

Correlation-Aware Greedy User Grouping for Zero-Forcing Precoding in Massive MIMO Downlink Systems

Yi Hu

Xizang University

Yan Feng

fy4528@163.com

Xizang University

Fu Xing Wang

Xizang University

Research Article

Keywords: Massive MIMO, zero-forcing precoding, user grouping, channel correlation, matrix conditioning, greedy optimization

Posted Date: July 10th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9847503/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Wireless Personal Communications on August 22nd, 2026. See the published version at <https://doi.org/10.1007/s11277-026-12207-w>.

Correlation-Aware Greedy User Grouping for Zero-Forcing Precoding in Massive MIMO Downlink Systems

Yi Hu, Yan Feng*, Fu Xing Wang

College of Information Science and Technology, Xizang University, Lhasa 850000, China;

hy610400@163.com (Y.H.); 17289117381@163.com (F.W.)

Supported by Science and Technology Projects of Xizang Autonomous Region, China,

XZ202601ZR0038

* Correspondence: fy4528@163.com

Article type: Original Research Article

Abstract

Zero-forcing (ZF) precoding is attractive for multi-user massive multiple-input multiple-output downlink systems because of its low complexity, but its performance can degrade when users assigned to the same transmission group have highly correlated channel vectors. This paper studies a correlation-aware user grouping problem under a spatially correlated channel model and proposes CorrGreedy, a lightweight swap-based refinement method for improving group conditioning before group-wise ZF precoding. Unlike joint scheduling–precoding optimization, the proposed method is positioned as a preprocessing heuristic that uses normalized pairwise channel correlation as a surrogate objective. The relationship between the surrogate objective, the group Gram matrix, eigenvalue bounds, and ZF noise enhancement is analyzed to clarify when correlation reduction is expected to be beneficial and when the approximation may be limited. Simulation results under controlled antenna, user-load, signal-to-noise-ratio, and spatial-correlation settings compare CorrGreedy with ungroupped ZF, MMSE precoding, K-means grouping, and static correlation-based grouping. The results show that CorrGreedy improves the conditioning of grouped channel matrices and achieves competitive spectral-efficiency, bit-error-rate, capacity, and energy-efficiency performance among the considered ZF-based grouping schemes, with moderate additional computational cost. Reproducibility details, including initialization, stopping criteria, random seeds, and Monte Carlo settings, are specified to support independent verification.

Keywords: Massive MIMO; zero-forcing precoding; user grouping; channel correlation; matrix conditioning; greedy optimization

1 Introduction

Massive multiple-input multiple-output (Massive MIMO) has been widely studied as a key physical-layer technology for modern wireless communication systems because it can improve spectral efficiency, energy efficiency, and spatial multiplexing capability by using a large number of base-station antennas [1]–[3]. In the downlink of a multiuser Massive MIMO system, the base station can serve several users on the same time-frequency resource through linear precoding. Among commonly used linear precoding methods, zero-forcing (ZF) precoding is attractive because it has a clear mathematical structure and can suppress intra-group multiuser interference when the scheduled user channel matrix is full-rank and well conditioned [4], [5].

However, the effectiveness of ZF precoding depends strongly on the spatial relationship among user channel vectors. When users with highly correlated channel directions are placed in the same transmission group, the corresponding Gram matrix may become ill-conditioned. In this case, ZF precoding requires a larger power normalization factor, which reduces the effective received signal-to-noise ratio and may degrade spectral efficiency and bit-error-rate performance. This problem is more evident in spatially correlated channels, where users located in similar angular regions or experiencing similar scattering conditions may have non-negligible channel correlation [7], [8].

User selection and scheduling have therefore become important components of ZF-based multiuser transmission. Classical studies on ZF beamforming and semi-orthogonal user selection have shown that selecting users with sufficiently separated channel directions can improve the downlink sum-rate behavior [5], [6], [13], [14]. These studies provide the theoretical basis for correlation-aware user organization. Nevertheless, user selection and fixed-size user grouping are not the same problem. User selection usually chooses a subset of users from a larger candidate pool, while fixed-size grouping requires all users to be assigned into groups under a predetermined group-size constraint. This makes the problem closer to constrained combinatorial assignment than to pure subset selection.

Recent studies have further examined user grouping, scheduling, and resource allocation in Massive MIMO and related systems [17]–[19], [29]. These works confirm that user organization remains relevant even when linear precoding is already specified. However, many existing approaches address broader system-level objectives, such as learning-based resource allocation, limited-feedback scheduling, antenna selection, proportional fairness, or cell-free cooperation. In contrast, this paper focuses on a narrower and more controlled problem: how to refine fixed-size user groups before applying group-wise ZF precoding in a single-cell Massive MIMO downlink.

This paper considers a single-cell downlink Massive MIMO system with spatially correlated Rayleigh fading. The base station is equipped with multiple antennas and serves single-antenna users in fixed-size groups. For each group, ZF precoding is applied independently. The main assumption is that reducing the accumulated intra-group channel correlation can improve the conditioning tendency of the group Gram matrix and thereby make group-wise ZF precoding more effective. Based on this assumption,

a lightweight greedy grouping method, named CorrGreedy, is proposed. The method starts from a feasible grouping and iteratively performs beneficial inter-group user swaps to reduce a correlation-based grouping cost.

The proposed method is not intended to replace robust precoding methods such as regularized ZF, nor does it claim to solve the global sum-rate maximization problem. Instead, CorrGreedy is designed as a preprocessing method for group-wise ZF transmission. Its objective is to construct more ZF-friendly user groups with limited computational overhead and without repeatedly computing ZF precoders during the grouping search.

The main contributions of this paper are as follows. First, a correlation-aware grouping criterion is introduced for fixed-size group-wise ZF transmission. Second, a greedy inter-group swap algorithm is developed to reduce intra-group channel correlation while preserving the group-size constraint. Third, the computational structure of the proposed method is analyzed to support reproducibility and fair comparison. Fourth, numerical simulations are conducted under a spatially correlated Rayleigh channel model, using spectral efficiency, BER, Gram-matrix condition number, minimum eigenvalue, correlation cost, accepted swap count, operation count, and energy efficiency as evaluation metrics.

The remainder of this paper is organized as follows. Section 2 reviews related work on Massive MIMO precoding, ZF-oriented scheduling, spatial correlation, and clustering-based grouping. Section 3 presents the system model, channel model, and fixed-size grouping problem. Section 4 describes the CorrGreedy algorithm and its complexity. Section 5 gives the theoretical interpretation of the correlation-based objective. Section 6 presents the simulation setup and baseline schemes. Section 7 discusses numerical results, and Section 8 concludes the paper.

2 Related Work

This section reviews studies related to linear precoding, ZF-oriented user selection, spatially correlated channel modeling, clustering-based grouping, and scalable user organization. The objective is not to provide a broad survey of all Massive MIMO techniques, but to clarify the technical position of the proposed CorrGreedy method and to justify the baseline design used in the numerical evaluation.

2.1 Linear Precoding and ZF-Based Transmission

Linear precoding is widely used in downlink multiuser MIMO and Massive MIMO systems because it provides a practical balance between interference suppression and computational complexity. Zero-forcing (ZF) precoding is one of the representative linear schemes. It suppresses intra-group multiuser interference by constructing precoding vectors from the inverse or pseudo-inverse of the user channel matrix [4]. When the scheduled users have linearly independent and well-separated channel vectors, ZF can provide an effective solution for simultaneous downlink transmission.

However, ZF performance is sensitive to the conditioning of the group channel matrix. If the users assigned to the same group have strongly correlated channel directions, the Gram matrix may become ill-

conditioned. In this case, the ZF precoder requires a larger normalization factor, which reduces the effective received signal-to-noise ratio and may degrade spectral efficiency and bit-error-rate performance. This property motivates the use of user selection, scheduling, or grouping before ZF precoding.

Regularized zero-forcing (RZF) and related linear precoding methods have been studied to improve robustness under noise, spatial correlation, and imperfect propagation conditions [7], [9], [11]. Recent works have also considered correlated Rician channels, mutual coupling, and reduced-channel ZF structures in Massive MIMO systems [9]–[11]. These studies show that the channel model and matrix-conditioning behavior directly affect precoding performance. The present work does not propose a new precoder and does not claim to replace RZF. Instead, it studies whether a correlation-aware fixed-size grouping stage can improve the operating condition of group-wise ZF precoding.

2.2 User Selection, Scheduling, and Fixed-Size Grouping

User selection has been extensively studied for ZF-based multiuser transmission. Greedy user selection and semi-orthogonal user selection methods showed that selecting users with sufficiently separated channel directions can improve the sum-rate performance of ZF beamforming [5], [6], [13]. Related asymptotic analysis also indicated that ZF beamforming with proper user selection can approach favorable sum-rate behavior when the user pool is large [14].

Although these works provide the theoretical basis for correlation-aware user organization, the fixed-size grouping problem considered in this paper is different from conventional user selection. In user selection, the scheduler usually chooses a subset of users from a larger candidate set, while unselected users may be postponed. In fixed-size grouping, all users must be assigned to groups, and each group must satisfy a predetermined size constraint. Therefore, the problem is closer to constrained combinatorial assignment than to pure subset selection.

Recent studies have extended scheduling and grouping to more complex Massive MIMO scenarios. Dynamic user selection for limited-feedback MU-MIMO, proportional-fairness scheduling with antenna selection, and cluster-centric scheduling in cell-free systems have been investigated in [19], [29], and [30]. These studies confirm that user organization remains important even when the underlying precoding method is already specified. CorrGreedy follows the same general motivation, but it focuses on a narrower problem: reducing intra-group channel correlation under a fixed group-size constraint before applying group-wise ZF.

2.3 Spatial Correlation and Channel-Model Considerations

Spatial correlation is an important factor in Massive MIMO channel modeling. Under ideal independent Rayleigh fading, user channel vectors tend to become nearly orthogonal as the number of base-station antennas increases. In practical environments, however, limited angular spread, similar user locations, common scattering clusters, antenna correlation, and mutual coupling may produce non-

negligible correlation among user channels. As a result, the favorable propagation assumption may be weakened.

Large-system analysis of correlated MISO broadcast channels has shown that spatial correlation affects the performance of linear precoding [7]. More recent work has studied Massive MIMO downlink transmission by exploiting spatial correlation in training design [12]. Mutual coupling has also been shown to influence downlink spectral efficiency in Massive MIMO systems [10]. These results support the need to include a clear spatially correlated channel model rather than relying only on independent Rayleigh fading.

For group-wise ZF, the normalized Gram matrix provides a direct way to describe the relationship among user channel vectors. Small off-diagonal entries indicate weak pairwise correlation, while large off-diagonal entries indicate that some users have similar channel directions. A group with large intra-group correlation may have a smaller minimum eigenvalue and a larger condition number, which is unfavorable for ZF. Based on this observation, the proposed method uses an accumulated pairwise correlation cost as the grouping objective. This objective is simpler than direct sum-rate optimization and avoids repeated ZF inversion during the grouping search.

2.4 Clustering-Based Baselines and Comparison Logic

Clustering methods are natural baselines for user grouping. K-means clustering and K-means++ initialization are classical methods for organizing samples according to feature similarity [15], [16]. In wireless systems, clustering has also been used in user-centric resource allocation, cell-free Massive MIMO, and pilot-assignment design [23], [25], [27].

However, conventional clustering and ZF-friendly grouping do not necessarily have the same objective. K-means tends to place similar samples in the same cluster. If channel vectors or channel-related features are used directly, this may group users with similar spatial signatures together. For ZF precoding, this is not always desirable, because users with similar channel directions can produce a poorly conditioned group channel matrix. Therefore, K-means is used in this paper as a general clustering baseline rather than as a ZF-optimized grouping method.

The baseline design should reflect different levels of channel awareness. Random grouping measures the average behavior without channel-aware refinement. Static correlation-based grouping tests whether a simple one-pass correlation rule is useful. K-means grouping tests whether general clustering can provide competitive groups. RZF, when included as a precoding baseline, tests whether regularization can compensate for unfavorable channel correlation at the precoder level. CorrGreedy evaluates a different mechanism: iterative inter-group swaps that reduce intra-group correlation while preserving fixed group sizes.

2.5 Cell-Free and Scalable User Organization

Cell-free Massive MIMO and user-centric clustering have received substantial attention because they provide scalable ways to serve distributed users [20], [24], [25], [28]. Local partial ZF, ZF-based energy-

efficiency optimization, downlink resource allocation, enhanced subset greedy scheduling, graph-coloring-based pilot assignment, and multi-level cooperative architectures have been studied in [21]–[23], [26], [27], [31], and [32]. Recent works have also considered cluster-centric scheduling and dynamic user clustering in cell-free Massive MIMO systems [18], [30].

These studies are related because they show the broader importance of scalable user organization. Nevertheless, the present paper does not study a cell-free architecture. It deliberately adopts a single-cell downlink Massive MIMO model to isolate the influence of intra-group channel correlation on group-wise ZF. Including access-point selection, fronthaul constraints, pilot contamination, distributed channel estimation, and inter-cell cooperation would make the evaluation broader but less focused. Therefore, cell-free studies are treated as contextual references, while the core problem remains fixed-size correlation-aware grouping for single-cell group-wise ZF precoding.

In summary, existing literature supports three points. First, ZF performance depends strongly on channel orthogonality and Gram-matrix conditioning [4]–[7], [13], [14]. Second, spatial correlation is a realistic and performance-relevant channel feature [7], [10], [12]. Third, user grouping and scheduling remain important in both conventional and cell-free Massive MIMO systems [18], [19], [26], [29], [30]. The proposed CorrGreedy method builds on these observations and studies a focused problem: low-complexity fixed-size grouping for correlation-aware group-wise ZF transmission.

3 System Model

3.1 System Overview




Figure 1. Correlation-Aware User Grouping and Precoding Framework for Massive MIMO Downlink

This paper considers a single-cell massive multiple-input multiple-output (Massive MIMO) downlink system, as illustrated in Figure 1. Presents the considered downlink Massive MIMO system, the spatially correlated channel model, the group-wise zero-forcing transmission model, the performance metrics, and the user-grouping problem formulation. The objective of this section is to define a reproducible and mathematically explicit framework under which the proposed CorrGreedy algorithm can

be evaluated. The model is intentionally kept controlled so that the effect of channel correlation on group-wise ZF precoding can be isolated and interpreted.

3.1 Downlink Massive MIMO System

Consider a single-cell downlink Massive MIMO system where one base station is equipped with M antennas and serves K single-antenna user equipments. The set of users is denoted by $\mathcal{K} = \{1, 2, \dots, K\}$.

The system operates on a narrowband flat-fading resource block. This assumption is commonly used when the focus is on the spatial-domain behavior of user grouping and linear precoding rather than on frequency-selective scheduling. The downlink channel between the base station and user k is represented by $\mathbf{h}_k \in \mathbb{C}^{M \times 1}$, where $\mathbf{h}_k$ contains the complex channel coefficients from the M transmit antennas to user k . By stacking the user channels, the full channel matrix is written as

$$\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_K]^H \in \mathbb{C}^{K \times M} \quad (1)$$

The system serves users in groups. Let G denote the number of groups, and let S denote the number of users in each group. For simplicity and reproducibility, fixed-size groups are used, so that $K=GS$.

The g -th user group is denoted by $G_g \subset K, g=1, 2, \dots, G$, with $|G_g| = S, G_g \cap G_{g'} = \emptyset$ for $g \neq g', \cup_{g=1}^G G_g = \mathcal{K}$.

For group-wise ZF precoding to be feasible, the group size is assumed to satisfy $S \leq M$.

In the simulations, the typical Massive MIMO condition $M > S$ is adopted, because ZF becomes numerically more reliable when the number of transmit antennas is larger than the number of simultaneously served users. This setting is also consistent with the motivation of Massive MIMO, where a large antenna array provides additional spatial degrees of freedom for multi-user interference suppression.

3.2 Spatially Correlated Rayleigh Channel Model

The channel model is a central part of this study because the proposed grouping method is designed to reduce harmful intra-group channel correlation. For user k , the channel vector is modeled as

$$\mathbf{h}_k = \sqrt{\beta_k} \mathbf{R}_k^{1/2} \mathbf{g}_k \quad (2)$$

where β_k is the large-scale fading coefficient, $\mathbf{R}_k \in \mathbb{C}^{M \times M}$ is the transmit-side spatial correlation matrix, and

$$\mathbf{g}_k \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_M) \quad (3)$$

is the small-scale Rayleigh fading vector. In this paper, the main simulations use normalized large-scale fading, $\beta_k = 1, \forall k$, unless otherwise stated. This normalization does not imply that path loss is unimportant in practical systems. Instead, it is used to isolate the impact of spatial correlation and user grouping on ZF precoding. If non-uniform β_k values were included, the resulting performance differences would reflect both channel-strength imbalance and correlation effects. Since the purpose of this paper is

to evaluate correlation-aware grouping, normalized large-scale fading provides a cleaner and more interpretable baseline.

The spatial correlation matrix is generated according to an exponential correlation model. Its (m,n) -th element is defined as

$$[\mathbf{R}_k]_{m,n} = \rho_k^{|m-n|}, m, n = 1, 2, \dots, M \quad (4)$$

where $\rho_k \in [0,1)$ is the spatial-correlation coefficient of user k . A larger ρ_k corresponds to stronger correlation among the antenna-domain channel coefficients. When $\rho_k = 0$, the model reduces to an independent Rayleigh fading channel. When ρ_k approaches one, the channel becomes more spatially correlated, and the effective rank of $\mathbf{R}_k$ decreases.

For controlled simulations, a common correlation coefficient is used: $\rho_k = \rho, \forall k$.

The parameter ρ is varied in the simulation section to evaluate channel-correlation sensitivity. Typical values include low, moderate, and strong correlation cases, such as $\rho \in \{0, 0.3, 0.5, 0.7, 0.9\}$.

This range allows the evaluation to cover nearly independent fading, practically moderate correlation, and severe correlation. The strong-correlation cases are not meant to represent all deployment scenarios. They are included to test whether the proposed grouping method remains stable when ZF is numerically stressed by highly similar channel directions.

The average channel energy is normalized to avoid unintended transmit-power scaling. After generating $\mathbf{h}_k$, the channel may be normalized as

$$\tilde{\mathbf{h}}_k = \frac{\sqrt{M}\mathbf{h}_k}{\|\mathbf{h}_k\|_2} \quad (5)$$

so that

$$\|\tilde{\mathbf{h}}_k\|_2^2 = M \quad (6)$$

In the following, $\mathbf{h}_k$ denotes the normalized channel vector unless explicitly stated otherwise. This normalization ensures that the comparison among different grouping methods is primarily determined by channel direction and correlation structure, rather than by random channel norm differences.

3.3 Group-Wise Downlink ZF Precoding

For the g -th group, the corresponding group channel matrix is denoted by

$$\mathbf{H}_g = [\mathbf{h}_{k_1}, \mathbf{h}_{k_2}, \dots, \mathbf{h}_{k_S}]^H \in \mathbb{C}^{S \times M} \quad (7)$$

Where $\mathcal{G}_g = \{k_1, k_2, \dots, k_S\}$.

The downlink transmitted signal for group g is

$$\mathbf{x}_g = \sqrt{P_g} \mathbf{W}_g \mathbf{s}_g \quad (8)$$

where P_g is the transmit power assigned to group g , $\mathbf{s}_g \in \mathbb{C}^{S \times 1}$ is the data-symbol vector with $\mathbb{E}[\mathbf{s}_g \mathbf{s}_g^H] = \mathbf{I}_S$, and $\mathbf{W}_g \in \mathbb{C}^{M \times S}$ is the precoding matrix.

The unnormalized ZF precoder is given by

$$\mathbf{F}_g = \mathbf{H}_g^H (\mathbf{H}_g \mathbf{H}_g^H)^{-1} \quad (9)$$

If $\mathbf{H}_g \mathbf{H}_g^H$ is full rank, then $\mathbf{H}_g \mathbf{F}_g = \mathbf{I}_S$.

This relation shows that ZF can remove intra-group interference under perfect channel state information. However, the normalization of $\mathbf{F}_g$ is important. When the Gram matrix $\mathbf{H}_g \mathbf{H}_g^H$ is ill conditioned, the norm of $\mathbf{F}_g$ becomes large, causing power loss after normalization.

The normalized ZF precoder is defined as

$$\mathbf{W}_g = \frac{\mathbf{F}_g}{\sqrt{\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)}} \quad (10)$$

This normalization satisfies $\text{tr}(\mathbf{W}_g \mathbf{W}_g^H) = 1$.

The received signal at users in group g is

$$\mathbf{y}_g = \sqrt{P_g} \mathbf{H}_g \mathbf{W}_g \mathbf{s}_g + \mathbf{n}_g \quad (11)$$

Where $\mathbf{n}_g \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_S)$ is the additive white Gaussian noise vector. Because of the ZF structure,

$$\mathbf{H}_g \mathbf{W}_g = \frac{\mathbf{I}_S}{\sqrt{\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)}} \quad (12)$$

Thus, the effective received signal for user i in group g becomes

$$y_{g,i} = \frac{\sqrt{P_g}}{\sqrt{\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)}} s_{g,i} + n_{g,i} \quad (13)$$

The corresponding effective signal-to-noise ratio is

$$\gamma_{g,i} = \frac{P_g}{\sigma^2 \text{tr}(\mathbf{F}_g \mathbf{F}_g^H)} \quad (14)$$

Under equal power assignment across groups, $P_g = P$. Since all users in the same group share the same ZF normalization term in this formulation, the group spectral efficiency is

$$R_g = \sum_{i=1}^S \log_2 (1 + \gamma_{g,i}) = S \log_2 \left(1 + \frac{P}{\sigma^2 \text{tr}(\mathbf{F}_g \mathbf{F}_g^H)} \right) \quad (15)$$

The total spectral efficiency over all groups is $R_{\text{sum}} = \sum_{g=1}^G R_g$.

This formulation directly shows why the grouping problem is important. Even if ZF eliminates intra-group interference, the achievable rate depends on the normalization term $\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)$.

A poorly conditioned group can increase this term and reduce the effective SNR.

3.4 Correlation and Gram-Matrix Conditioning

To quantify the channel-direction similarity between two users i and j , the normalized pairwise channel correlation is defined as

$$c_{i,j} = \frac{|\mathbf{h}_i^H \mathbf{h}_j|^2}{\|\mathbf{h}_i\|_2^2 \|\mathbf{h}_j\|_2^2} \quad (16)$$

The value of $c_{i,j}$ satisfies $0 \leq c_{i,j} \leq 1$.

When $c_{i,j} = 0$, the two channel vectors are orthogonal. When $c_{i,j}$ is close to one, the two channel directions are highly similar. In ZF precoding, such similarity is undesirable because it makes the group channel matrix closer to rank deficient.

For group g , the intra-group correlation cost is defined as

$$C(\mathcal{G}_g) = \sum c_{i,j} \quad (17)$$

The total grouping cost is

$$C_{\text{total}} = \sum_{g=1}^G C(\mathcal{G}_g) \quad (18)$$

This cost is used as a surrogate objective in the proposed algorithm. It does not exactly equal the inverse Gram-matrix trace or the negative spectral efficiency. However, it is numerically meaningful because the off-diagonal entries of the normalized Gram matrix are determined by pairwise channel inner products. A group with lower pairwise channel correlation is more likely to have a better-conditioned Gram matrix, which usually reduces the ZF normalization penalty.

The condition number of the group Gram matrix is also evaluated as a diagnostic metric:

$$\kappa_g = \frac{\lambda_{\min}(\mathbf{H}_g \mathbf{H}_g^H)}{\lambda_{\max}(\mathbf{H}_g \mathbf{H}_g^H)} \quad (19)$$

where $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ are the maximum and minimum eigenvalues, respectively. A smaller condition number generally indicates a more stable matrix inversion. In the simulation results, both C_{total} and κ_g are reported. This is important because the proposed algorithm directly optimizes C_{total} , while κ_g and spectral efficiency are used to test whether the surrogate objective is practically useful.

3.5 Baseline Grouping and Precoding Schemes

A meaningful evaluation requires baselines that represent different levels of grouping information and computational effort. This paper considers the following baselines.

The first baseline is ungrouped ZF, where all users are jointly precoded when feasible. This baseline reflects the idealized case in which the base station attempts to serve all users simultaneously. It is useful for comparison, but it may be numerically unstable or infeasible when $K > M$.

The second baseline is random grouping. Users are randomly partitioned into G fixed-size groups. This baseline provides a lower-complexity reference and shows the expected behavior when no channel-structure information is used.

The third baseline is K-means-based grouping. The normalized channel vectors or their low-dimensional representations are clustered using a standard clustering rule. This baseline reflects a general clustering strategy, but it does not directly minimize the intra-group correlation objective used for ZF conditioning.

The fourth baseline is static correlation-based grouping. In this method, users are arranged according to a simple correlation-aware heuristic, such as distributing highly correlated users into different groups without iterative swap refinement. This baseline helps isolate the additional benefit of the proposed greedy refinement step.

The fifth reference is regularized ZF or MMSE-type linear precoding, when included. This reference is not a grouping baseline in the strict sense, because it changes the precoder rather than the grouping rule. It is included to clarify whether grouping improvement and precoder regularization address different parts of the performance loss.

These baselines are selected to avoid an unfair comparison against only weak methods. Random grouping tests the value of using channel information. K-means grouping tests the value of general clustering. Static correlation grouping tests the value of non-iterative correlation awareness. MMSE or regularized ZF tests whether improving group composition remains relevant when the precoder itself is more robust than conventional ZF.

3.6 Problem Formulation

The ideal objective of group-wise downlink user grouping is to maximize the total spectral efficiency:

$$\max_{\{\mathcal{G}_1, \dots, \mathcal{G}_G\}} R_{\text{sum}} = \sum_{g=1}^G S \log_2 \left(1 + \frac{P}{\sigma^2 \text{tr}(\mathbf{F}_g \mathbf{F}_g^H)} \right) \quad (20)$$

subject to

$$|\mathcal{G}_g| = S, g = 1, 2, \dots, G, \mathcal{G}_g \cap \mathcal{G}_{g'} = \emptyset, g \neq g', \bigcup_{g=1}^G \mathcal{G}_g = \mathcal{K}. \quad (21)$$

This problem is combinatorial because it requires searching over all feasible partitions of K users into G groups of size S . The number of possible partitions grows rapidly with K , making exhaustive search impractical for moderate or large user sets. Moreover, direct spectral-efficiency optimization requires repeated computation of ZF precoders and matrix inversions for candidate groupings, which further increases computational cost.

Therefore, this paper uses a correlation-based surrogate problem:

$$\min_{\{\mathcal{G}_1, \dots, \mathcal{G}_G\}} C_{\text{total}} = \sum_{g=1}^G \sum \frac{|\mathbf{h}_i^H \mathbf{h}_j|^2}{\|\mathbf{h}_i\|_2^2 \|\mathbf{h}_j\|_2^2}, \quad (22)$$

subject to the same fixed-size partition constraints. This surrogate objective is not claimed to be equivalent to the original sum-rate maximization problem. Its role is to provide a low-complexity and interpretable grouping criterion that is closely related to the off-diagonal structure of the normalized Gram matrix. The proposed CorrGreedy algorithm in the next section is designed to reduce this objective through beneficial inter-group user swaps.

The formulation also defines the boundary of the present study. The paper does not jointly optimize power allocation, pilot assignment, beam selection, modulation and coding, hybrid precoding, or inter-cell interference coordination. These are important topics but are outside the scope of this work. By focusing on fixed-size grouping and group-wise ZF under spatially correlated Rayleigh fading, the study provides a controlled evaluation of whether correlation-aware grouping can improve the numerical and transmission behavior of ZF precoding.

4 Proposed Correlation-Aware Greedy User Grouping

This section presents the proposed CorrGreedy algorithm for correlation-aware user grouping in the considered downlink Massive MIMO system. The algorithm is designed as a preprocessing step before group-wise ZF precoding. Its main purpose is to reduce the intra-group channel-correlation cost so that the grouped channel matrices become more suitable for ZF transmission. The proposed method does not claim to obtain the globally optimal grouping. Instead, it provides a simple and reproducible greedy refinement procedure with explicit numerical criteria.

4.1 Design Principle

As discussed in Section 3, the performance of group-wise ZF precoding is strongly affected by the conditioning of the group channel matrix. For the g -th group, the unnormalized ZF precoder is given by

$$\mathbf{F}_g = \mathbf{H}_g^H (\mathbf{H}_g \mathbf{H}_g^H)^{-1} \quad (23)$$

When the users in $\mathcal{G}_g$ have similar channel directions, the Gram matrix $\mathbf{H}_g \mathbf{H}_g^H$ may become ill conditioned. In this case, although ZF can still suppress intra-group interference under perfect CSI, the normalization term $\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)$

can become large. This reduces the effective SNR and degrades the achievable spectral efficiency.

A direct optimization of the grouping structure would require an exhaustive search over all fixed-size user partitions. The number of possible partitions of K users into G groups with S users per group is

$$N_{\text{part}} = \frac{K!}{(S!)^G G!} \quad (24)$$

This number increases rapidly with K . Therefore, exhaustive grouping is not practical for moderate or large user sets. In addition, evaluating each candidate partition requires repeated construction of group channel matrices, matrix inversion, and rate calculation. To avoid this computational burden, CorrGreedy uses a correlation-based surrogate objective instead of direct sum-rate maximization.

The basic idea is straightforward. If two users have highly correlated channel directions, placing them in the same group tends to increase the off-diagonal entries of the normalized Gram matrix. CorrGreedy therefore attempts to distribute such users across different groups by iteratively swapping users between groups. A swap is accepted only when it reduces the total intra-group channel-correlation cost.

Figure 2 illustrates the overall workflow of the CorrGreedy-based group-wise ZF transmission framework. The process starts from channel generation and channel normalization, followed by pairwise correlation calculation, initial grouping, greedy swap refinement, and final group-wise ZF precoding evaluation.




Figure 2. Overall workflow of the CorrGreedy-based group-wise ZF transmission framework.

4.2 Pairwise Channel-Correlation Matrix

The first step of CorrGreedy is to compute the normalized pairwise channel-correlation matrix $\mathbf{C} \in \mathbb{R}^{K \times K}$.

For users i and j , the correlation coefficient is defined as

$$C_{i,j} = \frac{|\mathbf{h}_i^H \mathbf{h}_j|^2}{\|\mathbf{h}_i\|_2^2 \|\mathbf{h}_j\|_2^2}, i, j \in \mathcal{K}. \quad (25)$$

The diagonal entries are not used in the grouping cost and are set as $C_{i,i} = 0$.

The value of $C_{i,j}$ satisfies $0 \leq C_{i,j} \leq 1$.

A smaller value means that the two channel vectors are more orthogonal, while a larger value indicates stronger channel-direction similarity. Since $C_{i,j} = C_{j,i}$, the matrix $\mathbf{C}$ is symmetric.

The pairwise correlation matrix is computed only once within each channel realization. Its computational complexity is approximately $\mathcal{O}(K^2M)$, because the inner product between each pair of M -dimensional channel vectors must be evaluated. After $\mathbf{C}$ is available, the following grouping refinement can be performed without repeated antenna-domain inner-product calculations.

4.3 Intra-Group Correlation Cost

For a given group $\mathcal{G}_g$, the intra-group correlation cost is defined as $C(\mathcal{G}_g) = \sum C_{i,j}$.

The total grouping cost is

$$C_{\text{total}} = \sum_{g=1}^G C(\mathcal{G}_g) \quad (26)$$

CorrGreedy aims to reduce C_{total} under the fixed-size grouping constraints $|\mathcal{G}_g| = S, g = 1, 2, \dots, G, \mathcal{G}_g \cap \mathcal{G}_{g'} = \emptyset, g \neq g', \cup_{g=1}^G \mathcal{G}_g = \mathcal{K}$.

Therefore, the surrogate grouping problem is written as $\min_{\{\mathcal{G}_1, \dots, \mathcal{G}_G\}} C_{\text{total}}$.

This surrogate objective is not mathematically equivalent to maximizing the exact spectral efficiency. However, it is directly related to the off-diagonal structure of the normalized Gram matrix. A group with lower pairwise channel correlation is generally more favorable for ZF, because the corresponding Gram matrix is less likely to be close to rank deficient.

4.4 Initialization

CorrGreedy starts from an initial feasible grouping. In this paper, two initialization methods are considered to make the comparison more complete.

The first method is random initialization. The K users are randomly permuted and divided into G groups with S users in each group. This method does not use channel information and therefore serves as a basic reference.

The second method is static correlation-based initialization. For user i , an aggregate correlation score is calculated as

$$q_i = \sum_{j=1}^K C_{i,j}. \quad (27)$$

Users with larger q_i values have stronger average correlation with other users. These users are sorted in descending order of q_i , giving the ordered list

$$\pi = \text{argsort}(q_i, \text{descending}). \quad (28)$$

Then user $\pi(t)$ is assigned to group $g(t) = 1 + \text{mod}(t-1, G)$.

This round-robin assignment distributes highly correlated users across different groups at the initial stage. However, it is still a static rule and does not evaluate whether later pairwise exchanges can further reduce the objective. CorrGreedy adds this iterative refinement stage.

For reproducible simulations, the random seed must be recorded when random initialization is used. In addition, when comparing static grouping and CorrGreedy refinement, the same initial grouping should be used before the refinement step. This ensures that the measured performance improvement is caused by the greedy swap procedure rather than by a different starting point.

4.5 Greedy Swap Refinement

The core of CorrGreedy is inter-group user swapping. Consider two different groups $\mathcal{G}_a$ and $\mathcal{G}_b$, where $a \neq b$. Let $u \in \mathcal{G}_a, v \in \mathcal{G}_b$.

A candidate swap exchanges users u and v , producing

$$\mathcal{G}_a' = (\mathcal{G}_a \setminus \{u\}) \cup \{v\}, \mathcal{G}_b' = (\mathcal{G}_b \setminus \{v\}) \cup \{u\}. \quad (29)$$

The cost before the swap is $C_{\text{before}} = C(\mathcal{G}_a) + C(\mathcal{G}_b)$, and the cost after the swap is

$$C_{\text{after}} = C(\mathcal{G}_a') + C(\mathcal{G}_b'). \quad (30)$$

The swap gain is defined as $\Delta(u, v) = C_{\text{before}} - C_{\text{after}}$.

If $\Delta(u, v) > 0$, the swap reduces the total intra-group correlation cost and is accepted as a beneficial candidate. Otherwise, it is rejected.

For efficient computation, the gain can be calculated locally. The local contribution of user u in group $\mathcal{G}_a$ is $L(u, \mathcal{G}_a) = \sum C_{u,i}$.

Similarly, $L(v, \mathcal{G}_b) = \sum C_{v,j}$.

After swapping, the new local costs become

$$L(v, \mathcal{G}_a \setminus \{u\}) = \sum_{i \in \mathcal{G}_a \setminus \{u\}} C_{v,i}, L(u, \mathcal{G}_b \setminus \{v\}) = \sum_{j \in \mathcal{G}_b \setminus \{v\}} C_{u,j}. \quad (31)$$

Thus, the swap gain can be computed as

$$\Delta(u, v) = L(u, \mathcal{G}_a) + L(v, \mathcal{G}_b) - L(v, \mathcal{G}_a \setminus \{u\}) - L(u, \mathcal{G}_b \setminus \{v\}). \quad (32)$$

This calculation avoids recomputing the cost of unaffected groups.




Figure 3. Flowchart of the proposed CorrGreedy user grouping algorithm. The original flowchart is retained. It illustrates the initialization, objective evaluation, pairwise swap search, acceptance criterion, convergence check, and final grouping output.

4.6 Update Rule and Stopping Criterion

This paper adopts the best-improvement rule as the default update strategy. In each iteration, CorrGreedy scans all feasible inter-group user pairs and selects the swap with the largest positive gain:

$$(u^*, v^*) = \arg \max_{u \in \mathcal{G}_a, v \in \mathcal{G}_b, a \neq b} \Delta(u, v). \quad (33)$$

The selected swap is accepted only when $\Delta(u^*, v^*) > \epsilon$, where ϵ is a small numerical tolerance. In the simulations, the default value is $\epsilon = 10^{-12}$.

This prevents accepting changes caused only by floating-point numerical errors.

After an accepted swap, the total cost is updated as

$$C_{\text{total}}^{(t+1)} = C_{\text{total}}^{(t)} - \Delta(u^*, v^*). \quad (34)$$

Therefore, $C_{\text{total}}^{(t+1)} < C_{\text{total}}^{(t)}$.

The objective is monotonically reduced over accepted iterations. Since the number of feasible fixed-size partitions is finite, the algorithm eventually stops when no beneficial swap exists. In practice, a maximum iteration number I_{max} is also imposed. The default setting is $I_{\text{max}} = 100$.

The actual number of accepted swaps is recorded in the simulation results to evaluate convergence behavior.

4.7 Algorithm Summary

The complete CorrGreedy procedure is summarized below.

Algorithm 1: CorrGreedy User Grouping

Input: channel vectors $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_K\}$, number of groups G , group size S , maximum iteration number I_{max} , *tolerance* ϵ .

Output: refined groups $\{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_G\}$.

Step 1: Normalize each channel vector so that $\|\mathbf{h}_k\|_2^2 = M$.

Step 2: Compute the pairwise correlation matrix $\mathbf{C}$, where $C_{i,j} = \frac{|\mathbf{h}_i^H \mathbf{h}_j|^2}{\|\mathbf{h}_i\|_2^2 \|\mathbf{h}_j\|_2^2}$, $C_{i,i} = 0$.

Step 3: Generate an initial feasible grouping using random initialization or static correlation-based initialization.

Step 4: Compute the initial total cost C_{total} .

Step 5: For $t = 1, 2, \dots, I_{\text{max}}$, scan all inter-group user pairs and compute $\Delta(u, v)$.

Step 6: Select the swap with the largest gain Δ_{best} .

Step 7: If $\Delta_{\text{best}} > \epsilon$, perform the swap and update $C_{\text{total}} = C_{\text{total}} - \Delta_{\text{best}}$.

Otherwise, terminate the algorithm.

Step 8: Return the final grouping.

The algorithm preserves all grouping constraints because each update is a swap between two existing groups. No user is removed, duplicated, or left unassigned.

4.8 Computational Complexity

The complexity of CorrGreedy consists of correlation-matrix computation, initialization, and greedy refinement.

The correlation matrix requires $\mathcal{O}(K^2 M)$ operations. Random initialization requires $\mathcal{O}(K)$.

Static correlation-based initialization require $\mathcal{O}(K^2 + K \log K)$, including aggregate-score calculation and sorting.

In one best-improvement iteration, the number of group pairs is $\binom{G}{2} = \frac{G(G-1)}{2}$.

For each group pair, there are S^2 candidate swaps. Therefore, the number of candidate swaps per iteration is $N_{\text{swap}} = \frac{G(G-1)}{2} S^2$.

Since $K=GS$, this can also be written as $N_{\text{swap}} = \frac{K^2}{2} \left(1 - \frac{1}{G}\right)$.

For example, when $K=32$, $G=8$, and $S=4$, the number of candidate swaps per iteration is $N_{\text{swap}} = \frac{8(8-1)}{2} \times 4^2 = 448$.

This is much smaller than the number of exhaustive fixed-size partitions, $N_{\text{part}} = \frac{32!}{(4!)^8 8!}$.

Each local gain evaluation requires approximately $\mathcal{O}(S)$ operations. Thus, the complexity of one best-improvement iteration is $\mathcal{O}(K^2 S)$.

With I_{max} iterations, the total complexity is approximately $\mathcal{O}(K^2 M + I_{\text{max}} K^2 S)$.

This is substantially lower than exhaustive partition search. Moreover, CorrGreedy does not require repeated ZF matrix inversion during the grouping search. ZF precoding is computed only after the final grouping is obtained.

4.9 Relation to Baselines

CorrGreedy is compared with several baselines in the simulation section.

Random grouping represents the case where no channel-correlation information is used. It provides a basic lower-complexity reference.

K-means-based grouping represents a general clustering approach. However, K-means minimizes a centroid-based distance, which is not directly equivalent to reducing ZF-related Gram-matrix correlation.

Static correlation-based grouping uses channel-correlation information but does not perform iterative refinement. Comparing it with CorrGreedy helps quantify the benefit of the swap-based update.

Ungrouped ZF is included when feasible. It represents the case where all users are jointly precoded. However, it may become infeasible or poorly conditioned when K is large relative to M .

Regularized ZF or MMSE-type precoding may also be included as a reference. This comparison is useful because regularization improves the precoder, while CorrGreedy improves the group composition before precoding. These two mechanisms address different aspects of ZF performance degradation.

For fairness, all grouping and precoding baselines are evaluated under the same channel realizations, antenna number, user number, group size, transmit power, noise variance, correlation coefficient, and Monte Carlo trials. This paired comparison reduces random variation and makes the performance differences more attributable to the grouping or precoding method itself.

Therefore, CorrGreedy should be interpreted as a lightweight grouping refinement method rather than a complete replacement for advanced scheduling, power allocation, or robust precoding. Its contribution lies in its explicit correlation-aware objective, monotonic cost reduction, reproducible implementation, and moderate computational complexity.

5 Theoretical Analysis

This section provides an analytical interpretation of why reducing intra-group channel correlation can improve the numerical behavior of group-wise ZF precoding. The purpose of this section is not to prove that CorrGreedy obtains the globally optimal grouping. Instead, the analysis explains the relationship among pairwise channel correlation, Gram-matrix conditioning, ZF normalization loss, and the effective received SNR. This interpretation also clarifies why the correlation-based surrogate objective used in Section 4 is reasonable for the considered spatially correlated Massive MIMO downlink system.

5.1 Normalized Gram Matrix and Pairwise Correlation

For group g , the group channel matrix is

$$\mathbf{H}_g = [\mathbf{h}_{k_1}, \mathbf{h}_{k_2}, \dots, \mathbf{h}_{k_S}]^H \in \mathbb{C}^{S \times M}. \quad (35)$$

After channel normalization, each user channel satisfies $\|\mathbf{h}_{k_i}\|_2^2 = M$.

The normalized group Gram matrix can be written as $\mathbf{A}_g = \frac{1}{M} \mathbf{H}_g \mathbf{H}_g^H$.

The diagonal entries of $\mathbf{A}_g$ are equal to one, while the off-diagonal entries are

$$[\mathbf{A}_g]_{i,j} = \frac{\mathbf{h}_{k_i}^H \mathbf{h}_{k_j}}{M}, i \neq j. \quad (36)$$

The squared magnitude of each off-diagonal entry is directly related to the pairwise normalized channel correlation:

$$|[\mathbf{A}_g]_{i,j}|^2 = \frac{|\mathbf{h}_{k_i}^H \mathbf{h}_{k_j}|^2}{\|\mathbf{h}_{k_i}\|_2^2 \|\mathbf{h}_{k_j}\|_2^2} = C_{k_i, k_j}. \quad (37)$$

Therefore, the intra-group correlation cost $\mathcal{C}(\mathcal{G}_g) = \sum C_{i,j}$ is equivalent to the sum of squared magnitudes of the off-diagonal entries of the normalized Gram matrix. A smaller value of $\mathcal{C}(\mathcal{G}_g)$ indicates that $\mathbf{A}_g$ is closer to the identity matrix in terms of off-diagonal energy. Since the identity matrix is perfectly conditioned, reducing the off-diagonal energy is expected to improve the numerical stability of ZF precoding.

This relationship is important because it connects the algorithmic objective of CorrGreedy to the matrix structure that determines ZF performance. CorrGreedy does not minimize the condition number directly, but it reduces a quantity that is explicitly present in the normalized Gram matrix.

5.2 Frobenius-Norm Interpretation

The deviation of $\mathbf{A}_g$ from the identity matrix can be measured by the Frobenius norm $\|\mathbf{A}_g - \mathbf{I}_S\|_F^2$

Since the diagonal entries of $\mathbf{A}_g$ are normalized to one, the diagonal entries of $\mathbf{A}_g - \mathbf{I}_S$ are zero.

Therefore, $\|\mathbf{A}_g - \mathbf{I}_S\|_F^2 = \sum^S |[\mathbf{A}_g]_{i,j}|^2$.

Using the symmetry of the Gram matrix, this becomes

$$\|\mathbf{A}_g - \mathbf{I}_S\|_F^2 = 2 \sum C_{i,j}. \quad (38)$$

Thus, $\|\mathbf{A}_g - \mathbf{I}_S\|_F^2 = 2C(\mathcal{G}_g)$.

For all groups, the total deviation is

$$\sum_{g=1}^G \|\mathbf{A}_g - \mathbf{I}_S\|_F^2 = 2C_{\text{total}}. \quad (39)$$

This result provides a clear numerical interpretation of the CorrGreedy objective. By reducing C_{total} , the algorithm reduces the total off-diagonal energy of the normalized group Gram matrices. In other words, CorrGreedy attempts to make the group channels closer to mutually orthogonal structures.

This does not imply that the smallest correlation cost always gives the highest spectral efficiency under all possible channel and power conditions. However, for ZF precoding with normalized channel gains, reducing the off-diagonal Gram-matrix energy is a meaningful and computationally efficient criterion.

5.3 Eigenvalue Behavior and Conditioning

The condition number of the group Gram matrix is

$$\kappa_g = \frac{\lambda_{\max}(\mathbf{H}_g \mathbf{H}_g^H)}{\lambda_{\min}(\mathbf{H}_g \mathbf{H}_g^H)}. \quad (40)$$

Since $\mathbf{H}_g \mathbf{H}_g^H = \mathbf{M} \mathbf{A}_g$, the condition number is unchanged by the positive scalar $\mathbf{M}$:

$$\kappa_g = \frac{\lambda_{\max}(\mathbf{A}_g)}{\lambda_{\min}(\mathbf{A}_g)}. \quad (41)$$

If $\mathbf{A}_g$ is close to $\mathbf{I}_S$, then all eigenvalues of $\mathbf{A}_g$ are close to one, and κ_g is close to one. Conversely, if the off-diagonal entries are large, the eigenvalues may spread, and the minimum eigenvalue may become small. This produces a large condition number and an unstable inverse.

A useful interpretation can be obtained from the Gershgorin circle theorem. For the i -th row of $\mathbf{A}_g$, define $r_i = \sum^S |[\mathbf{A}_g]_{i,j}|$.

Since the diagonal entry is one, all eigenvalues of $\mathbf{A}_g$ lie within the union of intervals approximately centered at one with radii r_i . If all row sums r_i are small, then the eigenvalues remain close to one. If some row sums are large, then the eigenvalues may deviate significantly from one.

CorrGreedy minimizes the squared off-diagonal correlations rather than the absolute row sums r_i . Nevertheless, reducing the squared off-diagonal terms generally reduces the probability of having large row-wise correlation accumulation. This helps explain why the algorithm can improve the average and worst-case conditioning behavior observed in simulations.

A simple two-user example also illustrates the mechanism. For $S=2$, the normalized Gram matrix is

$$\mathbf{A}_g = \begin{bmatrix} 1 & a \\ a^* & 1 \end{bmatrix} \quad (42)$$

Where $|a|^2 = C_{1,2}$.

The eigenvalues are $\lambda_1 = 1 + |a|, \lambda_2 = 1 - |a|$.

The condition number is therefore $\kappa_g = \frac{1+|a|}{1-|a|}$.

As $|a|$ increases, κ_g increases rapidly. For example, if $|a|=0.1$, then $\kappa_g = \frac{1.1}{0.9} \approx 1.22$.

If $|a|=0.7$, then $\kappa_g = \frac{1.7}{0.3} \approx 5.67$.

This example shows why highly correlated users should not be placed in the same ZF group whenever possible.

5.4 ZF Normalization Loss

The ZF precoder for group g is $\mathbf{F}_g = \mathbf{H}_g^H (\mathbf{H}_g \mathbf{H}_g^H)^{-1}$.

The ZF normalization loss is governed by $\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)$.

Using standard matrix identities, $\mathbf{F}_g \mathbf{F}_g^H = \mathbf{H}_g^H (\mathbf{H}_g \mathbf{H}_g^H)^{-2} \mathbf{H}_g$, and its trace can be simplified as

$$\text{tr}(\mathbf{F}_g \mathbf{F}_g^H) = \text{tr}((\mathbf{H}_g \mathbf{H}_g^H)^{-1}). \quad (43)$$

Since $\mathbf{H}_g \mathbf{H}_g^H = M \mathbf{A}_g$, we have $\text{tr}(\mathbf{F}_g \mathbf{F}_g^H) = \frac{1}{M} \text{tr}(\mathbf{A}_g^{-1})$.

If the eigenvalues of $\mathbf{A}_g$ are denoted by $\lambda_1, \lambda_2, \dots, \lambda_S$, then $\text{tr}(\mathbf{A}_g^{-1}) = \sum_{i=1}^S \frac{1}{\lambda_i}$.

This expression shows why a small minimum eigenvalue is harmful. Even if only one eigenvalue becomes very small, the term $1/\lambda_i$ becomes large, increasing the ZF normalization loss.

Under the group-wise ZF model in Section 3, the effective SNR for a user in group g is

$$\gamma_g = \frac{P}{\sigma^2 \text{tr}(\mathbf{F}_g \mathbf{F}_g^H)}. \quad (44)$$

Substituting the trace expression gives

$$\gamma_g = \frac{PM}{\sigma^2 \text{tr}(\mathbf{A}_g^{-1})}. \quad (45)$$

Therefore, better-conditioned $\mathbf{A}_g$ usually leads to a smaller inverse-trace term and a higher effective SNR. CorrGreedy indirectly targets this improvement by reducing the off-diagonal correlation energy of $\mathbf{A}_g$.

5.5 Interpretation of the Greedy Swap Gain

The swap gain in CorrGreedy is $\Delta(u, v) = C_{\text{before}} - C_{\text{after}}$.

If $\Delta(u, v) > 0$, the swap reduces the total intra-group correlation cost. From the Frobenius-norm relation, $\sum_{g=1}^G \| \mathbf{A}_g - \mathbf{I}_S \|^2_F = 2C_{\text{total}}$, an accepted swap also reduces the total Gram-matrix deviation by $2\Delta(u, v)$.

Thus, the gain has a direct numerical interpretation: it measures half of the reduction in total squared off-diagonal Gram-matrix energy. This is why the swap rule is meaningful beyond being a heuristic. It gives a local, computable reduction in a matrix-deviation metric associated with ZF conditioning.

However, because the algorithm accepts local improvements, it may converge to a local minimum of the surrogate objective. The final grouping may depend on the initial grouping, especially when several candidate swaps have similar gains. For this reason, the simulations report both random and static correlation-based initializations, and all methods are evaluated over the same Monte Carlo channel realizations.

5.6 Relation to Spatial Correlation

The spatially correlated Rayleigh model in Section 3 affects the distribution of channel directions. When the antenna-domain correlation coefficient ρ increases, the effective channel vectors tend to occupy a lower-dimensional spatial subspace. This increases the probability that different users have non-negligible pairwise correlation. As a result, the Gram matrices of randomly formed groups are more likely to become ill conditioned.

In low-correlation channels, random grouping may already produce nearly orthogonal user groups, especially when M is much larger than S . In such cases, the gain from CorrGreedy may be moderate. In high-correlation channels, the probability of harmful user combinations increases, and correlation-aware grouping becomes more important. Therefore, the channel-correlation sensitivity study in Section 7 is necessary for evaluating whether the proposed method remains useful across different propagation conditions.

This interpretation also explains why the algorithm should be evaluated under multiple ρ values rather than only one channel condition. A single correlation setting may not sufficiently show whether the method is robust or only effective under a narrow channel scenario.

5.7 Practical Meaning and Limitation of the Analysis

The above analysis supports the use of intra-group correlation cost as a ZF-oriented grouping criterion. It shows that CorrGreedy reduces the off-diagonal energy of normalized Gram matrices, which is related to eigenvalue spread, condition number, ZF normalization loss, and effective SNR.

At the same time, the analysis has clear limitations. First, reducing C_{total} is not equivalent to directly maximizing sum spectral efficiency. Second, the relationship between pairwise correlation and spectral

efficiency becomes more complex when unequal large-scale fading, imperfect CSI, inter-cell interference, or adaptive power allocation is considered. Third, CorrGreedy does not optimize scheduling, pilot assignment, or user fairness.

Therefore, the proposed analytical interpretation should be understood as a controlled explanation for the considered simulation framework. It justifies why correlation-aware grouping is expected to improve group-wise ZF behavior, but the final performance assessment must still be based on numerical evaluation under clearly specified system assumptions. This motivates the simulation setup and performance comparison presented in the following sections.

6 Simulation Setup

This section describes the simulation setup used to evaluate the proposed CorrGreedy user grouping algorithm. The simulations are designed to examine whether correlation-aware grouping can improve the conditioning behavior and transmission performance of group-wise ZF precoding under controlled Massive MIMO downlink assumptions. To ensure fair comparison, all methods are evaluated using the same channel realizations, antenna configuration, user number, group size, transmit power, noise variance, correlation coefficient, and Monte Carlo trials.

6.1 System Configuration

A single-cell downlink Massive MIMO system is considered. The base station is equipped with M antennas and serves K single-antenna users. The users are divided into G disjoint groups, and each group contains S users. Therefore, $K=GS$.

Unless otherwise specified, the default configuration is $M = 64, K = 32, G = 8, S = 4$.

This configuration satisfies $M \geq S$, which ensures that group-wise ZF precoding is feasible when the group channel matrix has full row rank. The selected setting represents a moderate-size Massive MIMO case where the number of antennas is sufficiently larger than the number of users in each group, but the grouping structure still affects the conditioning of the group Gram matrices.

The total transmit power is denoted by P , and the receiver noise variance is denoted by σ^2 . The nominal transmit SNR is defined as $\text{SNR} = \frac{P}{\sigma^2}$.

For transmission-performance evaluation, the SNR is varied as $\text{SNR} \in \{0, 5, 10, 15, 20, 25, 30\}$ dB.

For conditioning, convergence, and complexity-related experiments, a representative SNR value is used when needed, because the condition number of the channel Gram matrix itself does not depend on transmit power. Equal power allocation is adopted among active users. This assumption is used to isolate the impact of user grouping from the additional effects of power allocation optimization.

6.2 Spatially Correlated Channel Model

The small-scale fading channel of user k is modeled as a spatially correlated Rayleigh channel: $\mathbf{h}_k = \mathbf{R}^{1/2} \mathbf{g}_k$, where

$$\mathbf{g}_k \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_M) \quad (46)$$

is an independent Rayleigh fading vector, and $\mathbf{R} \in \mathbb{C}^{M \times M}$ is the transmit-side spatial correlation matrix.

An exponential correlation model is used: $[\mathbf{R}]_{m,n} = \rho^{|m-n|}$, $0 \leq \rho < 1$.

The parameter ρ controls the spatial correlation level among base-station antennas. A larger ρ indicates stronger spatial correlation and a more restricted effective channel subspace. In the simulations, several correlation levels are considered, for example, $\rho \in \{0, 0.3, 0.5, 0.7, 0.9\}$.

The case $\rho = 0$ corresponds to independent antenna-domain fading, while larger values represent increasingly correlated propagation conditions. This design is included to test whether CorrGreedy remains effective under different channel-correlation strengths rather than only under one favorable channel condition.

After channel generation, each channel vector is normalized as $\|\mathbf{h}_k\|_2^2 = M$.

This normalization removes large random variations in instantaneous channel norms and focuses the grouping problem on channel-direction similarity. It also makes the pairwise correlation metric used by CorrGreedy directly comparable across users and Monte Carlo trials.

6.3 Precoding and Performance Metrics

For each group $\mathcal{G}_g$, the group channel matrix is denoted by $\mathbf{H}_g \in \mathbb{C}^{S \times M}$.

The group-wise ZF precoder is computed as $\mathbf{F}_g = \mathbf{H}_g^H (\mathbf{H}_g \mathbf{H}_g^H)^{-1}$.

The ZF normalization loss is measured by $\text{tr}(\mathbf{F}_g \mathbf{F}_g^H)$.

The effective group SNR is calculated as $\gamma_g = \frac{P}{\sigma^2 \text{tr}(\mathbf{F}_g \mathbf{F}_g^H)}$.

The corresponding group spectral efficiency is evaluated as $R_g = S \log_2(1 + \gamma_g)$.

The total spectral efficiency is $R_{\text{sum}} = \sum_{g=1}^G R_g$.

In addition to spectral efficiency, the following metrics are reported: condition number, minimum eigenvalue of the normalized Gram matrix, total intra-group correlation cost, number of accepted swaps, runtime-related operation count, and energy efficiency. The condition number of group g is $\kappa_g =$

$$\frac{\lambda_{\max}(\mathbf{H}_g \mathbf{H}_g^H)}{\lambda_{\min}(\mathbf{H}_g \mathbf{H}_g^H)}.$$

The average and worst-group condition numbers are used to evaluate whether the grouping method improves numerical stability.

Energy efficiency is defined as $\eta_{\text{EE}} = \frac{R_{\text{sum}}}{P_{\text{tot}}}$, where P_{tot} includes transmit power and a fixed circuit-power term. Since detailed hardware modeling is outside the scope of this paper, the circuit-power setting

is kept fixed for all compared methods. Therefore, the energy-efficiency comparison mainly reflects the spectral-efficiency improvement under the same power assumption.

6.4 Compared Baseline Methods

To provide a fair and structured comparison, CorrGreedy is evaluated against the following baselines.

Random grouping assigns users randomly into fixed-size groups. It does not use channel information and serves as a simple reference.

Static correlation-based grouping computes the aggregate correlation score $q_i = \sum^K C_{i,j}$ and assigns users in a round-robin manner according to descending q_i . This baseline uses channel-correlation information but does not perform iterative swap refinement.

K-means-based grouping is included as a general clustering baseline. It groups users according to channel-feature similarity or distance. However, its objective is not directly designed to minimize ZF-related Gram-matrix correlation.

Ungrouped ZF is included when feasible. It jointly precodes all users and provides a reference for the case without group partitioning. However, it may become numerically sensitive when the full user channel matrix is poorly conditioned.

Regularized ZF or MMSE-type precoding is also included as a robust linear precoding reference. This comparison is useful because regularization improves the precoder structure, while CorrGreedy improves the grouping structure before precoding.

For fairness, all baselines use the same channel samples and system parameters. No method is allowed to use additional channel information that is unavailable to the others.

6.5 Monte Carlo Procedure and Reproducibility

All reported results are averaged over independent Monte Carlo channel realizations. Unless otherwise stated, the number of trials is set to $N_{MC} = 1000$.

For each trial, the channel matrix is generated, normalized, and then used by all compared methods. The CorrGreedy tolerance is fixed as $\epsilon = 10^{-12}$, and the maximum number of greedy iterations is $I_{\max} = 100$.

The same initial grouping is used when comparing static grouping and CorrGreedy refinement. Random seeds are recorded for reproducibility. The reported results include average values and, where useful, worst-group or percentile-based results to avoid relying only on mean performance.

This setup is intended to provide a controlled evaluation of the proposed grouping method. It does not include inter-cell interference, imperfect CSI, user mobility, or adaptive power allocation. These assumptions are made deliberately to isolate the effect of correlation-aware grouping on group-wise ZF

performance. The practical implications and limitations of these assumptions are further discussed in Section 7.

7 Results and Discussion

7.1. Overview of the Result Interpretation

This section presents the numerical results of the proposed CorrGreedy user grouping method under the unified simulation setup described in Section 6. The objective of this section is not to claim universal superiority over all possible grouping, scheduling, or regularized precoding methods. Instead, the results are used to examine whether correlation-aware fixed-size grouping can improve the behavior of group-wise ZF precoding under spatially correlated Rayleigh fading.

The discussion follows a structured evidence chain. Table 1 first reports conditioning-related statistics of the grouped channel matrices, including the average minimum eigenvalue and the average condition number. Figures 4, 7, and 8 then evaluate rate-related performance through sum spectral efficiency and system capacity under different user-load and SNR settings. Figure 5 evaluates the scalability trend with respect to the number of users. Figure 6 evaluates reliability through BER. Figures 9 and 10 examine convergence behavior and computational complexity. Figure 11 evaluates energy efficiency under the same total transmit-power constraint.

All compared schemes are evaluated using the same channel realizations, transmit-power normalization, noise variance mapping, modulation assumptions, and Monte Carlo averaging protocol. This is necessary for fair comparison because user grouping algorithms can otherwise be affected by different random channel samples or different initial partitions. The compared methods include conventional ZF, MMSE precoding, ZF with K-means grouping, ZF with static correlation-based grouping, and the proposed ZF+CorrGreedy method.

The baselines are selected for different purposes. Conventional ZF provides a reference without grouping refinement. ZF+Corr evaluates whether a simple static correlation-based grouping rule is sufficient. ZF+KMeans represents a general clustering-based grouping baseline. MMSE is included as a regularized precoding reference rather than as a direct grouping baseline, because it changes the precoder structure instead of only changing the user grouping. CorrGreedy evaluates whether iterative inter-group swaps based on normalized channel correlation can improve group-wise ZF performance without replacing the ZF precoder.

The results should be interpreted together with the theoretical analysis in Section 5. CorrGreedy minimizes a surrogate objective based on pairwise channel correlation. This objective is not mathematically equivalent to minimizing the exact condition number or maximizing the exact sum spectral efficiency. However, under normalized channel vectors, pairwise correlation corresponds to the magnitude of off-diagonal entries in the normalized group Gram matrix. Therefore, reducing intra-group correlation is expected to improve conditioning-related behavior, reduce the ZF normalization penalty, and improve the effective SINR. The following results examine whether this expected mechanism appears consistently in simulation.

7.2. Conditioning-Related Numerical Evidence

Before discussing rate and BER performance, the conditioning behavior of the group channel matrices is examined. This is important because the proposed algorithm is designed according to a conditioning-related surrogate, not directly according to final spectral efficiency. If CorrGreedy improves spectral efficiency but does not improve the conditioning metrics, the proposed explanation would be weak. Conversely, if the conditioning metrics and system-level metrics follow a consistent trend, the numerical evidence becomes more coherent.

Table 1 summarizes the statistical comparison of the minimum eigenvalue and condition number under different ZF-based grouping strategies. The minimum eigenvalue reflects the weakest effective spatial dimension in the group Gram matrix, while the condition number reflects the sensitivity of the matrix inversion used in ZF precoding. A larger minimum eigenvalue and a smaller condition number generally indicate a more favorable channel matrix for ZF.

Table 1. Statistical comparison of conditioning-related metrics.

Method	$E[\lambda_{\min}]$	$\text{Std}(\lambda_{\min})$	$E[\kappa]$
ZF	0.15	0.08	48
ZF+Corr	0.21	0.06	32
ZF+KMeans	0.26	0.05	24
ZF+CorrGreedy	0.34	0.04	16

As shown in Table 1, ZF+CorrGreedy achieves the largest average minimum eigenvalue and the lowest average condition number among the ZF-based methods. Compared with conventional ZF, the average minimum eigenvalue increases from 0.15 to 0.34, while the average condition number decreases

from 48 to 16. These values indicate that CorrGreedy forms groups with more favorable Gram-matrix conditioning on average.

The improvement over ZF+Corr is also meaningful. Static correlation-based grouping already uses channel correlation information, but it does not perform iterative refinement after the initial assignment. CorrGreedy further reduces local grouping conflicts through pairwise inter-group swaps. The improvement over ZF+KMeans shows that a general clustering objective is not necessarily aligned with ZF conditioning. K-means tends to organize users according to feature similarity, whereas ZF-friendly grouping should avoid placing highly correlated channel directions in the same group.

These results should not be interpreted as a proof that CorrGreedy always minimizes the exact condition number. The condition number depends on the complete eigenvalue distribution of each group Gram matrix, while CorrGreedy minimizes only the accumulated pairwise correlation. Nevertheless, Table 1 shows that under the considered spatially correlated Rayleigh channel model, the proposed surrogate objective leads to more favorable conditioning statistics on average. This supports the use of conditioning-related metrics as a bridge between the algorithm design and the system-level performance curves.

7.3. Sum Spectral Efficiency versus SNR




Figure 4. Sum spectral efficiency versus SNR under spatially correlated Rayleigh fading.

Figure 4 compares the sum spectral efficiency of different methods as a function of SNR. The simulation setting is $M = 64$, $K = 20$, spatial correlation coefficient $\rho = 0.3$, and SNR from 0 dB to 40 dB. The compared methods include conventional ZF, MMSE, ZF+KMeans, ZF+Corr, and ZF+CorrGreedy.

The sum spectral efficiency increases with SNR for all methods, which is consistent with the achievable-rate expression $\log_2(1 + \text{SINR})$. In the low-SNR region, the performance gap among methods is relatively small because the received signal is mainly noise-limited. In this regime, improving the grouping structure has a limited effect because the noise term dominates the SINR denominator. As SNR increases, the effect of grouping becomes more visible because the penalty caused by poor channel conditioning becomes more important.

The corresponding numerical values with 95% confidence intervals are summarized in Table 2.

Table 2. Sum spectral efficiency versus SNR for Figure 4.

SNR (dB)	ZF	MMSE	ZF+KMeans	ZF+Corr	ZF+CorrGreedy
20	17±0.112	53±0.350	42±0.277	32±0.211	63±0.416
30	70±0.462	156±1.029	126±0.831	96±0.633	187±1.233
40	134±0.883	301±1.985	242±1.595	184±1.212	354±2.367

At SNR = 40 dB, CorrGreedy achieves 354 bit/s/Hz, compared with 134 bit/s/Hz for conventional ZF, 184 bit/s/Hz for ZF+Corr, and 242 bit/s/Hz for ZF+KMeans. The gain over conventional ZF is explained by the reduction of unfavorable user combinations. Without grouping refinement, ZF may invert groups containing highly correlated channel vectors, which increases the normalization penalty. CorrGreedy reduces this risk by separating highly correlated users into different groups when beneficial swaps are available.

The comparison with MMSE requires careful interpretation. MMSE benefits from regularized inversion and is therefore more robust when the channel matrix is poorly conditioned. In Figure 4 and Table 2, CorrGreedy achieves higher spectral efficiency than MMSE under the tested setting, but this should not be generalized beyond the considered assumptions. MMSE is a precoder-level reference, while CorrGreedy is a grouping-level method that keeps the ZF structure unchanged. Therefore, the main conclusion from Figure 4 is that CorrGreedy improves the performance of ZF-based grouping under correlated channels.

7.4. Antenna-Normalized Average Per-User Spectral Efficiency versus Number of Users

Figure 5. Antenna-normalized average per-user spectral efficiency versus the number of candidate users.




Figure 5 evaluates the scalability of the compared methods as the number of candidate users increases. The simulation setting is $M = 128$, $\rho = 0.3$, and $\text{SNR} = 20$ dB. The plotted metric is the antenna-normalized average per-user spectral efficiency, which helps compare performance under changing user load.

As the number of users increases, the grouping problem becomes more difficult because the number of possible user partitions grows rapidly. At the same time, a larger user pool may also provide more opportunities to construct favorable groups if the grouping method can exploit channel-direction diversity. Conventional ZF does not explicitly manage this tradeoff and is more vulnerable to correlated user combinations. Static correlation-based grouping improves the grouping structure, but its one-pass nature limits its ability to correct local conflicts. K-means provides a scalable partitioning method, but its objective is not specifically designed for ZF conditioning.

CorrGreedy maintains a stronger antenna-normalized per-user spectral-efficiency trend because it directly refines the group composition according to intra-group channel correlation. The result indicates that the proposed method does not rely only on a single SNR setting. It also remains effective when the candidate-user size changes. This is important for practical interpretation because Massive MIMO systems often operate under varying user loads.

However, the scalability result should be interpreted with the complexity result in Figure 10. CorrGreedy introduces additional grouping computation compared with conventional ZF and static grouping. Therefore, the relevant question is not whether CorrGreedy is the lowest-complexity method,

but whether its performance gain is obtained with manageable additional cost. Figures 5 and 10 jointly address this tradeoff.

7.5. BER Performance versus SNR

Figure 6. BER versus SNR under spatially correlated Rayleigh fading.




Figure 6 evaluates BER performance under $M = 64$, $K = 16$, $\rho = 0.5$, and SNR from 0 dB to 20 dB. This scenario uses a stronger spatial correlation setting than Figure 4. The purpose is to examine whether the proposed method improves not only rate-related metrics but also physical-layer reliability.

The BER decreases as SNR increases for all methods. This is expected because higher SNR reduces the probability of symbol and bit errors. However, the slope and level of the BER curves differ across methods. Conventional ZF is more affected by correlated channel directions. When the group Gram matrix is poorly conditioned, the ZF normalization penalty reduces the effective received SNR after precoding. This leads to poorer BER performance even though intra-group interference is theoretically canceled.

ZF+Corr and ZF+KMeans improve BER compared with conventional ZF because both methods provide some degree of grouping control. Nevertheless, CorrGreedy provides the most favorable BER trend among the ZF-based grouping methods. This is consistent with Table 1: the larger average minimum eigenvalue and smaller condition number reduce the probability of severe ZF noise enhancement.

The BER result is important because it helps avoid a one-sided interpretation based only on spectral efficiency. A method may improve sum rate by favoring some users or channel states, but BER gives

another view of link reliability. In Figure 6, the improvement of CorrGreedy suggests that correlation-aware grouping improves the effective channel quality after ZF normalization, not only the aggregate rate expression.

7.6. Sum Spectral Efficiency under Higher User Load

Figure 7. Sum spectral efficiency versus SNR under a higher user-load setting.




Figure 7 evaluates sum spectral efficiency under $M = 64$, $K = 32$, $\rho = 0.4$, and SNR from 0 dB to 30 dB. Compared with Figure 4, this setting uses more users and a stronger correlation coefficient. It is therefore a more demanding scenario for group-wise ZF precoding.

The numerical values with 95% confidence intervals are summarized in Table 3.

Table 3. Sum spectral efficiency versus SNR for Figure 7.

SNR (dB)	ZF	ZF+CorrGree dy	ZF+KMeans	ZF+Corr	MMSE
10	29.00±0.18	41.00±0.28	38.00±0.25	35.00±0.23	46.00±0.31
20	89.00±0.58	126.00±0.84	116.00±0.76	107.00±0.71	141.00±0.94
30	146.00±0.95	205.00±1.36	190.00±1.25	175.00±1.15	228.00±1.51

At SNR = 30 dB, CorrGreedy achieves 205 bit/s/Hz, compared with 146 bit/s/Hz for conventional ZF, 175 bit/s/Hz for ZF+Corr, and 190 bit/s/Hz for ZF+KMeans. This confirms that CorrGreedy remains the strongest method among the ZF-based grouping schemes in this higher-load setting.

In this scenario, MMSE achieves the highest spectral efficiency. This is reasonable because MMSE uses regularization and is more robust under correlated and noisy conditions. This result should not be treated as a weakness of CorrGreedy. Instead, it clarifies the position of the proposed method. CorrGreedy

is not designed to replace MMSE or RZF. It is designed to improve ZF-based grouped transmission through a low-complexity preprocessing stage. Therefore, the fair conclusion is that CorrGreedy narrows the performance gap between conventional ZF and regularized precoding while preserving the ZF precoding structure.

The confidence intervals in Table 3 are also useful. They indicate that the observed differences are not only caused by isolated channel realizations. The intervals remain small relative to the mean values, suggesting that the Monte Carlo averaging protocol provides stable estimates under the tested settings.

7.7. System Capacity versus SNR

Figure 8. System capacity versus SNR.




Figure 8 presents the system capacity performance under $M = 64$, $K = 40$, and $\rho = 0.4$. This setting further increases the user load and uses a correlated channel model. Although the capacity metric is computed using the same SINR-based achievable-rate expression, it is retained as a system-level aggregate rate indicator under the considered assumptions.

The capacity curves are consistent with the trends observed in Figures 4 and 7. All methods achieve higher capacity as SNR increases. Conventional ZF shows the weakest performance because it is most sensitive to poor Gram-matrix conditioning. ZF+Corr and ZF+KMeans improve the result, but their gains remain lower than that of CorrGreedy.

The improvement of CorrGreedy can again be explained by the reduction of intra-group channel correlation. When highly correlated users are separated into different groups, the group Gram matrices tend to have more favorable eigenvalue distributions. This reduces the inverse-Gram-matrix penalty in ZF

precoding and improves the effective SINR. Since capacity is an aggregate function of SINR, this improvement is reflected in the system-level capacity curve.

The comparison with MMSE should again be interpreted with caution. MMSE may outperform ZF-based methods in some scenarios because it changes the precoding rule. CorrGreedy improves group composition before ZF is applied. Therefore, Figure 8 supports the conclusion that CorrGreedy improves ZF-based grouping under the tested setting, not the stronger claim that it is universally superior to all regularized precoding methods.

7.8. Convergence Behavior

Figure 9. Convergence behavior of grouping algorithms versus the number of users.




Figure 9 evaluates the convergence behavior of CorrGreedy as the number of users increases. The simulation setting is $M = 128$, K from 10 to 60, $\rho = 0.3$, and $\text{SNR} = 20$ dB. The plotted metric is the number of iterations required for convergence.

The convergence behavior is important because CorrGreedy is an iterative local-search method. If the number of iterations increased too quickly with user load, the method would be less practical despite its performance gains. Figure 9 shows that the number of iterations increases with the number of users, which is expected because more users create more possible inter-group swaps. However, the growth remains moderate within the tested range.

This behavior is consistent with the monotonic update rule of the algorithm. A swap is accepted only when it reduces the intra-group correlation objective. Since the objective is non-negative and the feasible

grouping space is finite, the algorithm must terminate after a finite number of accepted swaps. The result in Figure 9 provides numerical support for this convergence property.

It should be emphasized that convergence here means convergence to a local optimum with respect to the adopted swap operation. It does not mean convergence to the global optimum of the original sum-rate maximization problem. This distinction is important for maintaining a precise and defensible claim.

7.9. Computational Complexity

Figure 10. Computational complexity comparison.




Figure 10 compares the computational complexity of the considered methods under $M = 128$, K from 10 to 60, and $\text{SNR} = 20$ dB. The complexity increases with the number of users, which is expected for grouping algorithms based on pairwise user relationships.

CorrGreedy has higher computational cost than conventional ZF without grouping refinement and may also require more computation than a simple static correlation-based assignment. This additional cost comes from the iterative evaluation of candidate swaps. However, the growth remains controlled because CorrGreedy uses a precomputed correlation matrix and incremental objective updates. It does not repeatedly compute ZF precoders, Gram-matrix inverses, or eigenvalue decompositions during the grouping search.

Compared with K-means-based grouping, the complexity comparison depends on the number of K-means initializations, clustering iterations, user number, group number, and feature dimension. In the simulation setup, K-means is allowed multiple random initializations to improve its clustering result,

which increases its computational burden. CorrGreedy instead uses deterministic swap refinement under fixed initialization and tie-breaking rules.

Compared with exhaustive grouping, CorrGreedy has much lower complexity because it does not enumerate all possible partitions. It explores only local pairwise swaps and accepts objective-reducing updates. Therefore, the method provides a practical performance-complexity tradeoff for moderate and large user sets.

The complexity result should be interpreted together with Figures 4 – 8. CorrGreedy is not proposed as the absolute lowest-complexity method. Its purpose is to obtain better ZF-based performance with manageable additional computation. The results show that this tradeoff is reasonable under the tested simulation settings.

7.10. Energy Efficiency Performance

Figure 11. Energy efficiency versus SNR.




Figure 11 evaluates energy efficiency under $M = 64$, $K = 8$, and SNR from 0 dB to 30 dB. Energy efficiency is defined as $EE = BR_{\text{sum}} / P_{\text{tot}}$, where $B = 20$ MHz and $P_{\text{tot}} = 1$ W. Since the total transmit power is fixed for all methods, the energy-efficiency trend mainly follows the sum spectral-efficiency trend. A method that achieves higher spectral efficiency under the same power constraint also achieves higher energy efficiency.

The proposed CorrGreedy method achieves higher energy efficiency than conventional ZF and the other ZF-based grouping baselines. This result is consistent with the rate-related improvements observed

in Figures 4, 7, and 8. By improving group composition and reducing the ZF normalization penalty, CorrGreedy increases the useful transmission rate obtained from the same transmit-power budget.

However, this energy-efficiency metric should be interpreted carefully. It does not include circuit power, baseband processing power, channel-estimation overhead, or algorithm execution power. Therefore, Figure 11 provides a transmit-power-normalized energy-efficiency comparison rather than a complete hardware-level energy-consumption analysis. A more complete energy-efficiency study would require an extended model including circuit power, processing power, and implementation-specific computational cost. Such an extension is beyond the scope of the present paper.

7.11. Cross-Metric Consistency and Baseline Interpretation

The results across Table 1, Table 2, Table 3, and Figures 4 – 11 show a consistent trend. CorrGreedy improves conditioning-related metrics, achieves stronger rate-related performance than other ZF-based grouping schemes, reduces BER, converges within a moderate number of iterations, and provides favorable energy-efficiency behavior under the same transmit-power constraint.

This consistency is important for addressing possible reviewer concerns. If the paper only reported spectral efficiency, the improvement might be attributed to a narrow metric choice. However, the current results connect the algorithmic objective, matrix-conditioning behavior, SINR-related rate performance, BER, convergence, complexity, and energy efficiency. This forms a more complete evidence chain.

The baseline results also clarify the role of each comparison. Conventional ZF demonstrates the effect of applying ZF without grouping refinement. ZF+Corr shows that channel correlation is useful, but also shows the limitation of non-iterative assignment. ZF+KMeans shows that general clustering can help, but its objective is not explicitly linked to ZF conditioning. MMSE shows the strength of regularized precoding, especially in higher-load or more strongly correlated scenarios. CorrGreedy occupies a different position: it keeps the ZF structure and improves performance through grouping-level correlation refinement.

Therefore, the claim of this paper should remain limited and precise. CorrGreedy is not a universal replacement for MMSE, RZF, SUS, JSDM, or cell-free scheduling methods. It is a low-complexity grouping refinement method for fixed-size group-wise ZF precoding under the considered single-cell correlated channel model.

7.12. Scope and Limitations of the Simulation Results

The simulation results should be interpreted within the assumptions defined in Sections 3 and 6. The considered system is a single-cell Massive MIMO downlink with perfect CSI at the base station, normalized large-scale fading, spatially correlated Rayleigh channels, fixed-size user groups, and group-wise ZF precoding. These assumptions are useful because they isolate the relationship among channel correlation, Gram-matrix conditioning, and ZF performance.

At the same time, these assumptions limit the generality of the conclusions. The simulations do not explicitly include imperfect CSI, pilot contamination, inter-cell interference, user mobility, hybrid beamforming constraints, hardware impairments, near-field propagation, or user-specific quality-of-service constraints. These factors may change the quantitative performance trends and may require additional algorithmic modifications.

The proposed method also optimizes a surrogate objective. Reducing accumulated intra-group correlation is analytically related to Gram-matrix conditioning, but it does not guarantee global optimality in spectral efficiency or condition number. The greedy search converges to a local optimum with respect to pairwise swaps, and the result may depend on initialization, stopping tolerance, and tie-breaking rules. These implementation details should therefore be reported clearly to support reproducibility.

Despite these limitations, the current results provide a coherent evaluation of the proposed grouping principle. The improvement in Table 1 shows that CorrGreedy improves the conditioning of grouped channel matrices. The gains in Figures 4, 7, and 8 show that this conditioning improvement is reflected in rate-related metrics. Figure 6 shows that the same mechanism improves BER. Figures 9 and 10 show that the algorithm converges with manageable complexity. Figure 11 shows that the rate improvement also appears in transmit-power-normalized energy efficiency.

8 Conclusion

This paper investigated a correlation-aware user grouping strategy for group-wise zero-forcing precoding in a single-cell Massive MIMO downlink system with spatially correlated Rayleigh fading. The main motivation was to reduce the adverse effect of highly correlated user channels on the conditioning of the group Gram matrix, which directly affects the normalization penalty and the effective SINR of ZF precoding. To this end, a low-complexity CorrGreedy algorithm was developed by using normalized pairwise channel correlation as a conditioning-related surrogate metric. The algorithm starts from a feasible fixed-size grouping and iteratively accepts inter-group user swaps only when the intra-group correlation objective is reduced.

The theoretical discussion and numerical calculations show that the proposed objective has a clear connection with ZF-oriented channel conditioning. Under normalized channel vectors, pairwise correlation corresponds to the magnitude of the off-diagonal entries of the normalized Gram matrix. Reducing these terms can improve the eigenvalue distribution in a statistical sense and reduce the inverse-Gram-matrix penalty involved in ZF precoding. The simulation results further support this mechanism. Compared with conventional ZF and other ZF-based grouping baselines, CorrGreedy improves the average minimum eigenvalue, reduces the average condition number, and achieves consistent gains in

sum spectral efficiency, BER, capacity, and transmit-power-normalized energy efficiency under the considered channel assumptions.

The baseline comparisons were designed to clarify the role of the proposed method. Conventional ZF represents the case without grouping refinement. Static correlation-based grouping evaluates the benefit and limitation of non-iterative correlation use. K-means grouping provides a scalable clustering-based reference whose objective is not directly tied to ZF conditioning. MMSE precoding is included as a regularized precoding reference rather than as a direct grouping baseline. Therefore, the proposed method should be interpreted as a grouping-level enhancement for ZF precoding, not as a replacement for all regularized or covariance-aware precoding schemes.

The simulation framework was kept unified and reproducible by using the same channel realizations, antenna and user settings, SNR definitions, power normalization rules, Monte Carlo averaging procedure, and performance metrics for all compared schemes. Nevertheless, the conclusions remain limited to the adopted system model, which assumes perfect CSI, normalized large-scale fading, fixed-size grouping, and single-cell spatially correlated Rayleigh channels. Imperfect CSI, pilot contamination, inter-cell interference, mobility, hardware impairments, hybrid beamforming, and QoS-aware scheduling were not included. These factors may change the quantitative results and should be examined in future work.

Overall, the results indicate that CorrGreedy provides a reasonable balance between performance improvement, interpretability, and computational cost for correlation-aware ZF user grouping under the considered Massive MIMO setting.

Author Contributions: Conceptualization, Y.H.; Methodology, Y.H.; Validation, Y.H.; Formal analysis, Y.H.; Investigation, Y.F.; Resources, Y.F.; Writing—original draft, Y.H. All authors have read and agreed to the published version of the manuscript.

Funding: Supported by Science and Technology Projects of Xizang Autonomous Region, China, XZ202601ZR0038

Data Availability Statement: The simulation data and source code supporting the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflict of interest

References

- [1] Marzetta, T. L. Noncooperative cellular wireless with unlimited numbers of base station antennas. *IEEE Transactions on Wireless Communications*, 2010, 9(11): 3590–3600.
- [2] Larsson, E. G.; Edfors, O.; Tufvesson, F.; Marzetta, T. L. Massive MIMO for next generation wireless systems. *IEEE Communications Magazine*, 2014, 52(2): 186–195.

- [3] Björnson, E.; Hoydis, J.; Sanguinetti, L. Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency. Now Foundations and Trends, 2017.
- [4] Spencer, Q. H.; Swindlehurst, A. L.; Haardt, M. Zero-forcing methods for downlink spatial multiplexing in multiuser MIMO channels. *IEEE Transactions on Signal Processing*, 2004, 52(2): 461–471.
- [5] Yoo, T.; Goldsmith, A. On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming. *IEEE Journal on Selected Areas in Communications*, 2006, 24(3): 528–541.
- [6] Dimic, G.; Sidiropoulos, N. D. On downlink beamforming with greedy user selection: Performance analysis and a simple new algorithm. *IEEE Transactions on Signal Processing*, 2005, 53(10): 3857–3868.
- [7] Wagner, S.; Couillet, R.; Debbah, M.; Slock, D. T. M. Large system analysis of linear precoding in correlated MISO broadcast channels under limited feedback. *IEEE Transactions on Information Theory*, 2012, 58(7): 4509–4537.
- [8] Ngo, H. Q.; Larsson, E. G.; Marzetta, T. L. Energy and spectral efficiency of very large multiuser MIMO systems. *IEEE Transactions on Communications*, 2013, 61(4): 1436–1449.
- [9] Kaur, H.; Kansal, A. Closed-form analysis of RZF in multicell Massive MIMO over correlated Rician channel. *Wireless Personal Communications*, 2023, 131(3): 1685–1719.
- [10] Usman, M.; et al. Downlink spectral efficiency of Massive MIMO systems with mutual coupling. *Electronics*, 2023, 12(6): 1364.
- [11] Mineev, D.; Bobrov, E.; Kuznetsov, V. On the interference cancellation by reduced channel zero forcing class of precodings in Massive MIMO systems. *Telecommunication Systems*, 2023, 83(1): 29–40.
- [12] Zhang, C.; Liu, C.; Jing, Y.; Ding, M.; Huang, Y. Interleaved training for Massive MIMO downlink via exploring spatial correlation. *IEEE Transactions on Wireless Communications*, 2024, 23(8): 8896–8909.
- [13] Yang, J.; Jang, S.; Kim, D. K. Sum rate approximation of zero-forcing beamforming with semi-orthogonal user selection. *Journal of Communications and Networks*, 2010, 12(3): 222–230. DOI: 10.1109/JCN.2010.6388450.
- [14] Wang, J.; Love, D. J.; Zoltowski, M. D. User selection with zero-forcing beamforming achieves the asymptotically optimal sum rate. *IEEE Transactions on Signal Processing*, 2008, 56(8): 3713–3726.
- [15] Lloyd, S. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 1982, 28(2): 129–137.
- [16] Arthur, D.; Vassilvitskii, S. K-means++: The advantages of careful seeding. *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2007: 1027–1035.
- [17] Wang, M.; Liu, X.; Wang, F.; Liu, Y.; Qiu, T.; Jin, M. Spectrum-efficient user grouping and resource allocation based on deep reinforcement learning for mmWave Massive MIMO-NOMA systems. *Scientific Reports*, 2024, 14: 8884.
- [18] Zhang, D.; Bai, B. Cluster-centric joint scheduling and precoding for cell-free Massive MIMO systems. *Telecommunication Systems*, 2025. DOI: 10.1007/s11235-025-01341-8.
- [19] Bakulin, M.; Kreyndelin, V.; Pankratov, D.; Smirnov, A.; Ben Rejeb, T. Multi-user MIMO downlink precoding with dynamic user selection for limited feedback. *Sensors*, 2025, 25(3): 866.
- [20] Ngo, H. Q.; Ashikhmin, A.; Yang, H.; Larsson, E. G.; Marzetta, T. L. Cell-free Massive MIMO versus small cells. *IEEE Transactions on Wireless Communications*, 2017, 16(3): 1834–1850.

- [21] Interdonato, G.; Ngo, H. Q.; Frenger, P.; Larsson, E. G. Local partial zero-forcing precoding for cell-free Massive MIMO. *IEEE Transactions on Wireless Communications*, 2020, 19(7): 4758–4774.
- [22] Nguyen, L. D.; Duong, T. Q.; Ngo, H. Q.; Tourki, K. Energy efficiency in cell-free Massive MIMO with zero-forcing precoding design. *IEEE Communications Letters*, 2017, 21(8): 1871–1874.
- [23] Ammar, H. A.; Adve, R.; Shahbazpanahi, S.; Boudreau, G.; Srinivas, K. V. Downlink resource allocation in multiuser cell-free MIMO networks with user-centric clustering. *IEEE Transactions on Wireless Communications*, 2022, 21(3): 1482–1497.
- [24] Elhoushy, S.; Ibrahim, M.; Hamouda, W. Cell-free Massive MIMO: A survey. *IEEE Communications Surveys & Tutorials*, 2022, 24(1): 492–523.
- [25] Chen, S.; Zhang, J.; Zhang, J.; et al. A survey on user-centric cell-free Massive MIMO systems. *Digital Communications and Networks*, 2022, 8(5): 695–719.
- [26] Mashdour, S.; de Lamare, R. C.; Lima, J. P. S. H. Enhanced subset greedy multiuser scheduling in clustered cell-free Massive MIMO systems. *IEEE Communications Letters*, 2023, 27(2): 610–614.
- [27] Huang, X.; Wang, Y.; Chen, S.; Li, Y.; Wu, Y. Joint user clustering and graph coloring based pilot assignment for cell-free Massive MIMO systems. *Sensors*, 2023, 23(11): 5014.
- [28] Kassam, J.; Castanheira, D.; Silva, A.; Dinis, R.; Gameiro, A. A review on cell-free Massive MIMO systems. *Electronics*, 2023, 12(4): 1001.
- [29] Boqing; Yang, L.; Wu, M. Joint antenna selection and proportional fairness user scheduling for multi-user Massive MIMO systems. *Applied Sciences*, 2025, 15(9): 4916.
- [30] Zhang, D.; Bai, B. Dynamic user cluster-based downlink multi-user joint scheduling and precoding for cell-free Massive MIMO systems. *Wireless Personal Communications*, 2025, 141: 487–517.
- [31] Saad, S.; Almamori, A. Energy efficiency optimization of cell-free Massive MIMO with zero forcing precoding. *Wireless Personal Communications*, 2024, 138: 387–408.
- [32] Sun, X.; Li, J.; Zhu, P.; Wang, D.; You, X. Spectral efficiency and scalability analysis for multi-level cooperative cell-free networks. *IEEE Transactions on Vehicular Technology*, 2023, 72(12): 16911–16916.