem is NP-hard, we consider its continuous relaxation: we sort users according to their value-to-cost ratio $v_u/\bar{r}_u$ and greedily assign users to the OFDMA resource blocks, prioritizing those users with a higher value-to-cost ratio, where the cost corresponds to the number of required resource blocks. We note that, in this case, if user u cannot be served due to a lack of resources, serving a user h with $v_h/\bar{r}_h < v_u/\bar{r}_u$ can still improve performance.

C. The Safety Problem (SP)

In this case, our goal is to limit the maximum uncertainty over all users, e.g., to respect a safety requirement for industrial equipment. Our objective function is thus over the maximum VoI among unscheduled users, limiting the loss of valuable information. The problem can be posed as follows:

$$\arg \max_{x_u \in \{0,1\}^K} 1 - \max_{u \in \mathcal{U}} (1 - x_u) v_u \quad \text{s.t.} \quad \sum_{u \in \mathcal{U}} x_u \bar{r}_u \leq M. \quad (12)$$

A straightforward approach to this problem is to allocate resources to users in a greedy manner, sorting them in descending order of their VoI as long as the requested resource budget, calculated as $\sum_{u \in \mathcal{S}} x_u \bar{r}_u$, is sustainable. We note that, if user u cannot be served, then any other user h with $v_h < v_u$ will not contribute to improving the objective function.

D. Computational Complexity

For BOA, the complexity is determined by the user's sorting according to the value-to-cost ratio and the VoI for the sum-value problem (SVP) and safety problem (SP) cases, respectively. In the case of U users, the computational cost is $C_{\text{BOA}} = U \cdot \log(U)$, as in the case of fractional knapsack.

For SOA, the problem becomes much more complex because the (average) capacity per block of a user $u_1 \in \mathcal{U}$ depends on the specific user u_2 that has been paired with u_1 . We should then consider all the possible ways to combine users into pairs, considering also the possibility that some users remain alone. The complexity is therefore at least combinatorial with the number of users, and grows as $U^{U/2}$.

To make the optimization practical, we resort to a sub-optimal lower-complexity scheme. We start from the solution of the BOA problem, which returns a first group of users, whose aggregated data rate can be supported without sharing resource blocks. Then, we consider the non-scheduled user u^* with the highest value-to-cost ratio in SVP (or VoI in SP) and try to pair u^* with the users in the scheduled group. If the allocation is successful, we repeat the procedure iteratively until the residual resources become insufficient.¹ Hence, the complexity of this solution is $C_{\text{SOA}} = U^3$.

IV. EXPERIMENTAL RESULTS

In this section, we analyze the proposed VoI-aware resource allocation schemes against benchmark methodologies that do not consider VoI. We evaluate performance as a function of

¹This problem can be tackled with other heuristic algorithms [13], [14], but the study of efficient techniques to exploit spatial multiplexing in OFDMA is outside the scope of this paper and is left to future work.

capacity and user density, as well as the correlation between VoI and the size of the transmitted data. Finally, we show the benefits and drawbacks of using a SOA approach instead of BOA in allocating OFDMA resources.

A. Benchmark Strategies and Simulation Settings

As a benchmark for our analysis, we consider three allocation strategies that are agnostic to VoI and thus represent the state-of-the-art methodologies currently used in practical mobile communication systems. More specifically, we order the users according to three different criteria:

- *Smallest Packet First (SPF)*: sorted in ascending order of frame sizes l_u , $\forall u \in \mathcal{U}$;
- *Largest Packet First (LPF)*: sorted in descending order of frame sizes l_u , $\forall u \in \mathcal{U}$;
- *Random*: randomly ordered.

Resources are allocated to the users sequentially, starting from the first user in the sorted list and proceeding until all resources are exhausted. All the benchmark strategies exploit the average resource requirement $\bar{r}_u$ to avoid the complexity of the original allocation problems, with $\alpha = 1$. Note that the benchmark strategies neglect VoI and consequently behave the same in SVP and SP. In contrast, our approach, named *VoI-based*, yields different allocation strategies for the two problems.

To compare the different strategies under realistic conditions, we consider an OFDMA resource grid with $N_B = 12$ frequency blocks and $N_T \in [14, 80]$ time blocks, so that the total number of resources varies as $M \in [12, 948]$ during resource sweep analysis. The individual subcarrier bandwidth is set to $b = 15$ kHz, and the time slot duration is $t = 1/14$ ms. We consider the number of users ranging in the interval $[2, 40]$ to assess scalability, and set the number of antennas at the BS to $N = 4$, so that we can concurrently serve two users in the same resource block in the case of SOA.

To model the wireless channel, we assume that the U users are uniformly distributed within a ring-shaped cell region defined by an inner radius $r_{\min} = 5$ m and an outer radius $r_{\max} = 100$ m from the BS. In particular, we consider a log-distance large-scale path loss with an exponent $\gamma = 2.45$ at a carrier frequency of $f = 1$ GHz, a transmit power of 10 dBm, and a frequency-selective small-scale Rading fading.

To model the requirements of each user $u \in \mathcal{U}$, we consider that the packet size is given by $l_u = l_{\min} + (l_{\max} - l_{\min}) \cdot (\nu_u \cdot \rho + \zeta_u \cdot (1 - \rho))$, where ν_u and ζ_u are uniformly distributed in $[0, 1]$, $l_{\min} = 100$ bits and $l_{\max} = 1000$ bits. Instead, the VoI associated with the packet is given by $v_u = \nu_u \cdot \rho + \omega_u \cdot (1 - \rho)$, where ω_u is uniform in $[0, 1]$, while $\rho \in [0, 1]$ is a hyperparameter denoting how much v_u and l_u are correlated. In the rest of the section, whenever not specified, we consider $\rho = 0$, assuming that v_u is independent of the frame size l_u .

B. Sum Value Problem (SVP)

In this subsection, we evaluate the performance of the proposed under the SVP scenario, contrasting the VoI-based solution against the benchmarks, named *SPF*, *LPF* and *random*. We vary the number of resource blocks (M), the number

of users (U), and the VoI-to-packet-size correlation (ρ). The comparison metric is defined as the ratio between the aggregate VoI achieved via (11), and the total expected VoI when all users in the network are scheduled. Thus, the normalized performance score lies in the interval $[0, 1]$, where a value of one corresponds to the full delivery of all users' frames.

Fig. 1a illustrates the system performance score as a function of the total number of resources M . As expected, all strategies increase their performance as M increases because the system can accommodate more transmissions. Crucially, the proposed VoI-based strategy consistently outperforms all benchmarks, even when resources are scarce ($M < 300$). To assess scalability, in Fig. 1b we report the performance when the number of users U increases from 2 to 40. While for higher loads the performance score decreases for all methods, the margin of VoI-based allocation with respect to the benchmark strategies progressively widens, with a 10% gain with respect to SPF, the best performing benchmark in this case.


Finally, in Fig. 1c we report the performance curves when varying the correlation factor ρ between VoI and the packet size. We observe that the VoI-based method achieves the best performance across the whole range, although the gap with respect to the best-performing benchmark reduces for higher correlation, i.e., when frames with higher VoI tend to be longer. In this case, the LPF strategy, which prioritizes users with longer frames, implicitly maximizes the aggregate VoI, thus approaching the VoI-based performance. However, this benchmark performs poorly when the correlation index is small, while the proposed approach consistently outperforms all VoI-agnostic methods.

C. Safety Problem


Next, we consider the SP scenario, where the primary objective is to mitigate the risk of leaving critical, high-value frames unscheduled. In this case, we evaluate the performance score defined in (12), which still lies within the interval $[0, 1]$; however, lower values are associated with more effective resource allocation strategies. The results are studied by varying the same three parameters considered in the previous case, namely the number of resources (M), the number of users (U) and the correlation factor (ρ).

Fig. 2a highlights the performance score when varying the number M of available resource blocks. When the resource budget is low, the benchmarks fail catastrophically, often dropping high-value frames. Instead, the proposed VoI-based strategy prioritizes users in ascending order of values to systematically exhaust the resource budget, leaving unscheduled users with the lowest VoI. Therefore, our solution consistently outperforms all benchmarks, although the gain clearly vanishes when resources are sufficient to allocate almost all users.


Fig. 2b analyzes the impact of network congestion. As U increases, the likelihood of competing high-priority frames increases, making all solutions lose performance. However, the VoI-based method exhibits a more graceful performance drop than the benchmarks. Finally, in Fig. 2c, we study the effect of the correlation factor ρ . Once again, the VoI-based



(a) Performance vs resources, with $U = 16$ and $\rho = 0$.



(b) Performance vs users, with $M = 168$ and $\rho = 0$.



(c) Performance vs correlation factor, with $M = 168$ and $U = 16$.


Fig. 1: Sum Value Problem (SVP).

strategy outperforms VoI-agnostic methods for all the values of ρ . Moreover, the performance curves of the frame size-based heuristics, namely SPF and LPF, exhibit a strong dependence on the correlation coefficient, with neither consistently outperforming the other across all scenarios.


D. Spatial OFDMA

In this subsection, we perform a comparative analysis between BOA and SOA, evaluating the benefits of aggregating multiple users over the same time-frequency block by precoding ZF. Fig. 3 illustrates the comparison between the two approaches in the SVP and SP scenarios, while varying the number of resources (M). We observe that for SVP, SOA provides a pronounced performance improvement over BOA on all values of M . This is because SOA effectively increases the aggregate capacity of the resource blocks, allowing the delivery of a larger number of frames and therefore a larger aggregate VoI, within the same physical bandwidth.


In the SP problem, the results show that the multiplexing capability of SOA can dramatically reduce the VoI of unscheduled users, which can be vital for mission-critical networks.



(a) Performance vs resources, with $U = 16$ and $\rho = 0$.



(b) Performance vs users, with $M = 168$ and $\rho = 0$.



(c) Performance vs correlation factor, with $U = 16$ and $M = 168$.

Fig. 2: Safety Problem (SP).

Under BOA, the system is fundamentally bounded by one user per block, creating a difficult bottleneck when multiple critical users require simultaneous emergency transmissions. The SOA approach mitigates this limitation by sharing OFDMA blocks among paired users, freeing up remaining time-frequency resources to serve other nodes. This multi-domain optimization ensures that the maximum unserved VoI drops rapidly toward zero as M is sufficiently large.

V. CONCLUSIONS

In this work, we propose a framework that integrates the GoC paradigm into the optimization of OFDMA systems. We consider a general scenario where a central BS is tasked with transmitting data to a set of users, each associated with dual communication requirements expressed in terms of throughput and VoI, respectively. We analyze two different optimization scenarios capturing the characteristics of traditional high-throughput applications and mission-critical services, and compare a VoI-based approach against benchmarks that are agnostic to the value of the transmitted data.




Fig. 3: Comparing baseline and spatial resource allocation, with $U = 16$ and $\rho = 0$.

Our results demonstrate that traditional allocation techniques are highly inefficient when VoI is adopted as the primary performance metric, and that the adoption of spatial multiplexing techniques can yield significant additional performance gains. Designing end-to-end GoC systems, where each layer of the protocol stack is optimized according to the actual value of each piece of information, represents a fundamental priority for future telecommunication networks.

REFERENCES

- [1] H. Zhou, Y. Deng, X. Liu, N. Pappas, and A. Nallanathan, "Goal-oriented semantic communications for 6g networks," *IEEE Internet of Things Magazine*, vol. 7, no. 5, pp. 104–110, 2024.
- [2] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Netw.*, vol. 34, no. 3, pp. 134–142, 2019.
- [3] E. Mozaffariahrar, F. Theoleyre, and M. Menth, "A survey of Wi-Fi 6: Technologies, advances, and challenges," *Future Internet*, vol. 14, no. 10, p. 293, 2022.
- [4] T. M. Getu, G. Kaddoum, and M. Bennis, "A survey on goal-oriented semantic communication: Techniques, challenges, and future directions," *IEEE Access*, vol. 12, pp. 51 223–51 274, 2024.
- [5] J. Holm, F. Chiariotti, A. E. Kalør, B. Soret, T. B. Pedersen, and P. Popovski, "Goal-oriented scheduling in sensor networks with application timing awareness," *IEEE Trans. Commun.*, vol. 71, no. 8, pp. 4513–4527, 2023.
- [6] P. Talli, E. D. Santi, F. Chiariotti, T. Soleymani, F. Mason, A. Zanella, and D. Gündüz, "Pragmatic Communication for Remote Control of Finite-State Markov Processes," *IEEE J. Sel. Areas Commun.*, vol. 43, no. 7, pp. 2589–2603, 2025.
- [7] B. Clerckx, Y. Mao, Z. Yang, M. Chen, A. Alkhateeb, L. Liu, M. Qiu, J. Yuan, V. W. Wong, and J. Montojo, "Multiple access techniques for intelligent and multifunctional 6G: Tutorial, survey, and outlook," *Proc. IEEE*, vol. 112, no. 7, pp. 832–879, 2024.
- [8] F. Riera-Palou and G. Femenias, "Cluster-based cooperative MIMO-OFDMA cellular networks: Scheduling and resource allocation," *IEEE Trans. Veh. Technol.*, vol. 67, no. 2, pp. 1202–1216, 2017.
- [9] K.-H. Lee and D. Kim, "Cross-layer optimization for heterogeneous MU-MIMO/OFDMA networks," *Sensors*, vol. 21, no. 8, p. 2744, 2021.
- [10] G. Femenias and F. Riera-Palou, "Scheduling and resource allocation in downlink multiuser MIMO-OFDMA systems," *IEEE Trans. Commun.*, vol. 64, no. 5, pp. 2019–2034, 2016.
- [11] Y.-A. Tai, Y.-W. Liu, P.-K. Kao, T.-Y. Lee, C.-I. Hu, and C.-Y. Li, "LOMAS: Latency-driven OFDMA Scheduling Design in Wi-Fi 6 Networks," in *2025 IEEE 33rd International Conference on Network Protocols (ICNP)*. IEEE, 2025, pp. 1–11.
- [12] S. Martello and P. Toth, *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- [13] M. A. Sedaghat and R. R. Müller, "On User Pairing in Uplink NOMA," *IEEE Trans. Wireless Commun.*, vol. 17, no. 5, pp. 3474–3486, 2018.
- [14] M. S. Ali, H. Tabassum, and E. Hossain, "Dynamic User Clustering and Power Allocation for Uplink and Downlink Non-Orthogonal Multiple Access (NOMA) Systems," *IEEE Access*, vol. 4, pp. 6325–6343, 2016.