

# Vector Coded Caching Multiplicatively Boosts MU-MIMO Systems Under Practical Considerations

Hui Zhao and Petros Elia

**Abstract**—This work presents a first comprehensive analysis of the impact of vector coded caching (VCC) in multi-user multiple-input multiple-output (MU-MIMO) systems with multiple receive antennas and variable pathloss — two key factors that critically influence systems with inherent MU unicasting behavior. We investigate two widely adopted precoding strategies: (i) block-diagonalization (BD) at the transmitter combined with maximal ratio combining (MRC) at the receivers, and (ii) zero-forcing (ZF) precoding. Our analysis explicitly accounts for practical considerations such as channel fading, channel state information (CSI) acquisition overhead, and fairness-oriented power allocation.

Our contributions span both analytical and simulation-based fronts. On the analytical side, we derive analytical expressions for the achievable throughput under BD-MRC and ZF, highlighting the performance benefits of equipping multi-antenna users with cache-aided interference management. Specifically, we develop a low-complexity BD-MRC optimization method that leverages matrix structure to significantly reduce the dimensionality involved in precoding computation, followed by solving the associated max-min fairness problem through an efficient one-dimensional search. In the massive MIMO regime, an asymptotic expression for the achievable throughput over Rayleigh fading channels is also derived. Simulations validate our theoretical results, confirming that VCC delivers substantial performance gains over optimized cacheless MU-MIMO systems. For example, with 32 transmit antennas and 2 receive antennas per user, VCC yields throughput improvements exceeding 300%. These gains are further amplified under imperfect CSI at the transmitter, where VCC's ability to offload interference mitigation to the receivers ensures robust performance even in the face of degraded CSI quality and elevated acquisition costs.

**Index Terms**—Coded caching, max-min fairness, precoding, power allocation, and MU-MIMO.

## I. INTRODUCTION

Coded caching was first introduced in the context of the single-stream Broadcast Channel (BC) in [2], where  $K$  cache-aided receivers request different files from a shared content library. This approach exploits each user's ability to store a fraction  $\gamma \in [0, 1]$  of the library, theoretically achieving  $K\gamma + 1$  degrees of freedom (DoF), often referred to as the coded caching gain [3]–[5], and it roughly corresponds to the number of users served at a time. However, practical implementations faced challenges due to the combinatorial nature of the scheme

This work was supported in part by EURECOM's industrial members: ORANGE, BMW, SAP, iABG, Norton LifeLock; in part by ERC-PoC Project LIGHT (Grant 101101031) and EU H2030 project CONVERGE, by the French projects EEMW4FIX (ANR) and PERSEUS (PEPR-5G); and in part by the Huawei France funded Chair towards Future Wireless Networks. Part of this work was presented at the International ITG 26th Workshop on Smart Antennas and 13th Conference on Systems, Communications, and Coding (WSA&SCC 2023) [1].

Hui Zhao and Petros Elia are with the Communication Systems Department, EURECOM, 06410 Sophia Antipolis, France (email: hui.zhao@eurecom.fr; petros.elia@eurecom.fr).

in [2], which requires each file to be divided into  $\binom{K}{K\gamma}$  non-overlapping subfiles. As a result, finite file sizes impose a fundamental constraint, reducing the real DoF to  $\Lambda\gamma + 1$ , where  $\Lambda$  ( $\Lambda \ll K$ ), constrained by the file size, denotes the number of distinct cached contents — equivalently referred to as *cache states* — that can be stored at the users. In practice, these constraints frequently limit achievable DoF to small single-digit values [6].

To overcome these limitations, research has explored the integration of coded caching with multi-antenna transmission, leading to the concept of *multi-antenna coded caching*. Notable related contributions include the works of Sharatiapanahi et al. [7] and Naderializadeh et al. [8], as well as studies on physical-layer (PHY) multi-antenna coded caching [9]. Other investigations have examined the scalability of content delivery rates in massive multiple-input multiple-output (MIMO) and cell-free networks [10], [11], while recent efforts have focused on cache-aided precoders for mitigating inter-stream interference [12] and location-aware caching for wireless extended reality applications [13]. Further contributions can be found in [3], [5], [14]–[18] and related works.

Despite these advancements, the early multi-antenna coded caching strategies remained constrained by file-size limitations, and as a consequence, their effective DoF remained upper-bounded by  $\Lambda\gamma + Q$  [5], where  $Q$  represents the multiplexing gain enabled by multiple antennas. This limitation simply meant that the presented caching gains were dwarfed by conventional multiplexing gains, particularly in massive MIMO where  $\Lambda\gamma \ll Q$ , as extensively validated in field trials [19].

This limitation was overcome with *vector coded caching* (VCC) in [20], whose introduction significantly enhanced the role of caching in multi-antenna systems. Unlike previous XOR-based approaches, VCC does not rely on multicasting; instead, it employs a clique-based structure on vectors, achieving a real DoF of  $Q(\Lambda\gamma + 1)$  under file-size constraints. This represents a *multiplicative DoF gain* over the cacheless case (DoF  $Q$ ), surpassing the additive improvement of previous XOR-based methods (real DoF  $\Lambda\gamma + Q$ ). Notably, this performance enhancement comes with a manageable subpacketization cost. Recently, VCC, also referred to as *signal-level coded caching* [21], has evolved into a broad class of schemes extensively investigated from an information-theoretic perspective, typically focusing on the delivery-load–subpacketization tradeoff (e.g., [22]–[24]). Moving beyond the infinite-SNR DoF analysis, the study in [21] examined the delivery performance of VCC at finite SNRs over wireless fading channels through numerical optimization, incorporating user locations into the cache placement phase. Another line of research in [25] provided a first statistical analysis of VCC under practical SNR regimes, focusing on

the case of single-antenna receivers as well as symmetric Rayleigh fading channels, incorporating beamforming gains and CSI acquisition costs, thus providing early evidence of the significant potential of VCC in wireless systems, demonstrating a substantial multiplicative improvement in spectral efficiency over conventional (cacheless) multi-user (MU) multiple-input single-output (MISO) systems.

#### A. Limitations of Existing Works and Technical Challenges

Going beyond the preliminary study in [25], our work here examines for the first time critical factors pertaining to a more realistic implementation of VCC in wireless systems, including the impact of channel asymmetry due to pathloss, the emergence of multi-antenna receivers, and the need for power allocation. In particular, channel asymmetry among users is crucial due to the near-far bottleneck associated to coded caching techniques, where this bottleneck has been shown to significantly reduce spectral efficiency in finite-SNR regimes [26]–[29]. Similarly, the growing adoption of multi-antenna receivers [30]–[32] necessitates an investigation into how they enhance VCC performance, particularly under heterogeneous pathloss conditions. These practical considerations introduce new challenges, such as determining the optimal multiplexing gain  $Q^*$ , and thus balancing the intricate tradeoff between multiplexing and beamforming gains as well as CSI overhead [25]. Moreover, multi-antenna receiver combining introduces a complex interaction between transmit precoding and receiver processing [33]–[35].

This same channel asymmetry studied here, now also raises concerns regarding user fairness, typically addressed via max-min fairness (MMF) techniques [12], [21], which optimize power allocation to maximize the minimum transmission rate. In VCC, analyzing such fairness aspects is automatically more involved, with optimizations that tend to be even more complex due to the potentially larger number of users involved (MMF leads to a non-convex, NP-hard problem [36]), leading to potentially higher computational complexity, thus further complicating resource allocation.

Together, these challenges — including pathloss variations, heterogeneous receiver antenna configurations, multiplexing-beamforming tradeoffs, and MMF constraints — make joint optimization particularly difficult. Even with techniques like block diagonalization (BD) precoding and maximal ratio combining (MRC) [37]–[40], deriving closed-form spectral efficiency expressions remains extremely challenging. These challenges are tackled here, in the presence of cache-aided interference management.

#### B. Contributions and Paper Structure

Our work captures the above crucial ingredients of fading, pathloss, fairness-based power allocation, CSI overhead, linear precoding, and receiver combining, offering a first-ever comprehensive study of VCC in practical MU-MIMO settings. Given that Video-on-Demand (VoD) traffic accounts for over 70% of total network data consumption [41], efficient caching-based delivery mechanisms are essential in alleviating network congestion. By addressing these challenges, our

study provides a systematic framework for enhancing spectral efficiency in modern wireless networks through VCC. Our main contributions are as follows:

- 1) We provide a rigorous analytical framework for VCC under BD-MRC, focusing on both precoding and receiver-side processing. Specifically, we leverage the idempotent and Hermitian properties of projection matrices to design a low-complexity BD precoder. Notably, for a given transmitted symbol's precoding column, this approach allows us to replace the eigenvalue decomposition of an  $L \times L$  matrix (where  $L$  is the number of base-station (BS) antennas) with the decomposition of a matrix whose dimension is determined solely by the number of receive antennas at the target user. Since the BS typically has far more antennas than any individual user, this substantially reduces computational complexity. Furthermore, we derive a theoretical expression for the transmission rate of a specific user and subsequently optimize the power allocation across multiple symbols intended for that user to maximize the achievable transmission rate under the constraint of the user's total allocated transmission power.
- 2) We then analytically determine the optimal effective sum-rate (cf. Definition 1) of the BD-MRC based VCC under the MMF requirement. To address the NP-hard complexity of MMF optimization, we transform it into a simple one-dimensional search problem where the objective equality has a unique root. Furthermore, we analytically derive the upper and lower bounds of the root, which substantially constrain the search range for determining the root numerically, making the root-finding process more efficient.
- 3) We subsequently focus on Rayleigh fading channels with variable pathloss for each user and analyze the delivery performance of the optimized BD-MRC based VCC in the massive MIMO regime.<sup>1</sup> Despite the requirement of a one-dimensional search, the simplified objective equality significantly reduces the search complexity. Furthermore, assuming an equal number of receiving antennas per user, we derive a simple closed-form expression for the optimal effective sum-rate. This simple expression not only allows us to bypass tedious numerical optimization but also crisply reflects the impact of the various factors under consideration.
- 4) Additionally, we analyze the performance of Zero-Forcing (ZF) precoding for multi-antenna receivers in VCC over Rayleigh fading channels with pathloss, and derive closed-form expressions for lower and upper bounds on the effective sum-rate, in the presence of MMF. Our numerical results validate these bounds as accurate approximations, demonstrating that for practical receiver configurations in VCC, ZF precoding optimized based on pathloss closely matches the performance of

<sup>1</sup>We also consider the special case of single-antenna receivers under BD precoding and the MMF requirement, and derive the closed-form expression for the optimal sum-rate, which was rightfully identified as a hard problem in [38]. We refer to [42, Lemma 4.3] for more details.

optimized BD-MRC.<sup>2</sup>

*Notations:* For a positive integer  $a$ , we use  $[a]$  to denote the set  $\{1, 2, \dots, a\}$ .  $|\cdot|$  denotes either the cardinality of a set or the magnitude of a complex number.  $\mathbb{C}$  denotes the set of complex numbers, while  $\mathbb{E}\{\cdot\}$  and  $\text{Tr}\{\cdot\}$  represent the average operator and the trace operator respectively.  $\text{Diag}\{a_1, \dots, a_n\}$  denotes a diagonal matrix with diagonal elements  $a_1, \dots, a_n$ , while  $\Psi \setminus \phi$  denotes the set  $\Psi$  after removing the element  $\phi$ . We will use  $(x)^+ \triangleq \max(x, 0)$ . All vectors are column vectors. Furthermore, we use  $\mathbf{A}^T$ ,  $\mathbf{A}^*$  and  $\mathbf{A}^H$  to denote the non-conjugate transpose, conjugation and conjugate transpose of the matrix  $\mathbf{A}$  respectively.  $[\mathbf{A}]_{\ell, \vartheta}$  denotes the  $(\ell, \vartheta)$ -th element in  $\mathbf{A}$ . Also,  $\mathbf{0}_L \in \mathbb{C}^L$  denotes the vector with all zero elements, and  $\mathbf{I}_L \in \mathbb{C}^{L \times L}$  represents the identity matrix. We additionally use  $X \sim \mathcal{Y}$  to denote that  $X$  follows the statistical distribution  $\mathcal{Y}$ , and we use  $\mathcal{CN}(\mathbf{m}, \Sigma)$  to denote complex Gaussian distribution with mean vector  $\mathbf{m}$  and covariance matrix  $\Sigma$ . As  $L \rightarrow \infty$ , we write  $f(L) \simeq g(L)$  to indicate that  $\lim_{L \rightarrow \infty} |f(L) - g(L)| = 0$ . Finally, we use  $\|\cdot\|$  to represent the  $\ell_2$ -norm operator. Moreover, the main variables used in the paper are summarized in Table I.

*Paper Structure:* We introduce the system model in Section II. Then in Section III we elaborate on BD-MRC based VCC, and in Section IV we present its corresponding analysis. Subsequently, in Section V, we mathematically analyze the performance of the ZF precoder for multi-antenna receivers. Numerical results and performance comparisons are presented in Section VI, while Section VII concludes the paper.

## II. SYSTEM MODEL

We consider a single-cell downlink MU-MIMO system, where a BS equipped with  $L$  antennas serves  $K$  cache-aided multi-antenna users, each requesting a different file from a file library  $\mathcal{F}$  of  $N$  ( $N \geq K$ ) equal-sized files. The BS has full access to the library  $\mathcal{F}$ , while each user can only cache a fraction  $\gamma \in [0, 1]$  of the library content, which is to be done during off-peak hours. The scheme is presented in detail in Section III-A, and it directly builds on the VCC principles from [20], adapted here for the setting of multi-antenna receivers. Due to the aforementioned finite file-size constraints, we consider  $\Lambda$  ( $\Lambda \leq K$ ) different cache states, forcing  $B = \frac{K}{\Lambda}$  users to have the same cache content (as elaborated in [20]). Again as elaborated in [20], during the subsequent delivery phase, we consider  $G \triangleq \Lambda\gamma + 1$  user groups (each group consisting of users with the same cache contents) that are selected for service in each transmission round, and from each group, we select  $Q \in [B]$  users that will receive information during that same round. This indicates that there are as many as  $GQ$  users served at a time.

We use  $\Psi$  to denote the  $G$  groups selected for service during a specific transmission round. We also use  $\mathbf{U}_{\psi, k}$  to denote the  $k$ -th ( $k \in [Q]$ ) active user in group  $\psi \in \Psi$ . We further assume that each served user  $\mathbf{U}_{\psi, k}$  is equipped with

<sup>2</sup>This ZF precoding corresponds to the work of the conference version [1], which though constrains itself to a special case of an equal number of receiving antennas per user. Furthermore, the numerical comparisons with the optimized BD-MRC were also omitted in [1].

TABLE I: Notations of Important Variables

| Notation                   | Definition   |
|----------------------------|--|
| $K$                        | Total number of users  |
| $L$                        | Number of transmit antennas  |
| $P_{\text{tot}}$           | Maximum allowable transmit power   |
| $F$                        | Total information bits per file  |
| $\mathcal{F}$              | Library with each file of $F$ bits   |
| $N$                        | Total number of files in the library $\mathcal{F}$   |
| $W_n$                      | The $n$ -th file of library $\mathcal{F}$  |
| $W_n^{\mathcal{T}}$        | Subfile of $W_n$ labeled by $\Lambda\gamma$ -tuple $\mathcal{T} \subseteq [\Lambda]$   |
| $\gamma$                   | Normalized cache size relative to the library content  |
| $\Lambda$                  | Number of distinct cache states  |
| $Q$                        | Number of users for signal decoding in a user-group of VCC   |
| $Q'$                       | Number of users served at a time in cacheless MU-MIMO  |
| $G$                        | $= \Lambda\gamma + 1$ , nominal gain of coded caching  |
| $\mathcal{G}$              | Effective gain over cacheless MU-MIMO  |
| $\mathcal{G}^*$            | Effective gain with $Q'$ and $Q$ being independently optimized   |
| $\Psi$                     | Set of $G$ user groups, each group with a distinct cache state   |
| $B$                        | Number of users caching the same content   |
| $\beta_{\psi, k}$          | Large-scale fading and/or pathloss to user $\mathbf{U}_{\psi, k}$  |
| $d_{\psi, k}$              | File index required by user $\mathbf{U}_{\psi, k}$   |
| $M_{\psi, k}$              | Number of receive antennas at user $\mathbf{U}_{\psi, k}$  |
| $J_{\psi, k}$              | Number of signal symbols simultaneously sent to user $\mathbf{U}_{\psi, k}$  |
| $M_{\psi}$                 | Total number of receive antennas in user-group $\psi$  |
| $J_{\psi}$                 | Total number of signal symbols sent to user-group $\psi$   |
| $s_{\psi, k, q}$           | $q$ -th data symbol sent to user $\mathbf{U}_{\psi, k}$  |
| $\mathbf{s}_{\psi, k}$     | Data vector sent to user $\mathbf{U}_{\psi, k}$  |
| $\mathbf{P}_{\psi, k}$     | Diagonal matrix for power allocation to user $\mathbf{U}_{\psi, k}$  |
| $P_{\psi, k, q}$           | Transmit power allocated to the $q$ -th data symbol to user $\mathbf{U}_{\psi, k}$   |
| $\mathbf{V}_{\psi, k}$     | Precoding matrix for user $\mathbf{U}_{\psi, k}$   |
| $\mathbf{v}_{\psi, k, q}$  | $q$ -th column of $\mathbf{V}_{\psi, k}$   |
| $\mathbf{s}_{\psi}$        | Data vector sent to user-group $\psi$  |
| $\mathbf{V}_{\psi}$        | Precoding matrix for user-group $\psi$   |
| $\mathbf{H}_{\psi, k}$     | Channel matrix from transmitter to user $\mathbf{U}_{\psi, k}$   |
| $\mathbf{H}_{\psi}$        | $\triangleq [\mathbf{H}_{\psi, 1}, \dots, \mathbf{H}_{\psi, Q}]$   |
| $\mathbf{H}_{\psi, -k}$    | $\triangleq [\mathbf{H}_{\psi, 1}, \dots, \mathbf{H}_{\psi, k-1}, \mathbf{H}_{\psi, k+1}, \dots, \mathbf{H}_{\psi, Q}]$                  |
| $\mathbf{T}_{\psi, -k}$    | Projection matrix of the null-space of $\mathbf{H}_{\psi, -k}^*$   |
| $\mathbf{R}_{\psi, k}$     | Receiver-side matrix for signal combining at user $\mathbf{U}_{\psi, k}$   |
| $\mathbf{r}_{\psi, k, q}$  | $q$ -th column of $\mathbf{R}_{\psi, k}$   |
| $\mathbf{z}_{\psi, k}$     | AWGN at user $\mathbf{U}_{\psi, k}$ and $\mathbf{z}_{\psi, k} \sim \mathcal{CN}(\mathbf{0}_{M_{\psi, k}}, N_0 \mathbf{I}_{M_{\psi, k}})$ |
| $\xi_{G, Q}$               | CSI costs of simultaneously serving $GQ$ users in VCC  |
| $\lambda_{\psi, k, q}$     | $q$ -th largest (non-zero) eigenvalue of $\mathbf{H}_{\psi, k}^T \mathbf{T}_{\psi, -k} \mathbf{H}_{\psi, k}^*$                           |
| $\mathbf{t}_{\psi, k, q}$  | Eigenvector associated with the eigenvalue $\lambda_{\psi, k, q}$  |
| $\lambda_{\psi, k}^{\min}$ | $\triangleq \min_{q \in [J_{\psi, k}]} \{\lambda_{\psi, k, q}\}$   |
| $\lambda_{\psi, k}^{\max}$ | $\triangleq \max_{q \in [J_{\psi, k}]} \{\lambda_{\psi, k, q}\}$   |

$M_{\psi, k}$  receive antennas. We use  $M_{\psi} \triangleq \sum_{k \in [Q]} M_{\psi, k}$  to denote the total number of active receive antennas in group  $\psi \in \Psi$ . As a result of having multi-antenna receivers, the BS can simultaneously communicate a certain number  $J_{\psi, k}$  of symbols to  $\mathbf{U}_{\psi, k}$ , thereby enhancing the throughput to  $\mathbf{U}_{\psi, k}$ , where the maximum achievable value of  $J_{\psi, k}$  will naturally depend on  $L$  and  $M_{\psi, k}$ , as well as will depend on the adopted precoding scheme. We further use

$$\mathbf{s}_{\psi, k} \triangleq [s_{\psi, k, 1}, \dots, s_{\psi, k, J_{\psi, k}}]^T \in \mathbb{C}^{J_{\psi, k}},$$

$$\mathbf{P}_{\psi, k} \triangleq \text{Diag}\left\{\sqrt{P_{\psi, k, 1}}, \dots, \sqrt{P_{\psi, k, J_{\psi, k}}}\right\} \in \mathbb{C}^{J_{\psi, k} \times J_{\psi, k}},$$

to denote the data vector to  $\mathbf{U}_{\psi, k}$  and the corresponding power allocation matrix respectively, where the independent variables  $\{s_{\psi, k, q} : q \in [J_{\psi, k}]\}$  have zero-mean and unit-power. As we can now see,  $J_{\psi} \triangleq \sum_{k \in [Q]} J_{\psi, k}$  corresponds to the total

$$\mathbf{y}_{\psi,k} = \mathbf{H}_{\psi,k}^T \mathbf{V}_{\psi,k} \mathbf{P}_{\psi,k} \mathbf{s}_{\psi,k} + \underbrace{\mathbf{H}_{\psi,k}^T \sum_{k' \in [Q] \setminus k} \mathbf{V}_{\psi,k'} \mathbf{P}_{\psi,k'} \mathbf{s}_{\psi,k'}}_{\text{intra-group interference}} + \underbrace{\mathbf{H}_{\psi,k}^T \sum_{\phi \in \Psi \setminus \psi} \sum_{\vartheta \in [Q]} \mathbf{V}_{\phi,\vartheta} \mathbf{P}_{\phi,\vartheta} \mathbf{s}_{\phi,\vartheta}}_{\text{inter-group interference}} + \mathbf{z}_{\psi,k} \quad (2)$$

number of data symbols simultaneously sent to group  $\psi$ . We also note that each signal vector  $\mathbf{s}_{\psi,k}$  destined for  $\mathbf{U}_{\psi,k}$  will be pre-processed into  $\mathbf{V}_{\psi,k} \mathbf{P}_{\psi,k} \mathbf{s}_{\psi,k}$  by a precoding matrix

$$\mathbf{V}_{\psi,k} \triangleq [\mathbf{v}_{\psi,k,1}, \dots, \mathbf{v}_{\psi,k,J_{\psi,k}}] \in \mathbb{C}^{L \times J_{\psi,k}},$$

where each unit-norm vector  $\mathbf{v}_{\psi,k,q} \in \mathbb{C}^L$  precodes  $s_{\psi,k,q}$ .

Let us define

$$\begin{aligned} \mathbf{s}_{\psi} &\triangleq [\mathbf{s}_{\psi,1}^T, \dots, \mathbf{s}_{\psi,Q}^T]^T \in \mathbb{C}^{J_{\psi}}, \\ \mathbf{V}_{\psi} &\triangleq [\mathbf{V}_{\psi,1}, \dots, \mathbf{V}_{\psi,Q}] \in \mathbb{C}^{L \times J_{\psi}}, \end{aligned}$$

where  $\mathbf{s}_{\psi}$  denotes the vector of data destined for the  $Q$  served users in group  $\psi$ , and where  $\mathbf{V}_{\psi}$  represents the corresponding precoding scheme. The transmitted signal  $\mathbf{x}_{\psi} \in \mathbb{C}^L$  for the groups in  $\Psi$  then takes the form

$$\mathbf{x}_{\psi} = \sum_{\psi \in \Psi} \sum_{k \in [Q]} \mathbf{V}_{\psi,k} \mathbf{P}_{\psi,k} \mathbf{s}_{\psi,k} = \sum_{\psi \in \Psi} \mathbf{V}_{\psi} \mathbf{P}_{\psi} \mathbf{s}_{\psi} \quad (1)$$

where the diagonal matrix  $\mathbf{P}_{\psi} \in \mathbb{C}^{J_{\psi} \times J_{\psi}}$  contains the ordered diagonal elements of matrices  $\{\mathbf{P}_{\psi,k} : k \in [Q]\}$ , which is responsible for allocating power to the symbols in  $\mathbf{s}_{\psi}$ .

*Remark 1:* This so-called VCC approach, simply ‘collapses’ (by linearly combining) a carefully selected set of  $|\Psi| = G$  vectors into the single vector in (1) for one-shot transmission. Without caching, this would require  $G$  sequential transmissions.

Given  $\mathbf{x}_{\psi}$  in (1), the received signal vector  $\mathbf{y}_{\psi,k} \in \mathbb{C}^{M_{\psi,k}}$  at user  $\mathbf{U}_{\psi,k}$  takes the form of (2), shown at the top of this page, where  $\mathbf{z}_{\psi,k} \sim \mathcal{CN}(\mathbf{0}_{M_{\psi,k}}, N_0 \mathbf{I}_{M_{\psi,k}})$  denotes the Additive White Gaussian Noise (AWGN), and where  $\mathbf{H}_{\psi,k} \in \mathbb{C}^{L \times M_{\psi,k}}$  denotes the channel matrix from the BS to  $\mathbf{U}_{\psi,k}$ . As each cache-aided user  $\mathbf{U}_{\psi,k}$  knows — as we detail in Section III-A — the messages  $\{\mathbf{s}_{\phi,\vartheta} : \phi \in \Psi \setminus \psi, \vartheta \in [Q]\}$  intended by the active users in other cache states in  $\Psi$ , we can conclude that the inter-group interference in (2) can be removed by using the cached content in  $\mathbf{U}_{\psi,k}$  and the composite CSI  $\{\mathbf{H}_{\psi,k}^T \mathbf{V}_{\phi,\vartheta} \mathbf{P}_{\phi,\vartheta} : \phi \in \Psi \setminus \psi, \vartheta \in [Q]\}$ , the cost of which we will account for in our analysis.

Decoding at each user  $\mathbf{U}_{\psi,k}$  will involve the removal of the known cache-aided interference, and the subsequent processing of the remaining received signal by a receiver-side matrix  $\mathbf{R}_{\psi,k} \triangleq [\mathbf{r}_{\psi,k,1}, \dots, \mathbf{r}_{\psi,k,J_{\psi,k}}] \in \mathbb{C}^{M_{\psi,k} \times J_{\psi,k}}$ , which is channel-dependent and whose columns are power-normalized to 1. Thus, in the end, after employing cache content to remove inter-group interference, and after applying the decoding matrix  $\mathbf{R}_{\psi,k}$ , the signal vector for decoding at  $\mathbf{U}_{\psi,k}$  takes the form

$$\begin{aligned} \mathbf{y}'_{\psi,k} &= \mathbf{R}_{\psi,k}^H \mathbf{H}_{\psi,k}^T \mathbf{V}_{\psi,k} \mathbf{P}_{\psi,k} \mathbf{s}_{\psi,k} + \mathbf{z}'_{\psi,k} \\ &+ \mathbf{R}_{\psi,k}^H \sum_{k' \in [Q] \setminus k} \mathbf{H}_{\psi,k'}^T \mathbf{V}_{\psi,k'} \mathbf{P}_{\psi,k'} \mathbf{s}_{\psi,k'}. \end{aligned} \quad (3)$$

It is easy to see that  $\mathbf{z}'_{\psi,k} \triangleq \mathbf{R}_{\psi,k}^H \mathbf{z}_{\psi,k} \in \mathbb{C}^{J_{\psi,k}}$  still follows a multi-variate complex Gaussian distribution.

### III. BD-MRC BASED VECTOR CODED CACHING

In this section, we will first elaborate on VCC for multi-antenna receivers, and then proceed to design the coherent BD-MRC scheme. Finally, we will present the main performance metrics for analysis in this paper.

#### A. Signal-Level Vector Coded Caching for Finite SNR

We will build on the general vector-clique structure in [20], which effectively super-imposes  $G$  differently precoded data vectors. We will here though carefully choose the precoding schemes, while also carefully calibrating the dimensionality of our super-imposed vectors by deciding how many users to serve from or groups. This will allow us to control CSI costs and to efficiently allocate power across users efficiently, thus controlling two aspects that crucially affect the performance in practical SNR regimes [25].

We proceed to describe the cache placement phase and the subsequent delivery phase. For clarity, we summarize the main VCC design in **Algorithm 1**.

1) *Placement Phase:* The cache placement policy is from [20]. The first step involves the partition of each library file  $W_n$  into  $\binom{\Lambda}{\Lambda\gamma}$  non-overlapping equally-sized subfiles  $\{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [\Lambda], |\mathcal{T}| = \Lambda\gamma\}$ , each labeled by some  $\Lambda\gamma$ -tuple  $\mathcal{T} \subseteq [\Lambda]$ . As discussed in Section I, the number of cache states  $\Lambda$  is chosen to satisfy the file-size constraint. Subsequently the  $K$  users are *arbitrarily* separated into  $\Lambda$  disjoint groups  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{\Lambda}$ , where the  $g$ -th user-group  $\mathcal{D}_g \triangleq \{b'\Lambda + g\}_{b'=0}^{B-1} \subseteq [K]$  consists of  $B = \frac{K}{\Lambda}$  users<sup>3</sup>. All the users belonging to the same group are assigned the same cache state and thus proceed to cache *identical* content. In particular, for those in the  $g$ -th group, this content takes the form  $\mathcal{Z}_{\mathcal{D}_g} = \{W_n^{\mathcal{T}} : \mathcal{T} \ni g, \forall n \in [N]\}$ . We clarify that this grouping as well as the entire placement phase (which takes place rarely, and over off-peak hours), are naturally done before the users’ requests take place.

2) *Delivery Phase:* This phase starts when each user  $\mathbf{U}_{\psi,\kappa}$ ,  $\psi \in \Psi$  and  $\kappa \in [B]$ , simultaneously asks for its intended file, denoted here by  $W_{d_{\psi,\kappa}}$ ,  $d_{\psi,\kappa} \in [N]$ . The BS then selects  $Q \leq B$  users from each group<sup>4</sup>. By doing so, the BS decides to first ‘encode’ over the first  $\Lambda Q$  users, and to repeat the encoding process  $B/Q$  times<sup>5</sup>. To deliver to the  $\Lambda Q$  users, the BS employs  $\binom{\Lambda}{\Lambda\gamma+1}$  sequential transmission stages. During

<sup>3</sup>We will henceforth consider  $K$  to be a multiple of  $\Lambda$  for the sake of clarity of exposition, without though limiting the scope of the results. The general case can be readily handled (cf. [20]) with a slight performance loss (cf. [25]).

<sup>4</sup>It is not difficult to see that in the absence of caching, corresponding to  $G = 1$ , this  $Q$  would have played the role of the multiplexing gain.

<sup>5</sup>To see this, let us consider the example in Fig. 1, where  $G = 2$  groups (2 cache states), each with  $Q = 2$  active users, are selected for service in each transmission stage. The algorithm that we describe here will be first applied to the first  $\Lambda Q$  users, and then, after this delivery is done, the same algorithm will apply to the remaining  $\Lambda Q$  users, thus eventually satisfying all  $K$  users. Also note that a small amount of additional subpacketization can easily resolve the case where  $B/Q$  may not be an integer [20].

**Algorithm 1:** Vector Coded Caching (VCC)

**1 Procedure Placement Phase**

- 2 Partition each file  $W_n$  into  $\binom{\Lambda}{\Lambda\gamma}$  non-overlapping subfiles:  $W_n \rightarrow \{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [\Lambda], |\mathcal{T}| = \Lambda\gamma\}$
- 3 Divide  $K$  users into  $\Lambda$  groups, each with  $B = K/\Lambda$  users, indexed by  $[\Lambda] \triangleq \{1, 2, \dots, \Lambda\}$ .
- 4 All users in the  $g$ -th ( $g \in [\Lambda]$ ) group store identical content:  $\mathcal{Z}_{\mathcal{D}_g} = \{W_n^{\mathcal{T}} : g \in \mathcal{T}, \forall n \in [N]\}$
- 5 **End Procedure**

**6 Procedure Delivery Phase**

- 7 Each user requests a distinct file
- 8 **for**  $b = 1$  **to**  $B/Q$  **do**
- 9     Each encoding process serves a *unique* selection of  $Q \leq B$  users from each group
- 10    **for**  $stage t = 1$  **to**  $\binom{\Lambda}{\Lambda\gamma+1}$  **do**
- 11      In each stage, choose a *unique* subset  $\Psi \subseteq [\Lambda]$  with  $|\Psi| = \Lambda\gamma + 1$
- 12      Denote the requested file of user  $U_{\psi,k}$  by  $W_{d_{\psi,k}}$ , for  $\psi \in \Psi$  and  $k \in [Q]$
- 13      Transmit  $\{W_{d_{\psi,k}}^{\Psi \setminus \psi} : \psi \in \Psi, k \in [Q]\}$
- 14      Design transmit signal  $\mathbf{x}_{\Psi}$  in (1) by mapping  $W_{d_{\psi,k}}^{\Psi \setminus \psi} \rightarrow \mathbf{s}_{\psi,k}$
- 15      For received signal  $\mathbf{y}_{\psi,k}$  in (2), each user  $U_{\psi,k}$  cancels inter-group interference via cached content:  $\{W_n^{\Psi \setminus \phi} : \phi \in \Psi \setminus \psi, \forall n \in [N]\}$
- 16      Intra-group interference in (2) is mitigated via precoding and receiver processing (see (3))
- 17     Selected  $Q$  users per group obtain requested files
- 18 **End Procedure**

each such stage, the BS simultaneously serves a unique set  $\Psi$  of  $|\Psi| = \Lambda\gamma + 1$  groups, corresponding to a total of  $Q(\Lambda\gamma + 1)$  users served at a time (i.e., per transmission stage), during which all the subfiles  $\{W_{d_{\psi,k}}^{\Psi \setminus \psi} : \psi \in \Psi, k \in [Q]\}$  are sent simultaneously. After mapping each such subfile  $W_{d_{\psi,k}}^{\Psi \setminus \psi}$  into a complex-valued data vector  $\mathbf{s}_{\psi,k}$ , the transmission is as shown in (1). We here note that  $U_{\psi,k}$  has cached the subfiles  $\{W_n^{\Psi \setminus \phi} : \phi \in \Psi \setminus \psi, \forall n \in [N]\}$ , thereby enabling  $U_{\psi,k}$  to cancel the inter-group interference resulting from  $\{W_{d_{\phi,k}}^{\Psi \setminus \phi} : \phi \in \Psi \setminus \psi, \forall \phi \in \Psi \setminus \psi, k \in [Q]\}$ . Note that the inter-group interference cancellation also requires knowledge of the composite CSI (cf. (2)). On the other hand, the intra-group interference experienced at each user  $U_{\psi,k}$ , resulting from  $\{W_{d_{\psi,k'}}^{\Psi \setminus \psi} : k' \in [Q] \setminus k\}$ , can be handled — as we will do here — with linear precoding that ‘separates’ the signals of the users from the same group. At the end of the  $\binom{\Lambda}{\Lambda\gamma+1}$  transmission stages, all the  $\Lambda Q$  users obtain their intended files. By repeating this process  $B/Q$  times per group, across all groups, all the  $K$  users obtain all their intended files. We refer to [20] for more details.

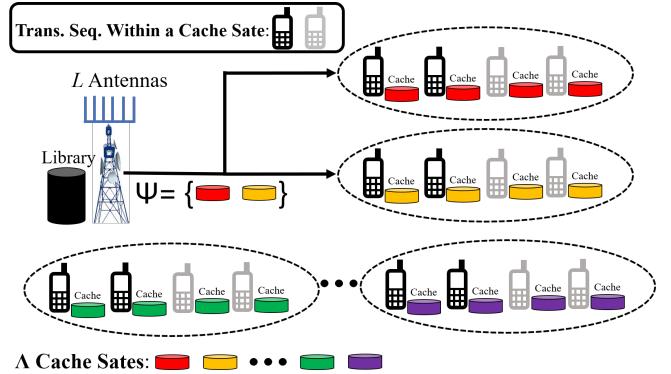


Fig. 1: Illustration for vector coded caching with  $G = 2$ ,  $B = 4$  and  $Q = 2$ .

**B. BD Precoding and MRC Combining**

As pointed out in [37], complete channel diagonalization (e.g., ZF) at the BS is suboptimal since each multi-antenna user is able to coordinate the processing of its own receiver outputs. We thus alternatively consider the well-known BD precoding method [37]. Toward this, we first define the matrix  $\mathbf{H}_{\psi,-k} \triangleq [\mathbf{H}_{\psi,1}, \dots, \mathbf{H}_{\psi,k-1}, \mathbf{H}_{\psi,k+1}, \dots, \mathbf{H}_{\psi,Q}] \in \mathbb{C}^{L \times (M_{\psi} - M_{\psi,k})}$ . To cancel the intra-group interference in (2),  $\mathbf{V}_{\psi,k}$  must lie in the null-space of  $\mathbf{H}_{\psi,-k}^*$  (denoted by  $\mathbb{O}\{\mathbf{H}_{\psi,-k}^*\}$ ) such that the product of  $\mathbf{H}_{\psi,-k}^T$  and  $\mathbf{V}_{\psi,k}$  is the zero matrix for any  $k \in [Q]$ . Thus, to successfully eliminate the inter-stream interference in  $U_{\psi,k}$ , we have that

$$J_{\psi,k} \leq \min \left\{ \text{Rank}(\mathbb{O}\{\mathbf{H}_{\psi,-k}^*\}), \text{Rank}(\mathbf{H}_{\psi,k}) \right\}, \quad (4)$$

which, for independent Rayleigh fading channels, implies that  $J_{\psi,k} \leq \min \{L - (M_{\psi} - M_{\psi,k}), M_{\psi,k}\}$ . We use  $Q_{\psi,\max}$  to denote the maximum  $Q$  in the user-group  $\psi$ , which equals  $\min \{\lfloor \frac{M+L-1}{M} \rfloor, B\}$  if all the served users have the same number of receive antennas  $M$ . As the users are equipped with different numbers of antennas, we determine  $Q_{\psi,\max}$  by solving  $\frac{\sum_{k \in [Q]} M_{\psi,k}}{1-L+\sum_{k \in [Q]} M_{\psi,k}} - Q = 0$ . As we here consider the simplified case of a uniform  $Q$  for each group, we accept the aforementioned maximum allowable  $Q$  that takes the form  $Q_{\max} = \min_{\psi \in \Psi} Q_{\psi,\max}$ .

Let  $\mathbf{T}_{\psi,-k} \triangleq \mathbf{I}_L - \mathbf{H}_{\psi,-k}^* (\mathbf{H}_{\psi,-k}^* \mathbf{H}_{\psi,-k})^{-1} \mathbf{H}_{\psi,-k}^T$  be the projection matrix which maps any vector  $\mathbf{m} \in \mathbb{C}^L$  into the null-space of  $\mathbf{H}_{\psi,-k}^*$ . We note that  $\mathbf{T}_{\psi,-k}^2 = \mathbf{T}_{\psi,-k} = \mathbf{T}_{\psi,-k}^H$  according to the idempotent and Hermitian properties of a projection matrix. The BD precoding matrix  $\mathbf{V}_{\psi,k} \in \mathbb{C}^{L \times J_{\psi,k}}$  dedicated to  $U_{\psi,k}$  can be written as

$$\mathbf{V}_{\psi,k} = \left[ \frac{\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,1}}{\|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,1}\|}, \dots, \frac{\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,J_{\psi,k}}}{\|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,J_{\psi,k}}\|} \right] \quad (5)$$

where the  $q$ -th ( $q \in [J_{\psi,k}]$ ) column (denoted by  $\mathbf{v}_{\psi,k,q}$ ) of  $\mathbf{V}_{\psi,k}$  precodes  $s_{\psi,k,q}$ . By using this BD precoding, the signal vector at  $U_{\psi,k}$  in (3) becomes

$$\mathbf{y}'_{\psi,k} = \mathbf{R}_{\psi,k}^H \mathbf{H}_{\psi,k}^T \mathbf{V}_{\psi,k} \mathbf{P}_{\psi,k} \mathbf{s}_{\psi,k} + \mathbf{z}'_{\psi,k}, \quad (6)$$

and its  $q$ -th element  $y'_{\psi,k,q}$  of  $\mathbf{y}'_{\psi,k}$  (i.e., which corresponds to decoding  $s_{\psi,k,q}$ ) is of the form

$$\begin{aligned} y'_{\psi,k,q} &= \mathbf{r}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q} \sqrt{P_{\psi,k,q}} s_{\psi,k,q} + z'_{\psi,k,q} \\ &+ \underbrace{\sum_{p \in [J_{\psi,k}] \setminus q} \mathbf{r}_{\psi,k,p}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,p} \sqrt{P_{\psi,k,p}} s_{\psi,k,p}}_{\text{inter-stream interference}}, \end{aligned} \quad (7)$$

where this inter-stream interference results from the other symbols intended for the same user. To maximize the power of the useful signal for decoding  $s_{\psi,k,q}$ , we have that  $|\mathbf{r}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q}|^2 \leq \|\mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q}\|^2$ , where the equality is achieved only if

$$\mathbf{r}_{\psi,k,q} = \theta \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q} = \theta \mathbf{H}_{\psi,k}^T \frac{\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}}{\|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}\|} \quad (8)$$

for a non-zero constant  $\theta$  according to the well-known Cauchy–Schwarz inequality. Here, we set  $\theta = 1/\|\mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q}\|$  to normalize  $\mathbf{r}_{\psi,k,q}$ ; this corresponds to the MRC receiver.

In the following, we will show how to select the vectors  $\mathbf{m}_{\psi,k,1}, \dots, \mathbf{m}_{\psi,k,J_{\psi,k}}$  in (5) to remove the inter-stream interference in (7). When  $\mathbf{m}_{\psi,k,1}, \dots, \mathbf{m}_{\psi,k,J_{\psi,k}}$  are the *orthogonal* eigenvectors of  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$ , the inter-stream interference in  $y'_{\psi,k,q}$  is completely eliminated, which can be easily observed by the fact that for  $p \neq q$ , we have that

$$\begin{aligned} \mathbf{r}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,p} &= \frac{\theta \mathbf{m}_{\psi,k,q}^H \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,p}}{\|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}\| \|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,p}\|} \\ &= \frac{\theta \lambda_{\psi,k,p} \mathbf{m}_{\psi,k,p}^H \mathbf{m}_{\psi,k,p}}{\|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}\| \|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,p}\|} = 0 \end{aligned} \quad (9)$$

where  $\lambda_{\psi,k,p}$  denotes the eigenvalue associated with the eigenvector  $\mathbf{m}_{\psi,k,p}$  of  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$ . Therefore, under the common Gaussian signaling and the coherent BD-MRC scheme, the signal-to-interference plus noise ratio (SINR) for decoding  $s_{\psi,k,q}$  at  $\mathbf{U}_{\psi,k}$  is of the form

$$\begin{aligned} \text{SINR}_{\psi,k,q}^{\text{BD-MRC}} &= \frac{P_{\psi,k,q} |\mathbf{r}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,q}|^2}{N_0 + \sum_{p \in [J_{\psi,k}] \setminus q} P_{\psi,k,p} |\mathbf{r}_{\psi,k,p}^H \mathbf{H}_{\psi,k}^T \mathbf{v}_{\psi,k,p}|^2} \end{aligned} \quad (10)$$

$$= \frac{P_{\psi,k,q} \mathbf{m}_{\psi,k,q}^H \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}}{N_0 \|\mathbf{T}_{\psi,-k} \mathbf{m}_{\psi,k,q}\|^2} \quad (11)$$

where we note that (10) is based on the received signal in (7), while (11) follows from (8) and (9).

Toward providing CSI estimates to the BS and the users, we will consider the common TDD uplink-downlink pilot transmission, as this applies to MU-MIMO systems. Thus for  $T$  being the coherence block period (in symbols), and for  $\Theta$  being the number of symbols per receiving antenna and per block used for pilot transmission, then the effective rate for  $\mathbf{U}_{\psi,k}$  is

$$R_{\psi,k} = \xi_{G,Q} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \text{SINR}_{\psi,k,q}^{\text{BD-MRC}} \right) \text{ nats/s/Hz} \quad (12)$$

where  $\xi_{G,Q} \triangleq 1 - \Theta(\sum_{\psi \in \Psi} \sum_{k \in [Q]} M_{\psi,k})/T$  accounts for CSI costs [38]. Without loss of generality, we sort  $\{s_{\psi,k,q} : q \in [J_{\psi,k}]\}$  in descending order according to the corresponding  $\{\text{SINR}_{\psi,k,q}^{\text{BD-MRC}} : q \in [J_{\psi,k}]\}$  at  $\mathbf{U}_{\psi,k}$ .

### C. Some Preliminary Results for BD-MRC

As we see, implementing BD-MRC requires us to eigen-decompose  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$  whose large size (comprising  $L \times L$  elements) renders it computationally consuming, especially in the massive MIMO regime. This motivates us to design a low-complexity approach to considerably ease the BD-MRC implementation in Lemma 1.

This design is motivated by the observation (which can be shown by using basic properties of projection matrices and basic algebraic manipulations) that the matrices  $\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*$  and  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$  share the same non-zero eigenvalues. Before presenting Lemma 1, we define the vectors  $\{\mathbf{t}_{\psi,k,q} \in \mathbb{C}^{M_{\psi,k}} : q \in [J_{\psi,k}]\}$  as the orthogonal eigenvectors of  $\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \in \mathbb{C}^{M_{\psi,k} \times M_{\psi,k}}$ , where the eigenvector  $\mathbf{t}_{\psi,k,q}$  is associated with the  $q$ -th largest (non-zero) eigenvalue  $\lambda_{\psi,k,q}$ .

**Lemma 1:** The precoding vector  $\mathbf{v}_{\psi,k,q}$  to  $s_{\psi,k,q}$  for  $\mathbf{U}_{\psi,k}$  under the BD-MRC scheme is given the form

$$\mathbf{v}_{\psi,k,q} = \frac{\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}}{\|\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}\|}, \quad (13)$$

and this yields a corresponding SINR for symbol  $s_{\psi,k,q}$  of the form

$$\text{SINR}_{\psi,k,q}^{\text{BD-MRC}} = \frac{P_{\psi,k,q}}{N_0} \lambda_{\psi,k,q}, \quad (14)$$

and an effective rate for  $\mathbf{U}_{\psi,k}$  of the form

$$R_{\psi,k} = \xi_{G,Q} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}}{N_0} \lambda_{\psi,k,q} \right). \quad (15)$$

*Proof:* The proof is relegated to Appendix I. ■

**Remark 2:** We can see that each user  $\mathbf{U}_{\psi,k}$  can remove interference to decode up to  $J_{\psi,k} \leq \text{Rank}(\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*) \leq M_{\psi,k}$  of its symbols. We can also see that under the constraint of a total power  $P_{\psi,k} = \sum_{q=1}^{J_{\psi,k}} P_{\psi,k,q}$  allocated to  $\mathbf{U}_{\psi,k}$ , water-filling (cf. [43, Ch. 10]) can maximize the effective rate of that user.

**Remark 3:** We note that the design in Lemma 1 involves an eigendecomposition of a much smaller  $M_{\psi,k} \times M_{\psi,k}$  matrix with a computational complexity of  $\mathcal{O}(M_{\psi,k}^3)$ . Assuming each user has  $M$  receive antennas and  $J_{\psi,k} = M$ , the complexity of (13) is  $\mathcal{O}(M^3 + L^2 + ML)$ , while that of  $\mathbf{V}_{\psi,k}$  is  $\mathcal{O}(M^3 + ML^2 + M^2L)$ . In contrast, eigendecomposing  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$  requires  $\mathcal{O}(L^3)$ , and the cost of (5) is  $\mathcal{O}(L^3 + ML^2)$ . Since  $L \gg M$  in practice, the overall scaling is asymptotically  $\mathcal{O}(L^2)$  for Lemma 1 and  $\mathcal{O}(L^3)$  for (5), yielding an order- $L$  reduction in computational complexity.

Let us now address the balance between multiplexing and beamforming gains, by first noting that such optimal balancing requires either real-time adjustment of the number of transmitted symbols ( $J_{\psi,k}$ ), or, when we always use the maximum allowable number of transmitted symbols for  $\mathbf{U}_{\psi,k}$ , requires water-filling to dynamically allocate different powers to our symbols. We will here consider the maximum allowable value of  $J_{\psi,k}$  transmitted symbols to  $\mathbf{U}_{\psi,k}$ , i.e., we will consider

$$J_{\psi,k} = \text{Rank}(\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*) \leq M_{\psi,k}. \quad (16)$$

The following corollary characterizes the optimal power allocation for  $U_{\psi,k}$  by leveraging the water-filling algorithm, which distributes power across the transmitted symbols based on instantaneous channel conditions, ensuring adherence to the power constraint  $P_{\psi,k}$  and maximizing the effective rate in (15).

*Corollary 1:* The optimal power allocated to symbol  $s_{\psi,k,q}$  for maximizing the effective rate in (15), takes the form

$$P_{\psi,k,q} = \left( \frac{1}{\alpha_{\psi,k}} - \frac{N_0}{\lambda_{\psi,k,q}} \right)^+, \quad (17)$$

where, under a power constraint  $P_{\psi,k}$ , the Lagrange multiplier  $\alpha_{\psi,k}$  is the solution to

$$\sum_{q=1}^{J_{\psi,k}} \left( \frac{1}{\alpha_{\psi,k}} - \frac{N_0}{\lambda_{\psi,k,q}} \right)^+ = P_{\psi,k}. \quad (18)$$

Then, the optimal effective rate for  $U_{\psi,k}$  takes the form

$$R_{\psi,k}^*(P_{\psi,k}) = \xi_{G,Q} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \left( \frac{\lambda_{\psi,k,q}}{N_0 \alpha_{\psi,k}} - 1 \right)^+ \right). \quad (19)$$

*Proof:* The proof is direct from the water-filling algorithm. ■

#### D. MMF Framework and Main Performance Metrics

Let us now further consider MMF, where the minimum effective rate among the simultaneously served users is maximized via power allocation under a specific precoding scheme. The MMF problem for serving the users  $\{U_{\psi,k} : \psi \in \Psi, k \in [Q]\}$  can be now formulated as

$$\mathcal{S}_1 \left\{ \begin{array}{l} \max_{\mathcal{P}_{\Psi}} \min_{\psi \in \Psi} \min_{k \in [Q]} R_{\psi,k} \\ \text{s.t. } P_t = \sum_{\psi \in \Psi} \sum_{k \in [Q]} P_{\psi,k} \leq P_{\text{tot}} \end{array} \right. \quad (20)$$

where  $P_{\text{tot}}$  denotes the maximum allowable transmit power at the BS, and  $\mathcal{P}_{\Psi} \triangleq \{P_{\psi,k} : \psi \in \Psi, k \in [Q]\}$  denotes the set of used powers.

*Remark 4:* The water-filling algorithm in Corollary 1 is applied only across the streams of a given user  $U_{\psi,k}$  to maximize that user's effective rate under its allocated power  $P_{\psi,k}$ , i.e., the total power used to transmit the signal streams to  $U_{\psi,k}$ . User fairness is ensured separately by the MMF formulation in (20), which maximizes the minimum effective rate among the concurrently served  $GQ$  users in VCC.

Let us now formally define some important metrics of interest, which will be analyzed in Sections IV and V.

*Definition 1: (Effective sum-rate).* For a  $(G, Q)$ -VCC scheme, its effective (instantaneous) sum-rate is denoted by  $R(G, Q)$  and is defined as the total effective rate (after accounting for CSI costs) summed over the  $GQ$  simultaneously served users. Moreover,  $R^*(G, Q)$  denotes the effective sum-rate optimized under the MMF criterion (cf. (20)).

*Definition 2: (Effective gain over MU-MIMO).* For a given set of SNR,  $L$  and  $M_{\psi,k}$  ( $\forall \psi \in \Psi, k \in [Q]$ ) resources, and a fixed underlying precoder class, the effective gain, after accounting for CSI costs, of the  $(G, Q)$ -VCC scheme over the cacheless scenario (corresponding to  $G = 1$ , and an operating multiplexing gain  $Q'$ ), will be denoted as  $\mathcal{G} \triangleq \frac{\mathbb{E}_{h,r}\{R^*(G, Q)\}}{\mathbb{E}_{h,r}\{R^*(1, Q')\}}$ , where  $\mathbb{E}_{h,r}\{R^*(G, Q)\}$  describes the rate  $R^*(G, Q)$  averaged over channel fading and pathloss.

We also call  $\mathcal{G}^* \triangleq \frac{\max_Q \mathbb{E}_{h,r}\{R^*(G, Q)\}}{\max_{Q'} \mathbb{E}_{h,r}\{R^*(1, Q')\}}$  as the effective gain of optimized rates, where  $Q'$  and  $Q$  are also optimized, independently for the cacheless and the cache-aided scenario.

#### IV. PERFORMANCE ANALYSIS ON BD-MRC BASED VECTOR CODED CACHING

This section begins by analyzing and solving the MMF problem under the BD-MRC based VCC framework, which is reformulated into a one-dimensional linear search problem. In addition, we derive analytical upper and lower bounds for the optimal MMF solution. We then examine two special scenarios: (i) the case where each served user is sent the same number of symbols in each transmission round, and (ii) the massive MIMO scenario with Rayleigh fading channels.

##### A. Main Results

Let us first recall (cf. (9)) that the eigenvalue set  $\{\lambda_{\psi,k,q} : \psi \in \Psi, k \in [Q], q \in [J_{\psi,k}]\}$  is a function of the channel gains but not of the power allocation policy. Thus, by using the effective rate expression in (15), the MMF optimization problem in (20) under the BD-MRC scheme for downlink power allocation, can be transformed into

$$\mathcal{S}_2 \left\{ \begin{array}{l} \max_{\mathcal{P}_{\Psi}} \min_{\psi \in \Psi} \min_{k \in [Q]} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}}{N_0} \lambda_{\psi,k,q} \right) \\ \text{s.t. } P_t = \sum_{\psi \in \Psi} \sum_{k \in [Q]} \sum_{q \in [J_{\psi,k}]} P_{\psi,k,q} \leq P_{\text{tot}}. \end{array} \right. \quad (21)$$

Obviously,  $P_t$  should reach its upper-bound  $P_{\text{tot}}$  when the optimum is achieved. The following theorem solves the optimization problem in (21) to derive the optimal effective (instantaneous) sum-rate  $R_{\text{BD-MRC}}^*$  and the effective gain of optimized rates (cf. Definition 2). Before presenting the theorem, let us note that  $f_{\psi,k}^{-1}(\cdot)$  will denote the inverse function of  $R_{\psi,k}^*(P_{\psi,k})$  in (19), which is a monotonically increasing function w.r.t.  $P_{\psi,k}$ . We proceed with the theorem.

*Theorem 1:* The effective sum-rate  $R_{\text{BD-MRC}}^*$  under optimal power allocation for the MMF problem is the solution to

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} f_{\psi,k}^{-1} \left( \frac{R_{\text{BD-MRC}}^*}{GQ} \right) = P_{\text{tot}}, \quad (22)$$

and the optimal effective rate for each simultaneously served user is identical, given by  $R_{\text{BD-MRC}}^*/(GQ)$ . The corresponding effective gain of optimized rates under the BD-MRC scheme is given by

$$\mathcal{G}_{\text{BD-MRC}}^* = \frac{\max_{Q \in [Q_{\text{max}}]} \mathbb{E}_{h,r}\{R_{\text{BD-MRC}}^*(G, Q)\}}{\max_{Q' \in [Q'_{\text{max}}]} \mathbb{E}_{h,r}\{R_{\text{BD-MRC}}^*(1, Q')\}}. \quad (23)$$

Furthermore, this optimal sum-rate is bounded as

$$\tilde{R}_{\text{BD-MRC}}^* \leq R_{\text{BD-MRC}}^* \leq \hat{R}_{\text{BD-MRC}}^*$$

where  $\tilde{R}_{\text{BD-MRC}}^*$  and  $\hat{R}_{\text{BD-MRC}}^*$  are respectively the solutions to

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} \frac{J_{\psi,k} N_0}{\lambda_{\psi,k}^{\min}} \left( \exp \left( \frac{\tilde{R}_{\text{BD-MRC}}^*}{\xi_{G,Q} J_{\psi,k} GQ} \right) - 1 \right) = P_{\text{tot}}, \quad (24)$$

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} \frac{J_{\psi,k} N_0}{\lambda_{\psi,k}^{\max}} \left( \exp \left( \frac{\hat{R}_{\text{BD-MRC}}^*}{\xi_{G,Q} J_{\psi,k} GQ} \right) - 1 \right) = P_{\text{tot}}, \quad (25)$$

where  $\lambda_{\psi,k}^{\min} \triangleq \min_{q \in [J_{\psi,k}]} \{\lambda_{\psi,k,q}\}$  and  $\lambda_{\psi,k}^{\max} \triangleq \max_{q \in [J_{\psi,k}]} \{\lambda_{\psi,k,q}\}$ . *Proof:* As the power allocation among the symbols intended by  $U_{\psi,k}$  does not affect the power allocation to other served users, the effective rate for  $U_{\psi,k}$  must reach its optimal bound under the power constraint  $P_{\psi,k}^*$  (optimal power allocated to  $U_{\psi,k}$  for MMF), which has been solved in Corollary 1. As  $R_{\psi,k}^*(P_{\psi,k})$  in (19) is a monotonically increasing function w.r.t.  $P_{\psi,k}$ , we can conclude that the simultaneously served users must have the same effective rate (equaling  $R_{\text{BD-MRC}}^*/G/Q$ ) under the optimal power allocation in (21). To see this, simply consider a contradictory case where, when the optimum for (21) is achieved, if there exists one user whose effective rate is higher than the smallest rate, this user can “borrow” some power to the user with the smallest rate until their rates are the same, without affecting the rates for other users, which enhances the smallest effective rates, and which is contradictory to the optimal power allocation assumption. Considering  $\sum_{\psi \in \Psi} \sum_{k \in [Q]} P_{\psi,k}^* = P_{\text{tot}}$  and the inverse function of  $R_{\psi,k}^*(P_{\psi,k})$ , we can obtain (22).

To derive the lower-bound for  $R_{\psi,k}$ , considering (15), we first have that

$$R_{\psi,k} \geq \xi_{G,Q} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}}{N_0} \lambda_{\psi,k}^{\min} \right) \quad (26)$$

which is then substituted into (21) as the objective function. Using similar analysis as that leading to (22), we can easily obtain the identity for  $\tilde{R}_{\text{BD-MRC}}^*$ . The upper-bound  $\hat{R}_{\text{BD-MRC}}^*$  follows the same procedure as  $\tilde{R}_{\text{BD-MRC}}^*$ , with a difference being that upper bounding  $R_{\psi,k}$  now uses  $\lambda_{\psi,k}^{\max}$ . ■

As shown in Fig. 2, the numerical results under different system parameter settings exhibit an excellent match with the simulation results, which confirms the accuracy of Theorem 1.

*Remark 5:* We quickly ask the reader to note that the derived sum-rate bounds  $\tilde{R}_{\text{BD-MRC}}^*$  and  $\hat{R}_{\text{BD-MRC}}^*$  considerably facilitate the use of the simple one-dimensional binary search to numerically establish  $R_{\text{BD-MRC}}^*$  in (22).

### B. Special Case I: Symmetric Case

We also have the following, which considers the commonly assumed symmetric case where each scheduled user is sent an equal number of symbols.

*Corollary 2:* In the symmetric case where  $J_{\psi,k} = J$  for any  $\psi \in \Psi$  and  $k \in [Q]$ , the lower and upper bounds to the optimal sum-rate  $R_{\text{BD-MRC}}^*$  take the forms

$$\tilde{R}_{\text{BD-MRC}}^* \triangleq \xi_{G,Q} G Q J \ln \left( 1 + \frac{P_{\text{tot}}/(N_0 J)}{\sum_{\psi \in \Psi} \sum_{k \in [Q]} (\lambda_{\psi,k}^{\min})^{-1}} \right), \quad (27)$$

$$\hat{R}_{\text{BD-MRC}}^* \triangleq \xi_{G,Q} G Q J \ln \left( 1 + \frac{P_{\text{tot}}/(N_0 J)}{\sum_{\psi \in \Psi} \sum_{k \in [Q]} (\lambda_{\psi,k}^{\max})^{-1}} \right). \quad (28)$$

*Proof:* The derivation of  $\tilde{R}_{\text{BD-MRC}}^*$  and  $\hat{R}_{\text{BD-MRC}}^*$  is direct after setting  $J_{\psi,k} = J$ ,  $\forall \psi \in \Psi, k \in [Q]$  in Theorem 1. ■

*Remark 6:* We note that Lemma 1, Corollary 1, Theorem 1 and Corollary 2, are all valid for any propagation channel model, including Rayleigh fading, Rician- $K$  fading, etc. The same results also hold for scenarios that involve a non-full-rank

channel matrix product  $\mathbf{H}_{\psi,-k}^T \mathbf{H}_{\psi,-k}^*$  (e.g., Keyhole channels) for  $\psi \in \Psi$  and  $k \in [Q]$ , in which case we apply the pseudo-inverse of  $\mathbf{H}_{\psi,-k}^T \mathbf{H}_{\psi,-k}^*$  in the projection matrix  $\mathbf{T}_{\psi,-k}$ .

### C. Special Case II: Massive MIMO Over Rayleigh Channels

In the following, we consider the massive MIMO regime [44] entailing a very large  $L$ , generally corresponding to the case of  $L \gg M_{\psi} = \sum_{k \in [Q]} M_{\psi,k}$  for  $\forall \psi \in \Psi$ . This setting captures certain technological trends, as well as simplifies exposition, by simplifying analysis of the eigenvalues of  $\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \in \mathbb{C}^{M_{\psi,k} \times M_{\psi,k}}$ . We here consider independent Rayleigh fading channels where  $J_{\psi,k} = M_{\psi,k}$  and where the elements of  $\mathbf{H}_{\psi,k}$  follow the i.i.d. complex Gaussian distribution with zero-mean and variance  $\beta_{\psi,k}$ , where  $\beta_{\psi,k}$  accounts for the large-scale fading and/or pathloss. Lemma 2 distills the result of Theorem 1 to the massive MIMO case, and the reported results are naturally independent of instantaneous channel fading.

*Lemma 2:* In the massive MIMO regime, the effective instantaneous sum-rate  $R_{\text{BD-MRC}}^*$  under the BD-MRC scheme and optimal power allocation for MMF over independent Rayleigh fading channels, is the solution to

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} \frac{N_0 M_{\psi,k} \left( \exp \left( \frac{R_{\text{BD-MRC}}^*}{\xi_{G,Q} M_{\psi,k} G Q} \right) - 1 \right)}{\beta_{\psi,k} (L - M_{\psi} + M_{\psi,k})} = P_{\text{tot}} \quad (29)$$

with the corresponding power allocation policy that yields the asymptotically optimal  $R_{\text{BD-MRC}}^*$ , taking the form

$$P_{\psi,k,q} = \frac{N_0 M_{\psi,k} \left( \exp \left( \frac{R_{\text{BD-MRC}}^*}{\xi_{G,Q} M_{\psi,k} G Q} \right) - 1 \right)}{\beta_{\psi,k} (L - M_{\psi} + M_{\psi,k})}. \quad (30)$$

For the case when  $M_{\psi,k} = M$  for  $\forall \psi \in \Psi, k \in [Q]$ , then

$$\begin{aligned} R_{\text{BD-MRC}}^*(G, Q) \\ \simeq \xi_{G,Q} G Q M \ln \left( 1 + \frac{P_{\text{tot}}(L - (Q-1)M)}{N_0 M \sum_{\psi \in \Psi} \sum_{k \in [Q]} \beta_{\psi,k}^{-1}} \right). \end{aligned} \quad (31)$$

*Proof:* The proof is relegated to Appendix II. ■

The numerical results shown in Fig. 3, under various system parameter settings, closely match the simulation results, confirming the accuracy of Lemma 2.

## V. PERFORMANCE ANALYSIS ON SIMPLE ZF BASED VECTOR CODED CACHING

In this section, we analyze the effective sum-rate achieved by the cache-aided downlink schemes of Section III-A for the case of the ZF linear precoder over independent Rayleigh fading channels where  $J_{\psi,k} = M_{\psi,k}$  for  $\psi \in \Psi$  and  $k \in [Q]$ . In contrast to the coherent BD-MRC design where both the precoding/combining matrix optimization and power allocation should be adjusted according to the instantaneous channel state (recall Theorem 1), we here (cf. Theorem 2) simply perform channel matrix inversion without any precoder optimization, and we simply calibrate the power allocation as a function of the (large-scale) pathloss which of course changes much slower than fading does. The numerical results in Section VI show that the performance gap between the simpler ZF precoder

and the fully optimized BD-MRC scheme is negligible for the practical setting of  $L \gg M_\psi$  (cf. Fig. 2).

Following the conventional ZF precoding for single-antenna receivers (cf. [45]), we completely separate the transmitted  $M_\psi = \sum_{k \in [Q]} M_{\psi,k}$  symbol streams such that there is no inter-stream interference. Therefore, the  $M_{\psi,k}$  symbols simultaneously sent to  $\mathbf{U}_{\psi,k}$  are fully separated (using complete channel diagonalization at the BS), and user  $\mathbf{U}_{\psi,k}$  independently decodes the intended  $M_{\psi,k}$  symbols without inter-stream interference. This corresponds to a decoding matrix  $\mathbf{R}_{\psi,k} = \mathbf{I}_{M_{\psi,k}}$  at user  $\mathbf{U}_{\psi,k}$ . We use  $\mathbf{H}_\psi \triangleq [\mathbf{H}_{\psi,1}, \dots, \mathbf{H}_{\psi,Q}] \in \mathbb{C}^{L \times M_\psi}$  to represent the channel matrix between the BS and the  $Q$  active users in user-group  $\psi$ . Here, the precoding matrix for the multi-antenna receivers in user-group  $\psi$  in (1) under the ZF variant scheme is designed as

$$\mathbf{V}_\psi = \mathbf{H}_\psi^* (\mathbf{H}_\psi^T \mathbf{H}_\psi^*)^{-1} \circ \mathbf{D}_\psi \in \mathbb{C}^{L \times M_\psi} \quad (32)$$

where  $\circ$  denotes the Hadamard product (element-wise product). In (32),  $\mathbf{D}_\psi \in \mathbb{C}^{L \times M_\psi}$  is the normalization matrix which guarantees a unit norm for each column of  $\mathbf{V}_\psi$ , and whose  $\ell$ -th ( $\ell \in [M_\psi]$ ) column takes the form  $\mathbf{D}_{\psi,\ell} \triangleq \mathbf{1}_L \left( [\mathbf{H}_\psi^T \mathbf{H}_\psi^*]_{\ell,\ell} \right)^{-1/2}$ , where  $\mathbf{1}_L \in \mathbb{C}^L$  is the vector with all elements equaling 1.

*Remark 7:* Computing the ZF precoder  $\mathbf{V}_\psi$  in (32) has complexity  $\mathcal{O}(M_\psi^3 + LM_\psi^2)$ . Under the common massive-MIMO regime with  $L \gg M_\psi$ , the complexity order is asymptotically  $\mathcal{O}(L)$ . In contrast, BD-MRC for the same group  $\psi$  scales asymptotically as  $\mathcal{O}(L^2)$ , as discussed in Remark 3. Thus, the ZF design asymptotically achieves an order- $L$  reduction in computational complexity compared to performing BD-MRC via Lemma 1.

We first present the upper and lower bounds for the effective rate at any given user, averaged over channel fading.

*Proposition 1:* The effective rate  $\bar{R}_{\psi,k}^{\text{ZF}}$  averaged over channel fading at user  $\mathbf{U}_{\psi,k}$  for  $\psi \in \Psi, k \in [Q]$  under ZF-based precoding, is bounded as  $\tilde{R}_{\psi,k}^{\text{ZF}} \leq \bar{R}_{\psi,k}^{\text{ZF}} \leq \hat{R}_{\psi,k}^{\text{ZF}}$ , where

$$\tilde{R}_{\psi,k}^{\text{ZF}} \triangleq \xi_{G,Q} \sum_{q=1}^{M_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}(L - M_\psi)\beta_{\psi,k}}{N_0} \right), \quad (33)$$

$$\hat{R}_{\psi,k}^{\text{ZF}} \triangleq \xi_{G,Q} \sum_{q=1}^{M_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}(L - M_\psi + 1)\beta_{\psi,k}}{N_0} \right). \quad (34)$$

*Proof:* The proof is relegated to Appendix III. ■

Since directly solving the MMF optimization problem in (20) is computationally prohibitive, especially in large-scale settings with many served users and large antenna arrays, we take an alternative approach by optimizing strict upper and lower bounds on the effective sum-rate. Specifically, instead of solving the original MMF problem directly, we formulate two separate optimization problems by treating the upper and lower bounds derived in Proposition 1 as the objective functions within the MMF framework in Theorem 2. This allows us to obtain analytically optimized bounds on the sum-rate after MMF optimization, thereby providing rigorous performance guarantees. Furthermore, when all users are equipped with the same number of receive antennas, the optimized bounds lead to

closed-form expressions, significantly simplifying performance analysis. We now proceed with the theorem.

*Theorem 2:* The optimal MMF-constrained effective sum-rate  $\bar{R}_{\text{ZF}}^*$  is bounded as  $\tilde{R}_{\text{ZF}}^* \leq \bar{R}_{\text{ZF}}^* \leq \hat{R}_{\text{ZF}}^*$ , where  $\tilde{R}_{\text{ZF}}^*$  and  $\hat{R}_{\text{ZF}}^*$  are respectively the solutions to<sup>6</sup>

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} \frac{N_0 M_{\psi,k} \left( \exp \left( \frac{\tilde{R}_{\text{ZF}}^*}{\xi_{G,Q} GQM_{\psi,k}} \right) - 1 \right)}{\beta_{\psi,k}(L - M_\psi)} = P_{\text{tot}}, \quad (35)$$

$$\sum_{\psi \in \Psi} \sum_{k \in [Q]} \frac{N_0 M_{\psi,k} \left( \exp \left( \frac{\hat{R}_{\text{ZF}}^*}{\xi_{G,Q} GQM_{\psi,k}} \right) - 1 \right)}{\beta_{\psi,k}(L - M_\psi + 1)} = P_{\text{tot}}. \quad (36)$$

In the symmetric case of  $M_{\psi,k} = M$  for  $\forall \psi \in \Psi, k \in [Q]$ , the closed-form expressions for  $\tilde{R}_{\text{ZF}}^*$  and  $\hat{R}_{\text{ZF}}^*$  are respectively

$$\tilde{R}_{\text{ZF}}^* = \xi_{G,Q} GQM \ln \left( 1 + \frac{P_{\text{tot}}(L - QM)}{MN_0 \sum_{\psi \in \Psi} \sum_{k \in [Q]} \beta_{\psi,k}^{-1}} \right), \quad (37)$$

$$\hat{R}_{\text{ZF}}^* = \xi_{G,Q} GQM \ln \left( 1 + \frac{P_{\text{tot}}(L - QM + 1)}{MN_0 \sum_{\psi \in \Psi} \sum_{k \in [Q]} \beta_{\psi,k}^{-1}} \right). \quad (38)$$

*Proof:* The proof follows a similar procedure to the BD-MRC case. We first show that, at the optimum of the corresponding MMF problem, all users attain the same lower bound on their effective rate. Then, by analyzing the lower bound, we derive the transmit power allocated to each user and sum over all served users to match the total power  $P_{\text{tot}}$ . This leads to (35). In the symmetric case where all users have the same number of receive antennas  $M$ , a closed-form expression for the lower bound in (37) can be further obtained. The upper bound can be derived using similar steps. ■

As shown in Figs. 2, 5, and 6, the tightness of the derived bounds in Theorem 2 in estimating the actual performance is validated under various system configurations.

It is important to note that while ZF precoding is particularly beneficial in the large-antenna regime, the results presented in Theorem 2 are not limited to large  $L$ . The derived upper and lower bounds remain valid for a wide range of antenna configurations, making them applicable to both moderate and large-scale systems. These results offer an efficient yet mathematically rigorous way to evaluate MMF-optimized sum-rate behavior while reducing computational complexity.

By comparing  $\tilde{R}_{\text{ZF}}^*$  and  $\hat{R}_{\text{ZF}}^*$  with  $R_{\text{BD-MRC}}^*$  (cf. (31)) in the symmetric setting, we obtain the following limit results:

$$\lim_{L \rightarrow \infty} \frac{|R_{\text{BD-MRC}}^* - \tilde{R}_{\text{ZF}}^*|}{\xi_{G,Q} GQM} = \lim_{L \rightarrow \infty} \frac{M}{L - QM} = 0, \quad (39)$$

$$\lim_{L \rightarrow \infty} \frac{|R_{\text{BD-MRC}}^* - \hat{R}_{\text{ZF}}^*|}{\xi_{G,Q} GQM} = \lim_{L \rightarrow \infty} \frac{M - 1}{L - QM + 1} = 0. \quad (40)$$

<sup>6</sup>Unlike in Theorem 1 for BD-MRC, where the MMF optimization is performed before averaging over channel fading, we change the order of operations in  $\tilde{R}_{\text{ZF}}^*$ . Specifically, we first take the expectation over channel fading and then perform the MMF optimization. This reformulation allows us to design the power allocation strategy based on large-scale fading and path loss. While this approach simplifies the optimization process, it introduces an approximation compared to directly solving the MMF problem before averaging over channel fading, as done in the BD-MRC case.

TABLE II: Simulation Parameters (cf. [46], [48])

| Parameters                              | Value   |
|---|---|
| AWGN spectral density                   | -174 dBm/Hz   |
| Spectrum bandwidth for each user        | 20 MHz  |
| Carrier frequency                       | 2 GHz   |
| Macro-cell size                         | Inner radius of 35 m and outer radius of 500 m  |
| Micro-cell size                         | Inner radius of 10 m and outer radius of 100 m  |
| Pathloss exponent $\eta_0$              | $\eta_0 = 3.76$ in Macro-cell and $\eta_0 = 3$ in Micro-cell  |
| Attenuation regularization $l_0$        | $l_0 = 10^{-3.53}$ in Macro-cell and $l_0 = 10^{-3.7}$ in Micro-cell  |
| Pathloss model $\beta_{\psi,k}$         | $\beta_{\psi,k} = l_0 r_{\psi,k}^{-\eta_0}$ where $r_{\psi,k}$ is the distance from the BS to user $U_{\psi,k}$ |
| Maximum transmit power $P_{\text{tot}}$ | Typical values are around 33 dBm in Micro-cell, and 40 dBm in Macro-cell  |

These results lead to the following remark.

*Remark 8:* In the massive MIMO regime (i.e.,  $L \rightarrow \infty$ ), the delivery performance gap between simple ZF and BD-MRC is negligible when the number of receive antennas per user  $M$  is small. This conclusion is further supported by the numerical results in Section VI (e.g., Figs. 2 and 5).

## VI. NUMERICAL RESULTS

This section presents various numerical results that validate our analysis as well as provide clear comparisons<sup>7</sup>. We here focus on the case where each user has the same number of receiving antennas  $M$ . We consider users with relatively low mobility, assuming a coherence block of  $T = 15,000$  symbols, which corresponds to a coherence time of 0.075 seconds and a coherence bandwidth of 200 kHz.<sup>8</sup> We further consider CSI pilot length of  $\Theta = 10$ , which, under some ideal conditions, could provide near-perfect CSI at both the BS and the users [38]<sup>9</sup>. We consider independent Rayleigh fading channels and generate 1000 realizations of users' locations, based on the assumption of uniformly-distributed users across the cell.<sup>10</sup> Other simulation parameters are listed in Table II.

<sup>7</sup>The bounds in Theorem 2 have been demonstrated to be very tight in our conference version [1]. For clarity of presentation, we only plot the tight lower and upper bounds for ZF precoding as benchmarks to BD-MRC, while omitting the simulation results since they almost perfectly overlap with the bounds and are thus indistinguishable.

<sup>8</sup>For practical reference, at a carrier frequency of  $f_c = 2$  GHz in an urban Micro-cell scenario, the RMS delay spread (DS) is approximately 114 nanoseconds [46, Table 7.5-6]. Using the empirical relation  $B_c = 1/(50 \text{ DS})$  for a correlation coefficient of 0.9 [47, Eq. (1.23)], the coherence bandwidth  $B_c$  is about 180 kHz. Considering a pedestrian moving at speed  $v = 3$  km/h, the maximum Doppler shift is  $f_D = v f_c / c \approx 5.5$  Hz [47, Eq. (1.34)], where  $c$  denotes the light speed. This yields a coherence time of  $T_c = 0.423/f_D \approx 0.077$  seconds [47, Eq. (1.31)].

<sup>9</sup> $\Theta$  could be further decreased at higher SNR, thus further reducing the CSI overhead, resulting in even higher performance gains brought about by VCC.

<sup>10</sup>To the best of our knowledge, there are currently no existing multi-antenna coded caching schemes that significantly outperform conventional cacheless MU-MIMO in terms of spectral efficiency. Although this work mainly focuses on deriving analytical expressions that parametrize the delivery performance of VCC, we still provide numerical comparisons between VCC and existing multi-antenna coded caching schemes in Section VI-C.

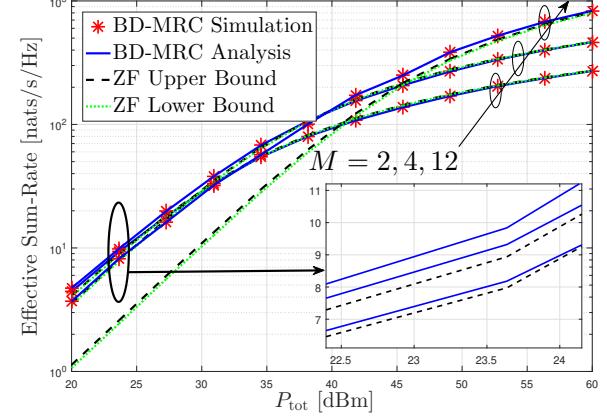


Fig. 2: Effective sum-rate versus  $P_{\text{tot}}$  in a Macro-cell under MMF with  $L = 64$ ,  $G = 5$ ,  $Q = 4$  and  $J = M$ .

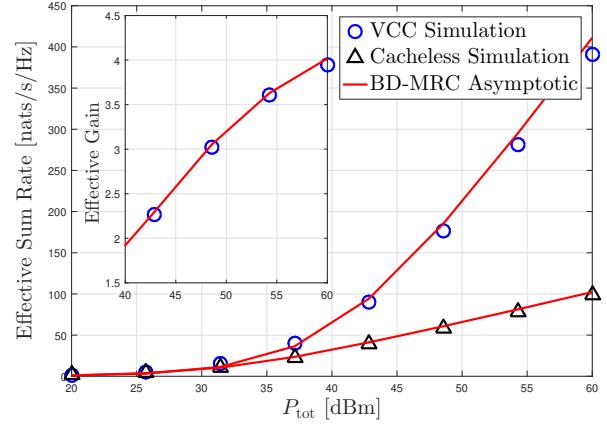


Fig. 3: Effective sum-rate and effective gain versus  $P_{\text{tot}}$  for  $M = 4$ ,  $L = 24$ ,  $G = 6$  and  $Q = Q' = 4$  in a Macro-cell under BD-MRC and MMF. Typical transmit power  $P_{\text{tot}}$  between 40–43 dBm.

### A. Numerical Comparisons With Fixed $Q$ and $Q'$

Fig. 2 plots the effective sum-rate versus  $P_{\text{tot}}$  under the MMF power allocation. The analytical results represented by blue solid lines, derived from Theorem 1, are obtained via a one-dimensional binary search for  $R_{\text{BD-MRC}}^*$ . For the simulation results represented by red asterisk symbols in the same figure, we use the built-in function "fminmax" in MATLAB to numerically solve the MMF optimization in (21). Interestingly, we note that the effective sum-rate increases with  $M$  when SNR is higher, while it decreases with  $M$  when SNR is relatively low, and where naturally spatial diversity is more impactful than multiplexing gain. For relatively small values of  $M$  (e.g.,  $M \leq 4$ ), simple ZF precoding exhibits reasonably good performance, being only slightly inferior to BD-MRC across the entire considered SNR range. This observation aligns with the analysis in Remark 8. However, for larger values of  $M$  (e.g.,  $M = 12$ ), its performance in the low-SNR regime deteriorates significantly, leading to a substantial performance gap compared to BD-MRC in this region.

Fig. 3 shows both the effective sum-rate and the corresponding gain versus  $P_{\text{tot}}$  for  $G = 6$ ,  $L = 24$ ,  $M = 4$  and  $Q = Q' = 4$  under BD-MRC and MMF in the Macro-cell.

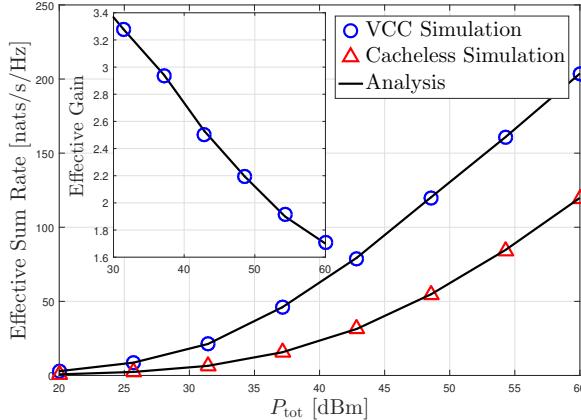


Fig. 4: Effective sum-rate and effective gain versus  $P_{\text{tot}}$  for  $L = 32$ ,  $M = 4$ ,  $Q = 2$ ,  $Q' = 8$  and (a very conservative)  $G = 4$  in a Macro-cell under BD-MRC and MMF.

This setup is motivated by the assumption that the precoder size remains unchanged in both the cacheless and VCC cases. As seen from the VCC transmit signal in (1), the system can sequentially precode  $G$  data vectors — each of dimension  $QM \times 1$  — and aggregate them for simultaneous transmission. In contrast, a conventional MU-MIMO system would transmit these vectors one after another. This structure allows VCC to support  $G$  times more users without requiring a larger precoder, making it compatible with existing transmitter architectures. In Fig. 3, the simulated results are generated following the same procedure as in Fig. 2, and the red curve corresponds to the asymptotic expression derived in Lemma 2. Even for moderate values of  $L$ , the asymptotic result closely matches the simulation, demonstrating its high accuracy. In the effective gain plot, it is observed that for practical transmit power levels  $P_{\text{tot}}$ , the spectral efficiency is nearly *doubled* compared to the conventional (cacheless) MU-MIMO system. Moreover, the effective gain continues to increase as  $P_{\text{tot}}$  grows.

The delivery performance with the same DoF (i.e.,  $Q' = GQ$ ) in VCC and in the cacheless scenario under BD-MRC and MMF is plotted in Fig. 4. In contrast to the increasing effective gain w.r.t.  $P_{\text{tot}}$  that we experience in Fig. 3, the effective gain in Fig. 4 decreases with  $P_{\text{tot}}$ , while still though allowing for good performance gains, especially in low SNR. This benefits from the fact that VCC introduces an extra multiplexing dimension for inter-user interference cancellation. This considerably alleviates the multiplexing load that should have been handled by multiple antennas, thus yielding greater beamforming gains for the users, which is particularly meaningful when requiring high beamforming gains in low SNR. Even at practical transmit power levels ( $P_{\text{tot}} \in [40, 43]$  dBm), the effective gain achieved by VCC remains above 230%, indicating a 1.3-fold improvement in spectral efficiency over the cacheless MU-MIMO system.

#### B. Numerical Comparisons With Optimized $Q$ and $Q'$

Now we shift to the case of the optimal balance between multiplexing and beamforming gains, which will be shown to be particularly meaningful in the medium to high SNR regime.

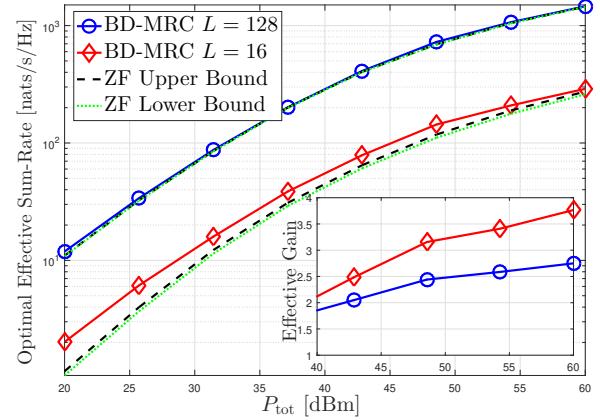


Fig. 5: Optimal effective sum-rate and effective gain versus  $P_{\text{tot}}$  for  $G = 6$ ,  $M = 4$  in a Macro-cell under MMF, where  $Q$  and  $Q'$  are independently optimized.

Fig. 5 plots the effective gain (corresponding to optimized rates — where  $Q$  and  $Q'$  are independently optimized) versus  $P_{\text{tot}}$ , for the Macro-cell case. While VCC provides a substantial gain over the corresponding traditional (cacheless) MU-MIMO scenario, this gain is lower than in the symmetric fading case (i.e., identical pathloss) studied in [25]. This reduction is primarily due to the severe near-far effect, which significantly constrains the effective sum-rate — particularly in large cells, even with optimal power allocation.<sup>11</sup> In contrast, things are very different in the Micro-cell setting (see Fig. 6) where the effective gain of optimized rates is again very notable. For example, in a Micro-cell setting, the recorded gain is 410% for the reasonable BS transmit power of  $P_{\text{tot}} = 33$  dBm. It is also worth noting that simple ZF precoding — with power allocation based on pathloss among the served users — yields a tight lower bound for the BD-MRC scheme when  $L$  is large and  $M$  is relatively small, as observed in Figs. 5–6. This observation is consistent with the analytical result in Remark 8.

#### C. Comparisons to Bit-Level Multi-Antenna Coded Caching

Since coded caching is essentially a form of network coding, most of VCC schemes (e.g., [20]–[25]) admit a PHY representation consistent with (1), characterized by a superposition of precoded  $L \times 1$  signal streams. In contrast, the conventional multi-antenna framework corresponds to *bit-level coded caching* (BCC) (e.g., [7], [10], [12]), where subfiles intended for a given user set are combined through bitwise XOR operations. This motivates a comparison between the PHY delivery performance of VCC and BCC. However, as noted in the Introduction, most existing multi-antenna coded caching schemes—typically bit-level—ignore the practical subpacketization constraint, which limits their achievable DoF to approximately  $\Lambda\gamma + Q$ . Under finite file sizes, [3] revisits the canonical multi-server coded caching framework [7], [12]

<sup>11</sup>The situation is further exacerbated by the uniform distribution of users within the Macro-cell, with a majority positioned near the cell edge. For instance, 64.32% of users are more than 300 meters from the BS, where, given a typical Macro-cell BS transmit power of  $P_t = 40$  dBm, their average SNR falls below 12.55 dB. Additional cases are detailed in Table 4.1 of [42].

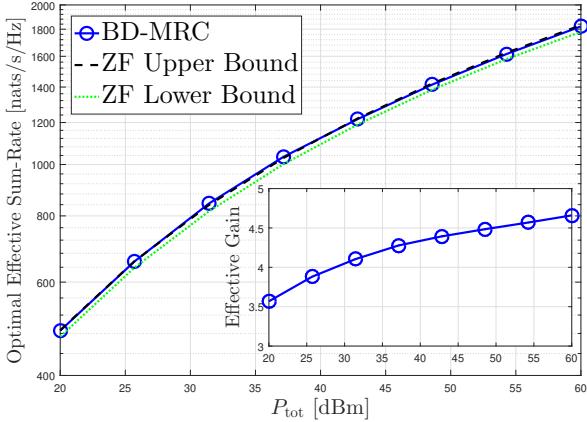


Fig. 6: Optimal effective sum-rate and effective gain versus  $P_{\text{tot}}$  for  $L = 32$ ,  $M = 2$ , and  $G = 6$  in a Micro-cell under MMF, where  $Q$  and  $Q'$  are independently optimized.

and proposes a reduced-complexity variant that preserves its main properties. We therefore adopt this *bit-level multi-server variant (MSV)* from [3] as the comparison baseline for VCC, and for simplicity, single-antenna users are considered, as also assumed in [3]. Furthermore, symmetric Rayleigh fading is assumed for all served users.

In the MSV scheme [3], the transmitter generates  $L$  signal streams, among which  $L - 1$  are dedicated to  $L - 1$  unicast users and one stream carries an XOR-coded (common) message intended for *another*  $G = \Lambda\gamma + 1$  users. Once the XOR message is decoded, each of these  $G$  users can reconstruct its desired subfile using the cached content. Under the equal power allocation policy, the transmit signal is given by

$$\mathbf{x}_{\text{MSV}} = \sqrt{\frac{P_{\text{tot}}}{L}} \mathbf{f}_0 c_0 + \sum_{k=1}^{L-1} \sqrt{\frac{P_{\text{tot}}}{L}} \mathbf{f}_k c_k, \quad (41)$$

where  $c_0$  represents the common XOR-coded message,  $c_k$  denotes the symbol for the  $k$ -th unicast user, and  $\mathbf{f}_0, \dots, \mathbf{f}_{L-1}$  are unit-norm precoding vectors of size  $L \times 1$ . The vector  $\mathbf{f}_0$  is designed to cancel the interference from the unicast signals at a *specific* multicast user (e.g., the first user), enabling that user to decode the XOR message. The remaining  $G - 1$  multicast users exploit their cached content to remove unicast interference before decoding the same XOR message. Consequently, the scheme simultaneously serves  $L - 1 + G = L + \Lambda\gamma$  users, which corresponds to its achievable DoF. Additional design details of the MSV framework can be found in [3].

Let  $\mathbf{h}_{\text{uc},k} \in \mathbb{C}^{L \times 1}$  ( $k \in [L - 1]$ ) and  $\mathbf{h}_{\text{mc},k'} \in \mathbb{C}^{L \times 1}$  ( $k' \in [G]$ ) denote the channels of unicast and multicast users, respectively. The precoding vectors in (41) are designed as

$$\begin{aligned} \mathbf{f}_0 &\in \mathbb{O}\{\mathbf{h}_{\text{uc},1}, \mathbf{h}_{\text{uc},2}, \dots, \mathbf{h}_{\text{uc},L-1}\}, \\ \mathbf{f}_k &\in \mathbb{O}\{\mathbf{h}_{\text{mc},1}, \mathbf{h}_{\text{uc},1}, \dots, \mathbf{h}_{\text{uc},k-1}, \mathbf{h}_{\text{uc},k+1}, \dots, \mathbf{h}_{\text{uc},L-1}\}, \end{aligned}$$

which can be implemented using BD precoding.

The transmission rate of the XOR-coded message is

$$R_{\text{mc}} = G \min_{k' \in [G]} \ln \left( 1 + \frac{P_{\text{tot}}}{N_0 L} |\mathbf{h}_{\text{mc},k'} \mathbf{f}_0|^2 \right), \quad (42)$$

where the minimum operator ensures successful decoding of the multicast message among all  $G$  target users. The transmission rate of the  $k$ -th unicast stream is

$$R_{\text{uc},k} = \ln \left( 1 + \frac{P_{\text{tot}}}{N_0 L} |\mathbf{h}_{\text{uc},k} \mathbf{f}_k|^2 \right). \quad (43)$$

Since the spatial multiplexing resources are fully utilized, no additional dimension is available for beamforming toward individual users, leading to

$$\mathbf{h}_{\text{mc},k'} \mathbf{f}_0 \sim \mathcal{CN}(0, 1), \quad \forall k' \in [G], \quad (44)$$

$$\mathbf{h}_{\text{uc},k} \mathbf{f}_k \sim \mathcal{CN}(0, 1), \quad \forall k \in [L - 1]. \quad (45)$$

The effective total rate of the MSV scheme is of the form

$$R_{\text{MSV}} = \xi_{L+\Lambda\gamma} \left( R_{\text{mc}} + \sum_{k=1}^{L-1} R_{\text{uc},k} \right), \quad (46)$$

where  $\xi_{L+\Lambda\gamma} \triangleq 1 - \Theta(L + \Lambda\gamma)/T$  represents the CSI overhead factor. Considering BD precoding in the cacheless baseline, the effective gain of MSV is defined as

$$G_{\text{MSV}} = \frac{\mathbb{E}\{R_{\text{MSV}}\}}{\max_{Q' \in [L]} \mathbb{E}\{R_{\text{BD}}(1, Q')\}}, \quad (47)$$

where the expectation is taken over channel realizations, and  $R_{\text{BD}}(1, Q')$  denotes the effective rate of BD precoding analyzed in Section IV for single-antenna receivers under symmetric channel conditions ( $\beta_k = 1$ ,  $\forall k \in [K]$ ). We note that the original MSV scheme in [3] does not adjust the number of unicast streams in (41) to optimize the multiplexing–beamforming tradeoff, as in VCC. To fully exploit the transmission capability of the MSV scheme under finite SNR, we further adjust the number of unicast streams  $Q_{\text{uc}} \leq L - 1$  in (41) to achieve an optimized multiplexing–beamforming tradeoff, referred to as the *modified MSV scheme*, where the optimal BD design follows the same procedures as we considered in VCC (cf. Section III-B). The corresponding effective gain is defined as

$$G'_{\text{MSV}} \triangleq \frac{\max_{Q_{\text{uc}} \in [L-1]} \mathbb{E}\{R_{\text{MSV}}(Q_{\text{uc}})\}}{\max_{Q' \in [L]} \mathbb{E}\{R_{\text{BD}}(1, Q')\}}. \quad (48)$$

As  $P_{\text{tot}} \rightarrow \infty$ , the optimal  $Q'$  in the cacheless BD baseline equals  $L$ , leading to the high-SNR limit for MSV as

$$G_{\text{MSV}} \longrightarrow \frac{L + \Lambda\gamma}{L}, \quad (49)$$

which approaches unity as  $L \rightarrow \infty$ . Even for moderate values, such as  $L = 32$  and  $G = 6$ , the high-SNR limit gain remains marginal, which equals  $\frac{32+5}{32} = 1.1563$  as shown in Fig. 7.

Fig. 7 compares the effective gains of MSV and BD-based VCC against the cacheless BD baseline. Both MSV and VCC exhibit increasing gains with SNR, yet VCC shows a notably sharper rise, implying a growing advantage over the cacheless baseline. In contrast, the effective gain of original MSV in [3] remains below 1 at finite SNRs, indicating that the cacheless baseline consistently outperforms original MSV. The main reasons are threefold: *i*) the DoF limitation of MSV, equal to  $L + \Lambda\gamma$ , providing only marginal gain for large  $L$  at finite SNRs; *ii*) the inherent worst-user bottleneck in multicasting, where the XOR rate is constrained by the weakest user (cf.

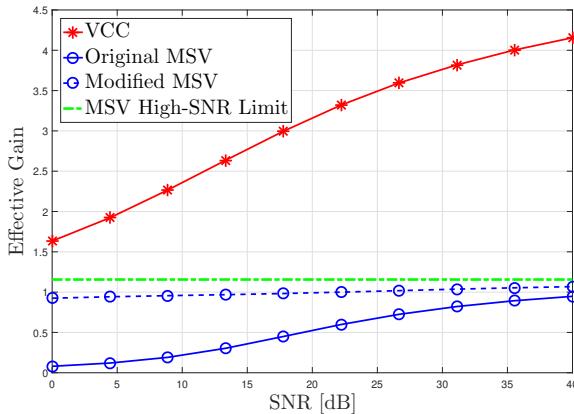


Fig. 7: Effective gain versus  $\text{SNR} \triangleq P_{\text{tot}}/N_0$  for  $L = 32$  and  $G = 6$  over symmetric Rayleigh fading channels

(42)); and *iii*) the additional CSI acquisition required for approximately  $\Lambda\gamma$  more users compared with the cacheless case, which further reduces the effective throughput. As illustrated in Fig. 7, while the modified MSV scheme achieves the optimal multiplexing–beamforming tradeoff at finite SNRs, its effective gain, though clearly improved, remains close to one, implying limited advantage over the cacheless baseline. This result corroborates the claim made in the Introduction that the caching gains achieved by bit-level multi-antenna coded caching are overshadowed by the conventional multiplexing gains.

#### D. VCC With Imperfect CSIT and CSIR

In this subsection, we present two parts. The first part considers the case with imperfect CSIT while assuming perfect CSI at the receiver (CSIR). The second part extends the study to include both imperfect CSIT and CSIR.

1) *Imperfect CSIT*: We now turn to the case of imperfect CSIT — a common scenario in practice due to pilot contamination, estimation errors, and feedback delays. While accurate CSI typically enhances spectral efficiency, Fig. 8 illustrates that VCC can, in some cases, achieve even higher gains under imperfect CSIT than in the perfect CSIT setting. This is because inter-group interference can be mitigated using cached content, independent of CSIT quality — a sharp contrast to conventional cacheless MU-MIMO, which relies heavily on accurate CSIT for interference management. Throughout, we assume perfect CSIR,<sup>12</sup> reflecting its more robust acquisition in practical systems. To simplify the analysis, we focus on single-antenna receivers under ZF precoding, aiming to assess whether the reported gains persist despite CSIT imperfections.

With single-antenna receivers, the channel between the BS and user  $U_{\psi,k}$  (for any  $\psi \in \Psi$  and  $k \in [Q]$ ) is now represented as a vector, denoted by  $\mathbf{h}_{\psi,k} \in \mathbb{C}^L$ . Let  $\hat{\mathbf{h}}_{\psi,k}$  denote the CSI estimate at the BS. Considering the linear MMSE estimation [49], we can express  $\mathbf{h}_{\psi,k}$  as  $\mathbf{h}_{\psi,k} = \hat{\mathbf{h}}_{\psi,k} + \tilde{\mathbf{h}}_{\psi,k}$ , where  $\tilde{\mathbf{h}}_{\psi,k}$  represents the estimation errors, whose elements follow the

<sup>12</sup>CSIR is typically more accurate than CSIT, as it is estimated directly from pilots at each receiver. In contrast, CSIT acquisition — via feedback in FDD or uplink pilots in TDD — is more error-prone due to delay, quantization, interference, and user coordination. This disparity further amplifies the relative gains of VCC.

i.i.d. Gaussian distribution with zero-mean and variance  $\tilde{\beta}_{\psi,k}$ . Naturally,  $\tilde{\mathbf{h}}_{\psi,k}$  is independent of the MMSE estimate  $\hat{\mathbf{h}}_{\psi,k}$ .

The BS treats  $\hat{\mathbf{h}}_{\psi,k}$  as the real channel vector and then follows the same paradigm as in (32) to perform ZF precoding. Let  $\hat{\mathbf{V}}_{\psi}$  denote the ZF matrix based on the MMSE estimates  $\{\hat{\mathbf{h}}_{\psi,k} : k \in [Q]\}$ . Following the channel training method in [38], after the BS finishes the CSI estimation, the BS issues an additional downlink training phase, where it sends several orthogonal training symbols along each of the precoding vectors to user  $U_{\psi,k}$  for any  $\psi \in \Psi$  and  $k \in [Q]$  for coherent detection. By doing so,  $U_{\psi,k}$  can obtain the coupling (composite) coefficients  $\{\sqrt{P_{\phi,k'}} \mathbf{h}_{\psi,k'}^T \hat{\mathbf{v}}_{\phi,k'} : \phi \in \Psi, k' \in [Q]\}$ , where  $\mathbf{v}_{\phi,k'}$  denotes the  $k'$ -th column of  $\hat{\mathbf{V}}_{\phi}$ . As the downlink training is assumed to be perfect,  $U_{\psi,k}$  detects these coupling coefficients perfectly (perfect composite CSIR), which enables it to completely cancel the inter-group interference by applying its cached content, as described in Section II. After removing the inter-group interference, the remaining signal at  $U_{\psi,k}$  is

$$y'_{\psi,k} = \sqrt{P_{\psi,k}} \mathbf{h}_{\psi,k}^T \hat{\mathbf{v}}_{\psi,k} s_{\psi,k} + z_{\psi,k} + \sum_{k' \in [Q] \setminus k} \sqrt{P_{\psi,k'}} \mathbf{h}_{\psi,k'}^T \hat{\mathbf{v}}_{\psi,k'} s_{\psi,k'}, \quad (50)$$

yielding a SINR when decoding  $s_{\psi,k}$  of the form

$$\text{SINR}_{\psi,k}^{\text{ZF}} = \frac{P_{\psi,k} |\mathbf{h}_{\psi,k}^T \hat{\mathbf{v}}_{\psi,k}|^2}{N_0 + \sum_{k' \in [Q] \setminus k} P_{\psi,k'} |\mathbf{h}_{\psi,k'}^T \hat{\mathbf{v}}_{\psi,k'}|^2}, \quad (51)$$

and an effective sum rate, under imperfect CSIT, of the form

$$R_{\text{sum}}^{\text{ZF}}(G, Q) = \left(1 - \frac{GQ\Theta}{T}\right) \sum_{\psi \in \Psi} \sum_{k \in [Q]} \ln(1 + \text{SINR}_{\psi,k}^{\text{ZF}}), \quad (52)$$

where  $\text{SINR}_{\psi,k}^{\text{ZF}}$  is given by (51). The effective sum rate of the corresponding cacheless case is  $R_{\text{sum}}^{\text{ZF}}(1, Q')$ , where  $Q'$  is the operational multiplexing gain under conventional ZF precoding. As shown in [42], power allocation for MMF optimization marginally improves the effective gain, despite the notable increase in the effective sum-rate. For simplicity, we consider equal power allocation among the users and then consider the effective gain as defined below

$$G'_{\text{ZF}} \triangleq \frac{\max_{Q \in [L]} \mathbb{E}\{R_{\text{sum}}^{\text{ZF}}(G, Q)\}}{\max_{Q' \in [L]} \mathbb{E}\{R_{\text{sum}}^{\text{ZF}}(1, Q')\}}, \quad (53)$$

where the expectation is averaged over channel states.

Fig. 8 is obtained by selecting the optimal multiplexing gain for each SNR value, while the effective gain is computed according to (53). As expected, under imperfect CSIT, VCC exhibits lower effective sum-rate than the perfect CSIT case in the medium-to-high SNR regime, since the system is interference-limited due to CSIT estimation errors. Surprisingly, however, the effective gain achieved by VCC under imperfect CSIT surpasses that of the perfect CSIT case in the medium-to-high SNR region. This phenomenon is primarily attributed to the fact that, in VCC, inter-group interference is completely canceled without relying on CSIT, significantly mitigating the impact of CSIT errors on system performance. In contrast, for the cacheless MU-MIMO system, the elimination of inter-user interference heavily depends on the quality of the CSIT, making its performance more sensitive to CSIT imperfections.

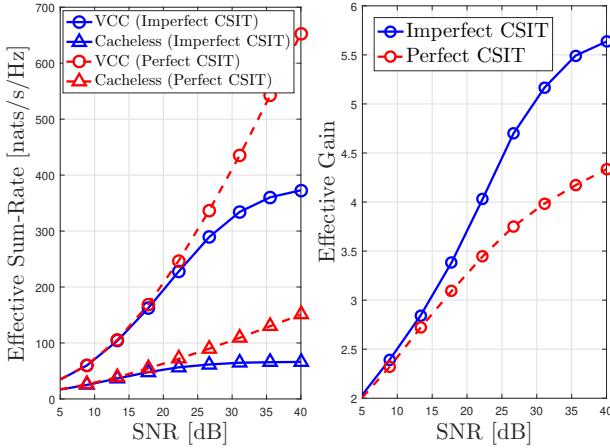


Fig. 8: Effective sum-rate and effective gain versus SNR for  $L = 16$ ,  $M = 1$  and  $G = 6$  under ZF precoding, where the users are statistically symmetric. Specifically,  $\beta_{\psi,k} = 1$ , and  $\tilde{\beta}_{\psi,k} = \tilde{\beta} = 0.01$  for any served user in the case of imperfect CSIT. SNR denotes the average transmit SNR  $\triangleq P_{\text{tot}}/N_0$ .

2) *Imperfect CSIT and CSIR*: In the following, we examine the impact of imperfect CSIT and CSIR on VCC delivery. Let  $A_{\psi,k}^{\phi,\vartheta} \triangleq \mathbf{h}_{\psi,k}^T \hat{\mathbf{v}}_{\phi,\vartheta}$  denote the coupling coefficient for all  $\psi, \phi \in \Psi$  and  $k, \vartheta \in [Q]$ . Under MMSE estimation at the receiver,  $A_{\psi,k}^{\phi,\vartheta} = \hat{A}_{\psi,k}^{\phi,\vartheta} + \tilde{A}_{\psi,k}^{\phi,\vartheta}$ , where  $\hat{A}_{\psi,k}^{\phi,\vartheta}$  is the estimate and  $\tilde{A}_{\psi,k}^{\phi,\vartheta} \sim \mathcal{CN}(0, \tilde{\beta}')$  is the error. For simplicity, we consider a symmetric case with identical error variance  $\tilde{\beta}'$  for all coefficients.

With imperfect CSIR, cache-aided interference cancellation is *not perfect* and leaves *residual interference*. The signal after imperfect cache-aided interference cancellation is

$$y_{\psi,k}'' = y_{\psi,k}' + \sum_{\phi \in \Psi \setminus \psi} \sum_{\vartheta \in [Q]} \sqrt{P_{\phi,\vartheta}} \tilde{A}_{\psi,k}^{\phi,\vartheta} s_{\phi,\vartheta}, \quad (54)$$

where  $y_{\psi,k}'$  is the signal for the imperfect-CSIT-only case (see (50)). Assuming equal power allocation and ZF precoding, the resulting SINR is

$$\text{SINR}_{\psi,k}^{\text{ZF}} = \frac{\frac{P_{\text{tot}}}{GQ} |\hat{A}_{\psi,k}^{\phi,\vartheta}|^2 + \frac{P_{\text{tot}}}{GQ} \tilde{\beta}'}{N_0 + I}, \quad (55)$$

where the interference term  $I$  accounts for residuals due to *both* imperfect CSIT and CSIR:

$$\begin{aligned} I &\triangleq \frac{P_{\text{tot}}}{GQ} \sum_{k' \in [Q] \setminus k} \mathbb{E}\{| \tilde{A}_{\psi,k'}^{\phi,\vartheta} |^2\} + \frac{P_{\text{tot}}}{GQ} \sum_{\phi \in \Psi \setminus \psi} \sum_{\vartheta \in [Q]} \mathbb{E}\{| \tilde{A}_{\psi,k}^{\phi,\vartheta} |^2\} \\ &= \frac{P_{\text{tot}}}{GQ} \tilde{\beta}' (GQ - 1). \end{aligned} \quad (56)$$

Substituting (55) into (52) and (53) yields the effective gain under simultaneous CSIT and CSIR imperfections.

Fig. 9 illustrates the effective gain of VCC under both imperfect CSIT and CSIR conditions with varying CSIR estimation accuracy. At low SNR, the effective gains under different CSIR qualities are close to that of the perfect-CSI case, consistent with the observation in Fig. 8. As the SNR increases, however, residual inter-user interference caused by imperfect CSIT and CSIR becomes more significant, leading to a gradual reduction in the effective gain. In practical systems, the CSIR estimation accuracy is usually much higher than that of CSIT, i.e.,  $\tilde{\beta}' \ll \tilde{\beta}$ , where  $\tilde{\beta}$  denotes the variance of CSIT

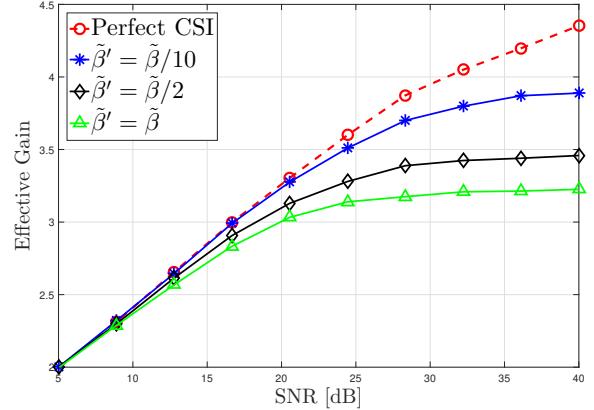


Fig. 9: Effective gain versus SNR under the same parameters as Fig. 8, with additional imperfect CSIR consideration.

errors. Therefore, the case of  $\tilde{\beta}' = \tilde{\beta}$  (green curve) represents a worst-case scenario. Even in this worst case, VCC still achieves *more than triple* the spectral efficiency of the cacheless baseline when SNR exceeds 20 dB. The numerical comparisons in Fig. 9 confirm that VCC maintains a substantial multiplicative spectral-efficiency gain over cacheless MU-MIMO systems, even in the presence of simultaneous CSIT and CSIR imperfections.

## VII. CONCLUSIONS AND FUTURE WORK

We investigated VCC under realistic conditions, including variable pathloss, multi-antenna receivers, MMF, and practical precoders and combiners. We derived analytical expressions for the effective sum-rate and gain under BD-MRC and ZF schemes, along with several closed-form results under simplifying assumptions. Numerical evaluations confirmed the tightness of these expressions and demonstrated the significant gains of VCC in realistic scenarios. These gains are particularly pronounced in Micro-cell environments, where we observed spectral efficiency improvements exceeding 300% over conventional (cacheless) MU-MIMO systems. Moreover, Fig. 7 compared the effective gains of VCC with the multi-server scheme and its variants, demonstrating its performance advantage over conventional multi-antenna coded caching approaches. Importantly, all reported gains in Figs. 5–9 represent improvements over optimized traditional MU-MIMO baselines, including any additional CSI acquisition costs.

In addition to the obvious multiplexing dimension for interference cancellation enabled by additional memory, VCC enhances beamforming gains which are especially valuable in low-SNR regimes. The technique also remains effective even under constraints such as a limited number of active signal streams (e.g.,  $GQM \leq 32$  in Fig 4). Beyond its theoretical advantages, VCC offers strong practical value. With VoD representing over 70% of network traffic and cacheable content comprising nearly 90% of data consumption, such caching-based strategies present a scalable solution to enhance spectral efficiency and reduce congestion in high-demand, content-heavy networks.

It is worth noting that the simulation results presented are somewhat conservative. In particular, the performance gains of the proposed cache-aided MU-MIMO scheme become

significantly more pronounced when imperfect CSIT is taken into account. This is a key strength of our approach: its inherent robustness and improved performance under CSIT imperfections, by virtue of the fact that VCC shifts much of the interference management to the receivers, inherently reducing per-user CSIT needs. A preliminary indication of these enhanced gains was illustrated in Fig. 8. Future work may focus on extending the current framework to explicitly characterize the impact of imperfect CSIT and CSIR. In particular, it would be valuable to analyze the performance under different CSI quality levels and investigate the trade-offs among CSI accuracy, feedback overhead, and system throughput. When VCC is applied in large-scale systems, several scalability challenges arise as the numbers of users and antennas increase. These include higher CSI overhead, power-efficiency concerns, and increased computational complexity; see [42] for further discussion. In addition, user synchronization—where users reveal their requests at different time instants—can create asymmetric signal structures during the delivery phase, leading to different dimensions of  $\mathbf{s}_\psi$  for various  $\psi \in \Psi$  in (1). A more comprehensive study of these issues, which forms a substantial research topic on its own, is deferred to future work.

#### APPENDIX I: PROOF OF LEMMA 1

We will find  $\mathbf{m}_{\psi,k,q}$  by decomposing a much smaller sized matrix  $\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}^2 \mathbf{H}_{\psi,k}^* = \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*$ . With  $\lambda_{\psi,k,q}$  and  $\mathbf{t}_{\psi,k,q}$  in place, we have

$$\mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q} = \lambda_{\psi,k,q} \mathbf{t}_{\psi,k,q}. \quad (57)$$

Multiplying  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*$  on both sides of (57) and using that  $\mathbf{T}_{\psi,-k} = \mathbf{T}_{\psi,-k}^2$  yields

$$\begin{aligned} \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} (\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}) \\ = \lambda_{\psi,k,q} (\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}), \end{aligned} \quad (58)$$

which shows that both  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}$  and  $\mathbf{m}_{\psi,k,q}$  are the eigenvectors associated to the  $q$ -th largest eigenvalue  $\lambda_{\psi,k,q}$  of  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$ . Furthermore, as  $\mathbf{t}_{\psi,k,q}$  and  $\mathbf{t}_{\psi,k,p}$  are orthogonal to each other for  $p \neq q$ , we have that

$$\begin{aligned} & \left( \mathbf{t}_{\psi,k,p}^H \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \right) \left( \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q} \right) \\ &= \mathbf{t}_{\psi,k,p}^H \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q} = \lambda_{\psi,k,q} \mathbf{t}_{\psi,k,p}^H \mathbf{t}_{\psi,k,q} = 0, \end{aligned}$$

which indicates that the eigenvectors  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}$  and  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,p}$  associated with different non-zero eigenvalues of  $\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}$  are also orthogonal to each other. Therefore, the precoding vector  $\mathbf{v}_{\psi,k,q}$  for  $\mathbf{U}_{\psi,k}$  in (5) can be rewritten as in (13).

Given the precoding vector  $\mathbf{v}_{\psi,k,q}$  in (13) and given an MRC receiver, the  $\text{SINR}_{\psi,k,q}^{\text{BD-MRC}}$  in (11) is derived as in (59), shown at the bottom of this page, where step (a) follows from (58), and which directly leads to (14).

$$\text{SINR}_{\psi,k,q}^{\text{BD-MRC}} = \frac{(\mathbf{t}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}) \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} (\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q})}{\|\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}\|^2 \frac{N_0}{P_{\psi,k,q}}} \stackrel{(a)}{=} \frac{\lambda_{\psi,k} (\mathbf{t}_{\psi,k,q}^H \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k}) (\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q})}{\|\mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \mathbf{t}_{\psi,k,q}\|^2 \frac{N_0}{P_{\psi,k,q}}} \quad (59)$$

#### APPENDIX II: PROOF OF LEMMA 2

For  $L \rightarrow \infty$  and finite  $M_{\psi,k}$ , we can use the Trace Lemma (cf. [50]) to derive that

$$\begin{aligned} \frac{1}{L} \frac{\mathbf{h}_{\psi,k,\ell}^T}{\sqrt{\beta_{\psi,k}}} \mathbf{T}_{\psi,-k} \frac{\mathbf{h}_{\psi,k,\ell}^*}{\sqrt{\beta_{\psi,k}}} \\ \xrightarrow{\text{a.s.}} \frac{1}{L} \text{Tr} \left\{ \mathbf{T}_{\psi,-k} \right\} \simeq 1 - \frac{M_\psi - M_{\psi,k}}{L}, \end{aligned} \quad (60)$$

$$\frac{1}{L} \mathbf{h}_{\psi,k,\ell}^T \mathbf{T}_{\psi,-k} \mathbf{h}_{\psi,k,\ell}^* \xrightarrow{\text{a.s.}} 0, \text{ for } \ell' \neq \ell, \quad (61)$$

where  $\mathbf{h}_{\psi,k,\ell}$  denotes the  $\ell$ -th column of  $\mathbf{H}_{\psi,k}$ , and where  $\xrightarrow{\text{a.s.}}$  denotes the almost sure convergence. In the above, we also consider that  $\text{Tr} \left\{ \mathbf{T}_{\psi,-k} \right\} = \text{Rank} \left\{ \mathbf{T}_{\psi,-k} \right\}$  because  $\mathbf{T}_{\psi,-k}$  is a projection matrix. Therefore, we can now derive that

$$\frac{1}{L} \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^* \simeq \beta_{\psi,k} \left( 1 - \frac{M_\psi - M_{\psi,k}}{L} \right) \mathbf{I}_{M_{\psi,k}}, \quad (62)$$

which reveals that all eigenvalues of  $\frac{1}{L} \mathbf{H}_{\psi,k}^T \mathbf{T}_{\psi,-k} \mathbf{H}_{\psi,k}^*$  become identical as  $L \rightarrow \infty$ . Therefore, we have the following asymptotic result

$$\lambda_{\psi,k,q} \simeq \beta_{\psi,k} (L - M_\psi + M_{\psi,k}). \quad (63)$$

Then, the effective rate in (15) has the asymptotic behavior

$$R_{\psi,k} \simeq \xi_{G,Q} \sum_{q=1}^{J_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}}{N_0} \beta_{\psi,k} (L - M_\psi + M_{\psi,k}) \right). \quad (64)$$

According to the properties of the water-filling algorithm, we know that equal-power allocation among  $\{s_{\psi,k,q} : q \in [J_{\psi,k}]\}$  is optimal, which in turn tells us that  $R_{\psi,k}^*(P_{\psi,k}) \simeq \bar{R}_{\psi,k}^*$  where  $\bar{R}_{\psi,k}^*$  is given by

$$\bar{R}_{\psi,k}^* \triangleq \xi_{G,Q} J_{\psi,k} \ln \left( 1 + \frac{P_{\psi,k}}{N_0 J_{\psi,k}} \beta_{\psi,k} (L - M_\psi + M_{\psi,k}) \right). \quad (65)$$

At this point, we use the asymptotic equivalence  $\bar{R}_{\psi,k}^*$  to substitute the objective function of the MMF optimization in (21), given by

$$\mathcal{S}_3 \left\{ \begin{array}{l} \max_{\mathcal{P}_\Psi} \min_{\psi \in \Psi} \min_{k \in [Q]} \bar{R}_{\psi,k}^* \\ \text{s. t. } P_t = \sum_{\psi \in \Psi} \sum_{k \in [Q]} P_{\psi,k} = P_{\text{tot}}. \end{array} \right. \quad (66)$$

Finally, as all users have the same effective rate under optimal power allocation for (66), we can use the total power constraint to derive (29)–(31) respectively.

#### APPENDIX III: PROOF OF PROPOSITION 1

For any  $\psi \in \Psi$ ,  $k \in [Q]$  and  $q \in [M_{\psi,k}]$ , let  $\mathbf{e}_{\psi,k,q} \in \mathbb{C}^{M_\psi \times 1}$  denote a vector with all zero elements except the  $k(q)$ -th element equalling 1, where  $k(q) \triangleq q + \sum_{k'=1}^{k-1} M_{\psi,k'}$ . The

precoding vector for symbol  $s_{\psi,k,q}$  under ZF precoding (cf. (1) and (32)) can be written as

$$\mathbf{v}_{\psi,k,q} = \mathbf{H}_{\psi}^* (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \mathbf{e}_{\psi,k,q} \sqrt{\left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)}}^{-1}. \quad (67)$$

After removing the inter-group interference and considering that  $\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^* (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} = \mathbf{I}_{M_{\psi}}$ , we simplify the received signal for decoding  $s_{\psi,k,q}$  in (3) under ZF precoding (cf. (32)) as

$$y'_{\psi,k,q} = \sqrt{\left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)}}^{-1} \sqrt{P_{\psi,k,q}} \mathbf{s}_{\psi,k,q} + z'_{\psi,k,q}. \quad (68)$$

Then, the effective average rate at  $\mathbf{U}_{\psi,k}$  is of the form

$$\bar{R}_{\psi,k}^{\text{ZF}} = \xi_{G,Q} \mathbb{E} \left\{ \sum_{q=1}^{M_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}/N_0}{\left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)}} \right) \right\} \quad (69)$$

$$\stackrel{(a)}{\geq} \xi_{G,Q} \sum_{q=1}^{M_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}/N_0}{\mathbb{E} \left\{ \left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)} \right\}} \right), \quad (70)$$

where (a) follows from using Jensen's inequality on the convex function  $\ln(1 + x^{-1})$ . Considering that  $\mathbb{E} \left\{ \left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)} \right\} = \frac{1}{\beta_{\psi,k}(L-M_{\psi})}$  (cf. [51]), we obtain the lower-bound in (33).

To obtain the upper-bound of  $\bar{R}_{\psi,k}^{\text{ZF}}$ , we re-use Jensen's inequality for the concave function  $\ln(1 + x)$  in (69), and obtain

$$\bar{R}_{\psi,k}^{\text{ZF}} \leq \xi_{G,Q} \sum_{q=1}^{M_{\psi,k}} \ln \left( 1 + \frac{P_{\psi,k,q}}{N_0} \mathbb{E} \left\{ \frac{1}{\left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)}} \right\} \right), \quad (71)$$

which induces the upper-bound in (34) by considering that  $\mathbb{E} \left\{ \left( (\mathbf{H}_{\psi}^T \mathbf{H}_{\psi}^*)^{-1} \right)_{k(q),k(q)} \right\} = \beta_{\psi,k}(L-M_{\psi}+1)$  (cf. [51]).

## REFERENCES

- [1] H. Zhao and P. Elia, "Vector coded caching substantially boosts MU-MIMO: Pathloss, CSI and power-allocation considerations," in *Proc. Int. 26th ITG Workshop on Smart Antennas and 13th Conf. on Systems, Commun., and Coding (WSA&SCC)*, Feb. 2023, pp. 37–42.
- [2] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [3] E. Lampiris and P. Elia, "Full coded caching gains for cache-less users," *IEEE Trans. Inf. Theory*, vol. 66, no. 12, pp. 7635–7651, Dec. 2020.
- [4] F. Engelmann and P. Elia, "A content-delivery protocol, exploiting the privacy benefits of coded caching," in *Proc. 15th Int. Symp. Model. Optim. Mobile, Ad Hoc Wireless Netw. (WiOpt)*, 2017, pp. 1–6.
- [5] E. Lampiris, A. Bazco-Nogueras, and P. Elia, "Resolving the feedback bottleneck of multi-antenna coded caching," *IEEE Trans. Inf. Theory*, vol. 68, no. 4, pp. 2331–2348, Apr. 2022.
- [6] Q. Yan, M. Cheng, X. Tang, and Q. Chen, "On the placement delivery array design for centralized coded caching scheme," *IEEE Trans. Inf. Theory*, vol. 63, no. 9, pp. 5821–5833, Sep. 2017.
- [7] S. P. Shariatpanahi *et al.*, "Multi-server coded caching," *IEEE Trans. Inf. Theory*, vol. 62, no. 12, pp. 7253–7271, Dec. 2016.
- [8] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, "Fundamental limits of cache-aided interference management," *IEEE Trans. Inf. Theory*, vol. 63, no. 5, pp. 3092–3107, May 2017.
- [9] S. P. Shariatpanahi, G. Caire, and B. H. Khalaj, "Physical-layer schemes for wireless coded caching," *IEEE Trans. Inf. Theory*, vol. 65, no. 5, pp. 2792–2807, May 2019.
- [10] K.-H. Ngo, S. Yang, and M. Kobayashi, "Scalable content delivery with coded caching in multi-antenna fading channels," *IEEE Trans. Wireless Commun.*, vol. 17, no. 1, pp. 548–562, Jan. 2018.
- [11] M. Bayat, R. K. Mungara, and G. Caire, "Achieving spatial scalability for coded caching via coded multi-point multicasting," *IEEE Trans. Wireless Commun.*, vol. 18, no. 1, pp. 227–240, Jan. 2019.
- [12] A. Tölli, S. P. Shariatpanahi, J. Kaleva, and B. H. Khalaj, "Multi-antenna interference management for coded caching," *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 2091–2106, Mar. 2020.
- [13] H. Bakhshzad Mahmoodi, M. Salehi, and A. Tölli, "Multi-antenna coded caching for location-dependent content delivery," *IEEE Trans. Wireless Commun.*, vol. 23, no. 1, pp. 408–422, Jan. 2024.
- [14] I. Bergel and S. Mohajer, "Cache-aided communications with multiple antennas at finite SNR," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 8, pp. 1682–1691, Aug. 2018.
- [15] X. Xu and M. Tao, "Modeling, analysis, and optimization of caching in multi-antenna small-cell networks," *IEEE Trans. Wireless Commun.*, vol. 18, no. 11, pp. 5454–5469, Nov. 2019.
- [16] M. J. Salehi *et al.*, "Low-complexity high-performance cyclic caching for large MISO systems," *IEEE Trans. Wireless Commun.*, vol. 21, no. 5, pp. 3263–3278, May 2022.
- [17] M. Salehi and A. Tölli, "Diagonal multi-antenna coded caching for reduced subpacketization," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2020.
- [18] M. Salehi, A. Tölli, and S. P. Shariatpanahi, "Subpacketization - beamformer interaction in multi-antenna coded caching," in *2020 2nd 6G Wireless Summit (6G SUMMIT)*, 2020, pp. 1–5.
- [19] "The industry's first independent benchmark study of 5G NR MU-MIMO," Signals Research Group, Tech. Rep., Sept. 2020.
- [20] E. Lampiris and P. Elia, "Adding transmitters dramatically boosts coded-caching gains for finite file sizes," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1176–1188, Jun. 2018.
- [21] H. Bakhshzad Mahmoodi, M. Salehi, and A. Tölli, "Low-complexity multi-antenna coded caching using location-aware placement delivery arrays," *IEEE Trans. Wireless Commun.*, vol. 23, no. 10, pp. 12687–12702, Oct. 2024.
- [22] E. Parrinello, A. Ünsal, and P. Elia, "Fundamental limits of coded caching with multiple antennas, shared caches and uncoded prefetching," *IEEE Trans. Inf. Theory*, vol. 66, no. 4, pp. 2252–2268, Apr. 2020.
- [23] E. Peter and B. S. Rajan, "Multi-antenna coded caching from a placement delivery array for shared caches," *IEEE Trans. Commun.*, vol. 70, no. 6, pp. 3627–3640, Jun. 2022.
- [24] T. Yang, K. Wan, M. Cheng, R. C. Qiu, and G. Caire, "Multiple-antenna placement delivery array for cache-aided MISO systems," *IEEE Trans. Inf. Theory*, vol. 69, no. 8, pp. 4855–4868, Aug. 2023.
- [25] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Vector coded caching multiplicatively boosts the throughput of realistic downlink systems," *IEEE Trans. Wireless Commun.*, vol. 22, no. 4, pp. 2683–2698, 2023.
- [26] M. Ji and R.-R. Chen, "Caching and coded multicasting in slow fading environment," in *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC)*, Mar. 2017, pp. 1–6.
- [27] B. Tegin and T. M. Duman, "Coded caching with user grouping over wireless channels," *IEEE Wireless Commun. Lett.*, vol. 9, no. 6, pp. 920–923, Jun. 2020.
- [28] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Wireless coded caching with shared caches can overcome the near-far bottleneck," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2021, pp. 350–355.
- [29] H. Zhao and A. Bazco-Nogueras and P. Elia, "Wireless coded caching can overcome the worst-user bottleneck by exploiting finite file sizes," *IEEE Trans. Wireless Commun.*, vol. 21, no. 7, pp. 5450–5466, Jul. 2022.
- [30] M. NaseriTehrani, M. Salehi, and A. Tölli, "Multicast transmission design with enhanced DOF for MIMO coded caching systems," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2024.
- [31] A. M. Girgis, O. Ercetin, M. Nafie, and T. ElBatt, "Coded caching and spatial multiplexing gains in MIMO interference networks," in *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC)*, Apr. 2019, pp. 1–6.
- [32] Y. Cao and M. Tao, "Treating content delivery in multi-antenna coded caching as general message sets transmission: A DoF region perspective," *IEEE Trans. Wireless Commun.*, vol. 18, no. 6, pp. 3129–3141, 2019.
- [33] M. N. Tehrani, M. Salehi, and A. Tölli, "Achievable DoF bounds for cache-aided asymmetric MIMO communications," *arXiv preprint arXiv:2501.10854*, Jan. 2025. [Online]. Available: <https://arxiv.org/abs/2501.10854>
- [34] M. N. Tehrani, M. Salehi, and A. Tölli, "Cache-aided MIMO communications: DoF analysis and transmitter optimization," *arXiv preprint arXiv:2407.15743*, Jul. 2024. [Online]. Available: <https://arxiv.org/abs/2407.15743>

[35] M. N. Tehrani, M. J. Salehi, and A. Tölli, “Enhanced achievable DoF bounds for cache-aided MIMO communication systems,” in *Proc. IEEE 25th Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, 2024, pp. 61–65.

[36] M. Sadeghi *et al.*, “Max-min fair transmit precoding for multi-group multicasting in massive MIMO,” *IEEE Trans. Wireless Commun.*, vol. 17, no. 2, pp. 1358–1373, Feb. 2018.

[37] Q. Spencer, A. Swindlehurst, and M. Haardt, “Zero-forcing methods for downlink spatial multiplexing in multiuser MIMO channels,” *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 461–471, Feb. 2004.

[38] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, “Multiuser MIMO achievable rates with downlink training and channel state feedback,” *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2845–2866, Jun. 2010.

[39] W. Ni and X. Dong, “Hybrid block diagonalization for massive multiuser MIMO systems,” *IEEE Trans. Commun.*, vol. 64, no. 1, pp. 201–211, Jan. 2016.

[40] S. Yang, H. Liu, Y. Zhou, Z. Ma, and P. Fan, “Joint precoding based on interference alignment and SVD for MU-MIMO OTFS with MRC detector,” *IEEE Wireless Commun. Lett.*, vol. 13, no. 10, pp. 2717–2721, Oct. 2024.

[41] “Cisco visual networking index: Global mobile data traffic forecast update, 2017–2022,” White Paper, Cisco, San Jose, CA, USA, Feb. 2019.

[42] H. Zhao, “High performance cache-aided downlink systems: Novel algorithms and analysis,” Ph.D. dissertation, Dept. Comput. Sci. Telecommun. Electron. Paris, Sorbonne University, Paris, 2022.

[43] A. Goldsmith, *Wireless Communications*. Cambridge Uni. Press, 2005.

[44] F. Rusek *et al.*, “Scaling up MIMO: Opportunities and challenges with very large arrays,” *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[45] G. Caire and S. Shamai, “On the achievable throughput of a multi-antenna Gaussian broadcast channel,” *IEEE Trans. Inf. Theory*, vol. 43, no. 7, pp. 1691–1706, Jul. 2003.

[46] “Study on channel model for frequencies from 0.5 to 100 GHz,” 3GPP, Tech. Rep. 38.901, version 16.1.0, Release 16, Dec. 2019.

[47] Y. S. Cho, J. Kim, W. Y. Yang, and C. G. Kang, *MIMO-OFDM Wireless Communications with MATLAB*. Singapore: John Wiley & Sons (Asia) Pte Ltd, 2010, IEEE Press.

[48] E. Björnson, M. Kountouris, and M. Debbah, “Massive MIMO and small cells: Improving energy efficiency by optimal soft-cell coordination,” in *Proc. Int. Conf. Telecommun. (ICT)*, May 2013, pp. 1–5.

[49] H. Q. Ngo, E. G. Larsson, and T. L. Marzetta, “Energy and spectral efficiency of very large multiuser MIMO systems,” *IEEE Trans. Commun.*, vol. 61, no. 4, pp. 1436–1449, Apr. 2013.

[50] R. Couillet and M. Debbah, *Random Matrix Methods for Wireless Communications*. Cambridge Uni. Press, 2011.

[51] K. K. Wong and Z. Pan, “Array gain and diversity order of multiuser MISO antenna systems,” *Int. J. Wireless Inf. Netw.*, vol. 15, no. 2, pp. 82–89, May 2008.