Contamination Aware Channel Estimation & Distributed Precoding For Cell Free Massive Mimo System With Mobility Support

by

Mohadese SHAJARI KOHAN

THESIS PRESENTED TO ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
IN PARTIAL FULFILLEMENT FOR A MASTER'S DEGREE
WITH THESIS IN ELECTRICAL ENGINEERING

M.A.Sc.

MONTREAL APRIL 2th, 2026

ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
UNIVERSITÉ DU QUÉBEC

BOARD OF EXAMINERS

THIS THESIS HAS BEEN EVALUATED

BY THE FOLLOWING BOARD OF EXAMINERS

Mr. Michel Kadoch, thesis supervisor
Department of Electrical Engineering, École de technologie supérieure

Mr. Abdelouahed Gherbi, President of the jury
Department of Software and IT Engineering, École de technologie supérieure

Mr. Jaafar Wael, jury
Department of Software and IT Engineering, École de technologie supérieure

THIS THESIS WAS PRESENTED AND DEFENDED

IN THE PRESENCE OF A BOARD OF EXAMINERS AND THE PUBLIC

ON MARCH, 24 2026

AT ÉCOLE DE TECHNOLOGIE SUPÉRIEURE

FORWARD

The rapid evolution of wireless communication systems toward beyond-5G and future network generations requires architectures that simultaneously deliver high spectral and energy efficiency, along with uniform quality of service, under realistic deployment conditions. Among emerging solutions, Cell-Free massive MIMO has attracted significant attention for its ability to eliminate traditional cell boundaries and enable cooperative transmission from distributed access points. However, translating theoretical gains into practical performance requires addressing several challenges, including pilot contamination, scalable precoding, support for mobility, and computational feasibility.

This thesis presents the development and evaluation of an integrated simulation framework that addresses these challenges within a unified Cell-Free massive MIMO environment. The work integrates contamination-aware channel estimation, distributed MMSE precoding, user-centric access-point selection with load balancing, and mobility-aware serving-set adaptation. The framework also supports dual-band operation across sub-6 GHz and millimetre-wave frequencies, enabling performance evaluation under heterogeneous propagation conditions.

The MATLAB-based simulation environment models realistic propagation effects, pilot reuse, temporal channel evolution, and user mobility. Extensive Monte Carlo simulations are conducted to assess the statistical robustness of the performance metrics. The results show that contamination-aware estimation substantially reduces pilot contamination, whereas distributed MMSE precoding achieves performance comparable to centralized processing with significantly lower fronthaul and computational requirements. Mobility simulations further confirm that dynamic updates to the serving set and handover mechanisms maintain service continuity as users move through the network.

The contributions of this work aim not only to improve system performance but also to provide practical implementation insights for scalable Cell-Free deployments. By integrating channel estimation, precoding, mobility management, and power allocation within a unified

framework, this thesis offers a comprehensive perspective on how distributed massive MIMO systems can operate efficiently under realistic constraints.

It is hoped that the methodologies and results presented in this work contribute to ongoing research efforts in practical large-scale distributed wireless networks and support future developments in next-generation communication systems.

ACKNOWLEDGMENTS

I want to express my sincere gratitude to my supervisor for the continuous guidance, technical insights, and constructive feedback provided throughout this research. Their expertise in wireless communications and Cell-Free massive MIMO systems was instrumental in shaping the direction of this thesis and in ensuring the technical rigour of the proposed methodologies.

I am also grateful to the members of the examination committee for their time, careful evaluation, and valuable comments, which helped improve the clarity and quality of this work.

I want to acknowledge the academic staff and faculty members of the Department of Electrical and Computer Engineering for providing a strong theoretical foundation and a supportive research environment. Their lectures, discussions, and academic resources significantly enhanced my understanding of advanced wireless communication systems.

Special thanks are extended to my colleagues and peers for their technical discussions, collaboration, and encouragement throughout the research process, particularly in simulation implementation, result validation, and performance analysis.

Finally, I would like to express my deepest appreciation to my family for their unwavering support, patience, and encouragement throughout my studies. Their motivation and understanding were essential to my completing this thesis.

Estimation de canal sensible à la contamination et précodage distribué pour les systèmes MIMO massifs sans cellule avec support de la mobilité

Mohadese SHAJARI KOHAN

RÉSUMÉ

Les systèmes MIMO massifs sans cellule (Cell-Free massive MIMO) constituent une architecture prometteuse pour les réseaux sans fil de cinquième génération et au-delà, en raison de leur capacité à offrir une couverture uniforme, une efficacité spectrale élevée et une meilleure équité entre les utilisateurs. En supprimant les frontières cellulaires traditionnelles et en faisant coopérer un grand nombre de points d'accès distribués, ces systèmes permettent d'exploiter pleinement la diversité spatiale. Toutefois, leur déploiement pratique est fortement limité par plusieurs défis majeurs, notamment la contamination des pilotes lors de l'estimation du canal, la complexité du précodage, les contraintes de fronthaul et la prise en charge de la mobilité des utilisateurs.

Ce mémoire propose un cadre intégré pour l'estimation du canal et le précodage distribué dans les systèmes MIMO Cell-Free massifs, en mettant l'accent sur la mitigation de la contamination par les pilotes et le soutien à la mobilité. Une méthode d'estimation du canal sensible à la contamination est développée, combinant une affectation spatiale des pilotes, un mécanisme d'annulation successive des interférences et une pondération adaptative tenant compte de la contamination résiduelle. Cette approche permet de réduire significativement l'impact des utilisateurs partageant les mêmes séquences de pilotes, améliorant ainsi la précision de l'estimation du canal par rapport à l'estimation LMMSE conventionnelle.

Par ailleurs, un schéma de précodage MMSE distribué avec régularisation adaptative est proposé afin d'atteindre des performances proches de celles du précodage MMSE centralisé, tout en limitant les exigences en termes de signalisation et de complexité de calcul. Le précodage est réalisé localement au niveau des points d'accès, en exploitant des informations de canal partielles, ce qui rend la solution compatible avec des architectures distribuées et

évolutives. Une stratégie d'allocation de puissance équitable est également intégrée afin d'améliorer l'efficacité énergétique et l'équité entre les utilisateurs.

La mobilité des utilisateurs est prise en compte à travers un modèle temporel de corrélation du canal et un mécanisme de sélection dynamique des points d'accès, basé sur la puissance du signal reçu et les contraintes d'équilibrage de charge. Le système est évalué par des simulations de Monte Carlo approfondies, incluant des scénarios bi-bandes (3,5 GHz et 28 GHz), des modèles de propagation réalistes et des utilisateurs mobiles à vitesse piétonne.

Les résultats montrent que le cadre proposé permet une réduction significative de la contamination des pilotes, une amélioration notable de l'efficacité spectrale et énergétique, ainsi qu'un maintien de performances stables en présence de mobilité. Ces contributions démontrent que les techniques développées dans cette thèse constituent une solution viable et performante pour le déploiement pratique de systèmes Cell-Free de MIMO à grande échelle dans les futurs réseaux sans fil.

Mots-clés: Cell-Free massive MIMO, estimation du canal, contamination des pilotes, précodage MMSE distribué, mobilité des utilisateurs, efficacité spectrale, efficacité énergétique, simulations de Monte Carlo.

Contamination Aware Channel Estimation & Distributed Precoding For Cell Free Massive Mimo System With Mobility Support

Mohadese SHAJARI KOHAN

ABSTRACT

Cell-Free massive multiple-input multiple-output (CF-mMIMO) systems have emerged as a promising architecture for next-generation wireless networks because they provide uniform coverage, high spectral efficiency, and improved user fairness by eliminating conventional cell boundaries. By enabling cooperation among a large number of geographically distributed access points (APs), CF-mMIMO systems exploit macro-diversity and spatial multiplexing gains. However, their practical deployment is constrained by several critical challenges, including pilot contamination during channel estimation, high computational complexity, fronthaul signalling limitations, and the need to support user mobility.

This thesis proposes an integrated framework for contamination-aware channel estimation and distributed precoding in CF-mMIMO systems with mobility support. A multi-stage contamination mitigation strategy is developed that combines spatial pilot assignment, successive interference cancellation, and contamination-aware linear minimum mean-square error (LMMSE) weighting. This approach explicitly accounts for the pilot-contamination structure and significantly improves channel-estimation accuracy compared with conventional LMMSE estimation.

To enable scalable and practical downlink transmission, a distributed minimum mean square error (MMSE) precoding scheme with signal-to-noise ratio (SNR)-adaptive regularization is proposed. The precoding is performed locally at each AP using partial channel state information, reducing fronthaul requirements while achieving performance comparable to centralized MMSE precoding. In addition, power allocation strategies are investigated to balance sum spectral efficiency, energy efficiency, and user fairness under per-AP power constraints.

User mobility is addressed through a temporal channel-correlation model and a dynamic access point selection and handover mechanism. The proposed mobility framework maintains service continuity while limiting handover frequency and load imbalance across APs. System performance is evaluated through extensive Monte Carlo simulations under realistic propagation conditions, including dual-band operation at sub-6 GHz and millimeter-wave frequencies.

Simulation results show that the proposed framework achieves substantial pilot-contamination reduction, near-centralized precoding performance, and stable operation under user mobility. These findings indicate that contamination-aware estimation combined with distributed MMSE precoding offers an effective and practical approach for deploying CF-mMIMO systems in future wireless networks.

Keywords: Cell-Free massive MIMO, pilot contamination, contamination-aware channel estimation, distributed MMSE precoding, user mobility, access point selection, power allocation, dual-band systems

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

List of Abbreviations and Acronyms

AP	Access Point
ASIC	Application-Specific Integrated Circuit
bps/Hz	bits per second per hertz
CF-mMIMO	Cell-Free massive multiple-input multiple-output
CPU	Central Processing Unit
CSI	Channel State Information
CV	Coefficient of Variation
DPC	Dirty Paper Coding
EE	Energy Efficiency
FFT	Fast Fourier Transform
FPGA	Field-Programmable Gate Array
GHz	Gigahertz
GPS	Global Positioning System

GPSDO	GPS-Disciplined Oscillator
ICI	Inter-Cell Interference
IoT	Internet of Things
I/Q	In-phase / Quadrature
J	Joule
km/h	kilometers per hour
K-factor	Rician K-factor
LMMSE	Linear Minimum Mean Square Error
LSFD	Large-Scale Fading Decoding
LoS	Line-of-Sight
MC	Monte Carlo
MIMO	Multiple-Input Multiple-Output
MMSE	Minimum Mean Square Error
MR	Maximum Ratio
mmWave	millimeter-wave

Mbit/J	Megabits per Joule
MHz	Megahertz
NMSE	Normalized Mean Square Error
NLOS	Non-Line-of-Sight
O-RAN	Open Radio Access Network
PON	Passive Optical Network
PTP	Precision Time Protocol
QoS	Quality of Service
RF	Radio Frequency
RIS	Reconfigurable Intelligent Surface
RSRP	Reference Signal Received Power
RZF	Regularized Zero-Forcing
SE	Spectral Efficiency
SF	Shadow Fading
SINR	Signal-to-Interference-Plus-Noise Ratio

SNR	Signal-to-Noise Ratio
sub-6 GHz	Sub-6 Gigahertz
TCO	Total Cost of Ownership
TDD	Time-Division Duplex
THP	Tomlinson–Harashima Precoding
TSN	Time-Sensitive Networking
UE	User Equipment
UMi	Urban Micro
URLLC	Ultra-Reliable Low-Latency Communication
ZF	Zero-Forcing

LIST OF SYMBOLS AND UNITS OF MEASUREMENT

h_{lk}	channel coefficient between access point l and user k (dimensionless)
w_{lk}	precoding weight applied at access point l for user k (dimensionless)
ρ_k	transmit power allocated to user k (watt, w)
$sinr_k$	signal-to-interference-plus-noise ratio of user k (decibel, db)
snr	signal-to-noise ratio (decibel, db)
β_{lk}	large-scale fading coefficient (dimensionless)
τ_p	pilot sequence length (symbols)
τ_c	channel coherence interval (symbols)
l	number of access points
k	number of users
n	number of antennas per access point
s	number of serving access points per user
f_c	carrier frequency (hertz, hz)
b	system bandwidth (hertz, hz)

r	achievable data rate (bits per second, bps)
se	spectral efficiency (bits per second per hertz, bps/hz)
ee	energy efficiency (bits per joule, bits/j)
p_{tx}	transmit power (watt, w)
σ^2	noise power (watt, w)
v	user velocity (kilometres per hour, km/h)

INTRODUCTION

Wireless communication systems continue to evolve to support the rapid growth of connected devices and data-intensive applications that require reliable, high-capacity connectivity. Conventional cellular networks, which rely on fixed cell boundaries and centralized base stations, struggle to deliver uniform service quality as user density and traffic demand increase. Users at cell edges often experience degraded performance due to interference and long access distances, resulting in non-uniform coverage and inefficient use of network resources.

Massive multiple-input multiple-output (MIMO) technology has significantly improved spectral efficiency by enabling simultaneous service to multiple users through spatial multiplexing. However, traditional cellular massive MIMO systems still suffer from limitations caused by cell boundaries and inter-cell interference. Cell-Free massive MIMO (CF-mMIMO) has therefore emerged as a promising alternative architecture, in which a large number of distributed access points cooperatively serve users without forming cells. By allowing users to be jointly served by multiple access points, CF-mMIMO improves coverage uniformity, enhances macro-diversity, and reduces the distance between users and the serving infrastructure.

Despite these advantages, practical deployment of CF-mMIMO systems presents several challenges. Accurate channel estimation remains difficult due to pilot contamination, which occurs when users reuse pilot sequences during channel estimation, leading to interference and degraded channel knowledge. In addition, scalable precoding strategies must carefully balance performance and computational complexity in distributed deployments, and user mobility requires continuous adaptation of serving access points to maintain reliable communication links.

Energy efficiency and user fairness must also be addressed, as distributed transmission from many access points increases infrastructure power consumption and may cause load imbalance across the network. Efficient power allocation and adaptive access point selection are therefore required to ensure stable system performance while preserving scalability.

To address these challenges, this work presents an integrated CF-mMIMO framework that combines contamination-aware channel estimation, distributed precoding, adaptive access point selection, and mobility-aware operation within a unified simulation environment. By modelling realistic propagation conditions and dynamic user movement, the study demonstrates that distributed cooperation among access points can sustain reliable performance, reduce the impact of pilot contamination, and remain practical to implement.

CHAPTER 1

RESEARCH OVERVIEW

1.1 Background and Motivation

The exponential growth of wireless data traffic and the increasing number of connected devices have created unprecedented challenges for modern communication systems. According to recent industry forecasts, global mobile data traffic is expected to nearly triple by 2028, while the number of connected devices will exceed 30 billion (Cisco, 2020). Traditional cellular networks, which rely on centralized base stations to serve users within defined cell boundaries, are approaching fundamental performance limits in spectral efficiency, energy efficiency, and coverage uniformity (Larsson et al., 2014).

Massive Multiple-Input Multiple-Output (massive MIMO) technology has emerged as a key enabler for fifth-generation (5G) wireless networks and beyond. By deploying a large number of antennas at the base station, massive MIMO systems can serve multiple users simultaneously on the same time-frequency resource through spatial multiplexing, thereby achieving significant gains in spectral efficiency (Marzetta, 2010). However, conventional cellular massive MIMO systems still suffer from several inherent limitations, including cell-edge performance degradation, inter-cell interference, and non-uniform quality of service (QoS) across the coverage area (Ngo et al., 2017).

Cell-Free massive MIMO (CF-mMIMO) represents a paradigm shift from traditional cellular architecture. In CF-mMIMO systems, a large number of geographically distributed Access Points (APs) cooperate to jointly serve all users in the coverage area, without creating cell boundaries (Ngo et al., 2017). Each AP is equipped with multiple antennas and connected to a Central Processing Unit (CPU) via fronthaul links. This distributed architecture offers several fundamental advantages over conventional cellular systems: (1) improved coverage uniformity

by eliminating cell boundaries, (2) enhanced macro-diversity gain through cooperation among distributed APs, (3) reduced access distance between users and serving APs, and (4) improved energy efficiency due to distributed transmission (Demir et al., 2021).

Despite these advantages, CF-mMIMO systems face significant technical challenges that must be addressed before practical deployment. One of the most critical challenges is pilot contamination, which occurs when multiple users transmit the same pilot sequences simultaneously during the channel estimation phase. Due to the limited coherence time of the wireless channel, the number of orthogonal pilot sequences is fundamentally limited, forcing pilot reuse among users. This pilot contamination severely degrades channel estimation quality and causes interference that does not vanish even as the number of AP antennas grows to infinity (Marzetta, 2010; Jose et al., 2011). In CF-mMIMO systems, where all APs must estimate channels to all users, the pilot contamination problem is even more severe than in conventional cellular massive MIMO.

Another critical challenge is supporting user mobility, which is essential for practical wireless systems. As users move, their channel characteristics change over time, requiring frequent channel re-estimation and potentially triggering handovers between serving APs. Mobility support in CF-mMIMO systems is particularly challenging because: (1) the distributed nature requires coordination among multiple APs for handover decisions, (2) channel aging effects become more pronounced with higher user velocities, and (3) the system must maintain service continuity as users transition between different sets of serving APs (Anand et al., 2023; D'Andrea et al., 2021; Zaher et al., 2024).

Furthermore, the computational complexity and signalling overhead of CF-mMIMO systems pose practical implementation challenges. Centralized precoding schemes, which collect all channel state information (CSI) at the CPU for joint processing, require substantial fronthaul capacity and impose a high computational burden on the CPU. Distributed precoding approaches, in which each AP performs local precoding based on partial CSI, reduce fronthaul requirements but may sacrifice performance (Björnson & Sanguinetti, 2020).

However, existing approaches in CF-mMIMO literature still exhibit important limitations that motivate this work. Standard LMMSE channel estimation, although widely adopted, remains highly sensitive to pilot contamination under practical pilot reuse conditions. Recent studies have shown that, even with advanced estimation techniques, pilot contamination continues to limit performance in large-scale CF-mMIMO deployments, especially under dense user scenarios and limited pilot resources (Björnson & Sanguinetti, 2020; Interdonato et al., 2019; Attarifar et al., 2021). In realistic configurations with pilot reuse, the contamination level can remain non-negligible, leading to degradation in channel estimation accuracy and overall system performance.

Moreover, most existing pilot assignment strategies do not fully exploit the spatial distribution of users, resulting in inefficient reuse patterns and increased interference. Spatially-aware pilot assignment and graph-based approaches have been investigated in recent works, yet they still face challenges in scalability and dynamic adaptation (Liu et al., 2020; Sabbagh et al., 2021).

In addition, while centralized MMSE precoding provides near-optimal performance, its implementation is constrained by high fronthaul requirements and computational complexity, which limits its scalability in large-scale CF-mMIMO systems (Björnson & Sanguinetti, 2020). Distributed precoding schemes offer a more practical alternative, but they typically suffer from performance degradation due to partial channel state information and lack of adaptive regularization mechanisms (Demir et al., 2021; Nayebi et al., 2017).

Furthermore, mobility is often simplified or neglected in many CF-mMIMO studies, despite its critical impact on channel variation, handover frequency, and service continuity. Recent investigations highlight that user mobility introduces additional challenges such as channel aging and dynamic AP association, which are not fully addressed in static system models (D’Andrea et al., 2021; Zaher et al., 2024).

Finally, power allocation strategies are frequently treated as secondary design aspects, although recent results demonstrate that optimized power allocation has a substantial impact on system performance, fairness, and energy efficiency in CF-mMIMO networks (Bashar et al., 2020; Van Chien et al., 2020).

Motivated by these limitations, this work focuses on the joint design of contamination-aware channel estimation, spatial pilot assignment, adaptive distributed MMSE precoding, and optimized power allocation, while explicitly incorporating user mobility into the system evaluation framework.

1.2 Problem Statement

Deploying CF-mMIMO systems in practical scenarios requires solving a set of tightly coupled design problems related to channel estimation, precoding, mobility management, and power allocation. These challenges are not independent, and their joint impact directly determines the achievable system performance.

Pilot Contamination and Channel Estimation:

In a system with k users and a limited number of orthogonal pilot sequences τp , pilot reuse becomes unavoidable when $K > \tau p$. This leads to pilot contamination, where the received pilot signals at each AP are corrupted by users sharing the same pilot. As a result, channel estimates become biased and the interference persists even with a large number of antennas. Recent studies confirm that pilot contamination remains a dominant performance bottleneck in CF-mMIMO systems, especially under dense user deployments and limited pilot resources (Björnson & Sanguinetti, 2020; Attarifar et al., 2021; Demir et al., 2021).

Moreover, conventional LMMSE estimation assumes Gaussian channel statistics and does not explicitly account for structured interference caused by pilot reuse. In realistic environments, the channel exhibits spatial correlation and frequency-dependent characteristics (sub-6 GHz and mmWave), which deviate from ideal Gaussian assumptions. Therefore, the estimation problem is not only limited by noise but also by structured pilot interference and non-i.i.d channel properties (Interdonato et al., 2019; Liu et al., 2020).

Precoding Design:

The system must design downlink transmit precoding at distributed APs, where each AP has

only partial CSI. This creates a trade-off between performance and implementation complexity. MR precoding can be implemented locally at each AP with low complexity but results in high inter-user interference. ZF and centralized MMSE precoding require global CSI at the CPU, leading to high fronthaul load and computational complexity.

In contrast, distributed MMSE precoding provides a scalable solution, but its performance strongly depends on the selection of the regularization parameter, which must adapt to the effective SNR and interference conditions across APs. Recent works have shown that adaptive and distributed precoding strategies are essential to approach centralized performance while maintaining scalability (Björnson & Sanguinetti, 2020; Demir et al., 2021; Bashar et al., 2020).

Mobility Support:

In practical scenarios, users move at pedestrian speeds, which introduces time variation in the wireless channel and affects both channel estimation and AP association. The system must continuously update the serving AP set, maintain reliable channel tracking, and manage handover events between AP clusters. Mobility introduces additional challenges such as channel aging and dynamic pilot reuse conflicts, which are not captured in static models. Recent studies highlight that mobility-aware CF-mMIMO design is still an open problem with significant performance impact (D'Andrea et al., 2021; Zaher et al., 2024).

Power Allocation:

Power allocation in CF-mMIMO is formulated as a constrained optimization problem, where transmit power must be distributed across users and APs while satisfying both total system power constraints and per-AP power limits**.** In this work, the allocation is performed at each AP based on local channel gains and precoding vectors, using a proportional fairness strategy with minimum power guarantees.

Specifically, each AP allocates power by combining:

- **a** minimum guaranteed power per active user
- **a** weighted allocation based on effective channel gain**,**

followed by a global normalization step to enforce system-wide power constraints.

This formulation reflects recent approaches where power allocation must balance sum-rate performance, fairness, and scalability under distributed processing constraints (Van Chien et al., 2020; Bashar et al., 2020). The problem is inherently complex because it involves multiple conflicting objectives and must be solved with low computational overhead in large-scale systems.

1.3 Research Objectives

The primary objective of this thesis is to develop and evaluate a practical and scalable Cell-Free massive MIMO (CF-mMIMO) system that improves performance under realistic deployment conditions. The work addresses key limitations such as pilot contamination, distributed processing constraints, and user mobility through a unified system design.
The specific objectives are defined as follows:

- To improve channel estimation robustness in CF-mMIMO systems, with emphasis on reducing the impact of pilot contamination under limited pilot resources and realistic channel conditions.
- To design distributed precoding schemes that approach the performance of centralized MMSE while operating with partial CSI at the APs and reduced fronthaul signalling, including the use of adaptive mechanisms.
- To incorporate user mobility into the CF-mMIMO framework, enabling dynamic AP selection and stable performance under time-varying channels.
- To develop power allocation strategies that jointly consider sum-rate performance and user fairness while satisfying per-AP and total power constraints in a distributed architecture.

1.4 Main Contributions

This thesis contributes to the design and evaluation of a practical CF-mMIMO system through several algorithmic and system-level developments. The main contributions are summarized as follows:

1. Contamination-Aware Channel Estimation Framework:

A contamination-aware channel estimation framework is developed by integrating spatial pilot reuse, successive interference cancellation (SIC), and modified LMMSE estimation within a unified processing structure.

While these techniques have been studied individually in prior works, the contribution of this work lies in their joint implementation with adaptive interference modeling and iterative refinement. In particular:

- a contamination graph is constructed to quantify interference between users sharing the same pilot,
- users are processed in descending order of contamination severity,
- SIC is applied with a reliability-based cancellation factor derived from NMSE, and
- the LMMSE estimator is modified to include a residual interference covariance matrix, instead of assuming only noise-limited conditions.

Furthermore, the estimation is refined through multiple SIC iterations, where cancellation aggressiveness is progressively adjusted.

This approach differs from standard LMMSE estimation by explicitly modeling structured pilot interference and incorporating estimation uncertainty into the cancellation process, thereby improving robustness under pilot reuse conditions.

2. Distributed MMSE Precoding with Adaptive Regularization:

A downlink distributed MMSE precoding scheme is implemented, where each AP computes its precoding vectors locally using partial CSI, without requiring global channel aggregation at the CPU.

The main contribution is the introduction of a multi-regime SNR-adaptive regularization strategy, where:

- the system SNR is estimated using effective channel gains and noise power,
- the operating regime is classified (very low, low, medium, high SNR), and
- the regularization parameter α is adjusted accordingly:
    - MR precoding is used in very low SNR,
    - strong regularization in low SNR,
    - moderate scaling in medium SNR,
    - and weak regularization (approaching ZF) in high SNR.

In addition, per-AP SNR estimation is used to adapt regularization locally, enabling heterogeneous behavior across APs.

This differs from conventional distributed MMSE implementations, which typically use a fixed regularization parameter, and therefore cannot adapt to changing channel and interference conditions.

The distributed formulation avoids global CSI exchange and reduces fronthaul signalling compared to centralized MMSE processing, since each AP only requires local channel information.

3. Mobility and Handover Management:

A mobility-aware system model is developed to incorporate user movement and dynamic AP association into CF-mMIMO operation.

The contribution includes:

- a time-correlated channel evolution model,
- a dynamic AP selection mechanism based on large-scale fading, and
- an RSRP-based handover strategy, where handover decisions are triggered when an alternative AP set provides higher received signal power.

The RSRP gain is defined as the difference between the received signal power before and after handover, allowing quantitative evaluation of handover decisions.

In addition, AP selection includes a load constraint (maximum number of users per AP) to prevent overload, while still prioritizing strong channel conditions.

This framework enables evaluation of CF-mMIMO performance under realistic mobility conditions, which is typically not addressed in static system models.

4. Distributed Power Allocation with Constraint Enforcement and Fairness Control

A distributed power allocation algorithm is developed based on per-AP optimization combined with global constraint enforcement.

The contribution is defined by a two-stage allocation structure:

- Stage 1 (Local allocation):
  - each AP allocates power based on effective channel gain $|h^H w|^2$,
  - a minimum guaranteed power per active user is enforced,
  - remaining power is distributed using proportional fairness weights ($\sqrt{gain}$)
- Stage 2 (Constraint enforcement):
  - per-AP power constraints are enforced through scaling,
  - total system power constraint is enforced through global normalization

In addition, fairness is explicitly evaluated using Jain's fairness index, ensuring that performance improvements do not come at the expense of user imbalance.

This approach differs from centralized optimization methods by enabling scalable implementation while maintaining fairness and strict constraint satisfaction

5. Comprehensive System Integration and Validation:

A unified simulation framework is developed to evaluate the interaction between channel estimation, precoding, mobility, and power allocation within a single system.

The contribution lies in enabling:

- consistent comparison between distributed and centralized processing,
- evaluation under both static and mobility scenarios, and
- analysis under realistic channel models including path loss, shadowing, and spatial correlation.

The framework allows systematic validation of algorithmic design choices under practical system constraints.

CHAPTER 2

LITERATURE REVIEW

2.1 Introduction

Cell-free massive multiple-input multiple-output (CF-mMIMO) has emerged as a promising wireless architecture to overcome the coverage limitations, inter-cell interference, and fairness issues inherent in conventional cellular networks. By replacing cell-centric operation with a user-centric serving paradigm, CF-mMIMO enables a large number of distributed access points (APs) to serve users over the same time–frequency resources in a cooperative manner. While this architecture offers substantial gains in spectral efficiency, energy efficiency, and user fairness, it also introduces new challenges in channel estimation, pilot contamination, signal processing, and system scalability.

Accurate channel state information (CSI) acquisition is a fundamental requirement for realizing the theoretical benefits of CF-mMIMO. However, the system's distributed nature, combined with the limited availability of orthogonal pilot sequences, leads to severe pilot contamination, which fundamentally limits performance even in the large-antenna regime. Unlike conventional cellular massive MIMO, pilot contamination in CF-mMIMO exhibits complex spatial characteristics due to overlapping serving sets and heterogeneous propagation conditions. As a result, classical estimation and mitigation techniques developed for cellular systems are not directly applicable and require substantial adaptation.

In parallel, the design of efficient precoding and signal-processing strategies is critical for balancing performance, computational complexity, and fronthaul requirements. Distributed and centralized processing architectures, as well as hybrid approaches such as large-scale fading decoding, have been extensively investigated to address the scalability challenges of CF-mMIMO. Linear precoding schemes, including maximum ratio, zero-forcing, and

minimum mean square error precoding, offer distinct trade-offs between interference suppression and implementation feasibility, particularly under imperfect CSI.

This chapter reviews the state of the art in CF-mMIMO with a focus on channel estimation, pilot contamination mitigation, and linear precoding strategies. The discussion highlights key theoretical insights, practical limitations, and recent advances reported in the literature, providing the necessary background and motivation for the proposed framework developed in the subsequent chapters.

2.2 Channel Estimation and Pilot Contamination

2.2.1 Fundamental Challenges in Channel Estimation

One of the fundamental challenges in CF-mMIMO systems is obtaining accurate CSI in the presence of pilot contamination, which becomes particularly severe when the number of users exceeds the number of available orthogonal pilot sequences. The distributed nature of CF-mMIMO exacerbates this challenge, as each AP must estimate the channels from all users, resulting in substantial estimation overhead. Marzetta (2010) first identified pilot contamination as the limiting factor in massive MIMO performance, showing that inter-cell interference from pilot reuse imposes a performance ceiling that cannot be overcome by simply adding more antennas. His analysis demonstrated that in the asymptotic regime where the number of antennas approaches infinity, the SINR saturates at a finite value determined by the pilot contamination level rather than increasing without bound. In the CF-mMIMO context, this problem manifests differently than in cellular massive MIMO, as users sharing the same pilot may be served by overlapping sets of APs, creating complex interference patterns that vary spatially across the network.

The impact of pilot contamination on CF-mMIMO performance has been quantified through extensive analysis and simulations. Studies have shown that in typical deployments with a pilot reuse factor of 2-4 (when the number of users is 2-4 times the number of orthogonal pilots), pilot contamination can degrade the achievable rate by 30-50% compared to the ideal case with perfect CSI (Ngo et al., 2017; Demir et al., 2021). The degradation is particularly severe for users located in areas with high AP density, where signals from multiple pilot-sharing users

combine constructively, creating strong contamination. Furthermore, the contamination effect exhibits a near-far phenomenon where users with stronger channel conditions contaminate the channel estimates of weaker users more severely than vice versa, leading to unfair performance distribution across the network. This understanding has motivated the development of advanced estimation techniques that explicitly account for and mitigate contamination effects.

2.2.2 Advanced Estimation Techniques

Recent advances in channel estimation for CF-mMIMO have focused on exploiting statistical channel properties and structured signal models to improve estimation accuracy under pilot contamination. Classical MMSE-based approaches remain a benchmark, where spatial correlation and large-scale fading are used to enhance estimation performance. More recent works extend these methods by incorporating adaptive estimation strategies and enhanced spatial modeling. For instance, Björnson and Sanguinetti (2020) and Demir et al. (2021) provide comprehensive frameworks for MMSE-based estimation under spatially correlated channels, showing that covariance-aware processing significantly improves estimation accuracy compared to conventional methods.

However, these approaches rely on accurate knowledge of channel statistics and assume relatively stable propagation conditions. In practical CF-mMIMO deployments, especially in heterogeneous environments combining sub-6 GHz and mmWave bands, channel statistics vary across users and access points. Recent studies highlight that such variability, together with imperfect CSI and hardware impairments, can significantly degrade estimation performance (Zaher et al., 2024; Chen et al., 2022).

To address pilot contamination more directly, recent research has focused on structured pilot assignment and interference-aware estimation techniques. Graph-based pilot allocation and spatial reuse strategies exploit user location and channel conditions to reduce pilot conflicts (Liu et al., 2020; Sabbagh et al., 2021). In addition, iterative interference cancellation approaches have been investigated, where contaminated pilot observations are progressively refined by subtracting estimated interference components. Attarifar et al. (2021) show that such

structured approaches can significantly improve estimation robustness compared to conventional MMSE methods.

Another important direction is the exploitation of channel sparsity and multi-dimensional structure. At higher frequencies, particularly in mmWave systems, channels exhibit angular sparsity, which can be leveraged using compressed sensing techniques. Chen et al. (2020) demonstrate that sparsity-based estimation can reduce pilot overhead while maintaining estimation accuracy. More recent works extend this idea using structured and hybrid estimation approaches, combining statistical and sparsity-based methods to improve performance under realistic conditions (Zhang et al., 2022).

Furthermore, recent measurement-based studies indicate that CF-mMIMO channels exhibit strong spatial variability and non-uniform propagation conditions across access points, which are not fully captured by idealized Gaussian models (Zaher et al., 2024). These findings highlight the need for estimation methods that explicitly account for spatial correlation, interference structure, and environmental dynamics.

Overall, despite significant progress, pilot contamination remains a fundamental limitation in CF-mMIMO systems, particularly under practical pilot reuse constraints and dense user deployments. Existing approaches either rely heavily on statistical assumptions, increase computational complexity, or require additional signalling overhead. Therefore, there is still a need for estimation frameworks that explicitly model structured interference and provide robust performance under realistic system conditions.

2.2.3 Pilot Assignment and Contamination Mitigation

Intelligent pilot assignment strategies have proven to be one of the most effective approaches for mitigating pilot contamination in CF-mMIMO systems. Chen and Björnson (2021) demonstrated that carefully designed pilot assignment algorithms can reduce contamination by up to 3 dB compared to random assignment by ensuring that users sharing the same pilot have minimal mutual interference (Liu et al., 2020; Sabbagh et al., 2021; Huang et al., 2023; Zhang et al., 2024). Their greedy algorithm considers both path-loss differences and serving-AP

overlap among users when assigning pilots, thereby exploiting the near-far effect to reduce interference. The algorithm iteratively assigns pilots to users, selecting at each step the user-pilot pair that minimizes the expected contamination based on large-scale fading coefficients (Ali et al., 2023; Chen et al., 2024). Simulation results show that this approach can improve the 5th percentile user rate by 40-60% compared to random pilot assignment, with the gains being most pronounced in heterogeneous deployments where path loss varies significantly across users (Wang et al., 2024).

Beyond pilot assignment, several other techniques have been developed to mitigate pilot contamination in CF-mMIMO. Pilot power control, where users adjust their pilot transmission power based on their channel conditions and interference levels, can reduce contamination by 20-30% while maintaining channel estimation quality for all users (Bashar et al., 2020; Van Chien et al., 2020; Chen et al., 2024). The key insight is that users with favourable channel conditions can reduce their pilot power without significantly affecting estimation accuracy, thereby reducing the contamination they impose on other users. Additionally, iterative channel estimation and data detection schemes have been proposed in which initial channel estimates are refined using detected data symbols, thereby increasing the number of observations available for channel estimation (Attarifar et al., 2021; Zhang et al., 2023). These iterative approaches can achieve performance within 1-2 dB of the perfect CSI case after 3-4 iterations, at the cost of increased computational complexity and latency.

2.2.4 Statistical CSI and Long-Term Channel Knowledge

Interdonato et al. (2019) proposed a framework for exploiting channel statistics and long-term channel knowledge in CF-mMIMO systems. Their work distinguishes between instantaneous CSI and statistical CSI, where statistical CSI includes channel covariance matrices, path loss, and spatial correlation. This information is not used for instantaneous detection, but for improving system-level design, such as pilot assignment, AP selection, and resource allocation.

In contrast to pilot assignment methods such as Chen and Björnson (2021), which mainly use large-scale fading for user grouping and pilot reuse, the approach of Interdonato et al. focuses on incorporating statistical CSI directly into the estimation process and system optimization.

In particular, covariance-aware MMSE estimation uses spatial correlation information to improve estimation accuracy, especially in scenarios with strong correlation such as urban environments.

The acquisition and tracking of statistical CSI introduce additional challenges in CF-mMIMO systems. Although statistical CSI varies slowly compared to instantaneous CSI, the distributed architecture requires coordination across multiple APs to maintain consistent estimates. Research has shown that statistical CSI typically remains valid for 100–1000 times longer than instantaneous CSI, corresponding to time scales of seconds to minutes in typical mobility scenarios (Björnson & Sanguinetti, 2020; Demir et al., 2021). Recent works have investigated scalable methods for statistical CSI acquisition and usage, including distributed covariance estimation and low-overhead statistical learning approaches (Demir et al., 2021; Zaher et al., 2024; Huang et al., 2023).

Furthermore, statistical CSI is used for long-term system optimization, including pilot reuse design, AP clustering, and power allocation. These parameters do not need to be updated at every coherence interval, but only when user positions or propagation conditions change significantly. This reduces signalling overhead compared to instantaneous CSI-based methods while maintaining stable system performance.

2.2.5 Adopted Estimation Framework and Contributions of This Work

While the fundamental principles of channel estimation and pilot contamination mitigation in CF-mMIMO are well established in the literature, this work adopts and extends these techniques to address practical limitations arising in dual-band operation, user mobility, and user-centric serving-set–based deployments. In this context, the *serving set* refers to the subset of access points (APs) that jointly serve each user based on large-scale fading conditions, as commonly considered in user-centric CF-mMIMO architectures (Ngo et al., 2017; Björnson & Sanguinetti, 2020; Demir et al., 2021).

The baseline estimation framework implemented in this study follows the covariance-aware LMMSE formulation presented in Ngo et al. (2017) and further developed by Björnson and

Sanguinetti (2020). In this formulation, channel estimation relies on statistical channel knowledge, including large-scale fading coefficients and spatial correlation matrices, rather than instantaneous channel realizations. This assumption is consistent with practical CF-mMIMO implementations, where statistical CSI can be estimated reliably over longer time scales.

Building on this baseline, this work introduces a contamination-aware channel estimation strategy that explicitly models pilot contamination at the AP–user link level. Unlike conventional LMMSE estimation, which incorporates contamination implicitly through the aggregate interference covariance, the proposed method constructs a contamination graph that quantifies mutual interference between users sharing the same pilot based on large-scale fading across all serving APs. Similar interference-aware formulations have been considered in recent works on pilot assignment and interference modelling (Liu et al., 2020; Sabbagh et al., 2021; Attarifar et al., 2021).

Based on this graph, users are ordered according to their contamination severity, defined as the total interference they receive from other users sharing the same pilot. Users with higher contamination are processed first in the estimation procedure. This ordering is motivated by the fact that highly contaminated users introduce stronger interference components in the received pilot signals. Estimating these users first allows their dominant interference contributions to be identified and partially removed before processing less contaminated users. In contrast, starting from weakly contaminated users would leave strong interference components unaccounted for, reducing the effectiveness of successive interference cancellation.

Successive interference cancellation (SIC) is then applied during the pilot observation phase, where previously estimated channels are subtracted from the received pilot signals. The cancellation is not full but weighted based on estimation reliability, which is quantified using the normalized mean square error (NMSE) of previously estimated channels. This approach is inspired by iterative interference cancellation techniques used in recent CF-mMIMO and multi-user detection studies (Attarifar et al., 2021; Zhang et al., 2023).

Furthermore, the proposed estimator incorporates contamination-aware covariance weighting, in which the residual interference covariance is dynamically adjusted based on the estimated reliability of previously cancelled users. This differs from conventional LMMSE estimation, where interference is treated as a static covariance term, and enables partial interference suppression without amplifying estimation errors.

Another extension of this work is the joint treatment of sub-6 GHz and millimetre-wave bands within a unified estimation framework. While many studies focus on a single frequency band, recent works highlight the importance of multi-band CF-mMIMO systems for achieving both coverage and capacity (Zaher et al., 2024; Demir et al., 2021). In this work, contamination-aware estimation is applied consistently across both bands while accounting for their different spatial correlation characteristics.

Overall, this work does not redefine the fundamentals of LMMSE estimation, but extends existing methods toward a contamination-aware, dual-band, and mobility-consistent framework. By explicitly modelling pilot interference, introducing contamination-based user ordering, and incorporating reliability-aware interference cancellation, the proposed approach improves robustness to pilot contamination while remaining consistent with established CF-mMIMO system models.

2.3 Precoding Strategies and Signal Processing

2.3.1 Distributed versus Centralized Processing

The choice between distributed and centralized signal processing represents a fundamental trade-off in CF-mMIMO system design, impacting performance, complexity, and scalability. In fully centralized processing, all received signals are forwarded to the CPU, where joint detection and precoding are performed, enabling optimal signal processing at the cost of massive fronthaul requirements and computational complexity. Conversely, distributed processing performs signal processing locally at each AP, requiring only the exchange of user data and control information over the fronthaul. However, it may sacrifice performance due to the lack of cooperation among APs. Björnson et al. (2020) conducted a comprehensive

comparison showing that distributed MMSE precoding can achieve 70-80% of the performance of centralized processing while reducing fronthaul requirements by 90% and computational complexity by 60%, making it the preferred choice for practical deployments.

The development of hybrid processing schemes that balance between centralized and distributed approaches has emerged as a promising direction. These schemes perform initial processing at the APs to reduce dimensionality and suppress interference locally, followed by centralized combining at the CPU to exploit the remaining correlation and diversity. Large-scale fading decoding (LSFD), in which signals are combined based on large-scale fading coefficients rather than on instantaneous channels, is one such approach that achieves near-optimal performance with significantly reduced complexity. Studies have shown that LSFD can achieve 90-95% of the optimal performance while requiring only statistical CSI at the CPU, eliminating the need to forward instantaneous channel estimates over the fronthaul (Nayebi et al., 2017; Björnson & Sanguinetti, 2020; Demir et al., 2021; Interdonato et al., 2019; Zaher et al., 2024). This approach is particularly attractive for deployments with limited fronthaul capacity or when serving highly mobile users, where instantaneous CSI quickly becomes outdated.

2.3.2 Linear Precoding Schemes

While early CF-mMIMO studies primarily focused on simple maximum-ratio (MR) precoding due to its suitability for distributed implementation, recent research has explored more sophisticated linear precoding schemes that better balance performance and complexity. MR precoding, in which each AP transmits along the channel vector to maximize the received signal strength, achieves a performance with minimal coordination among APs. However, MR precoding does not attempt to suppress interference, limiting its effectiveness in dense deployments with significant user overlap. Performance analysis shows that MR precoding achieves sum rates of 20-40 bps/Hz in typical CF-mMIMO deployments, with significant variations depending on user density and pilot contamination levels (Ngo et al., 2017; Björnson & Sanguinetti, 2020; Interdonato et al., 2019; Bashar et al., 2020).

Zero-forcing (ZF) and regularized zero-forcing (RZF) precoding, which explicitly nulls inter-user interference, can significantly improve performance in interference-limited scenarios. However, implementing ZF in CF-mMIMO requires global CSI and matrix inversion operations that scale with the total number of users and APs, making centralized implementation challenging for large networks. Distributed implementations of ZF have been proposed that approximate the global ZF solution through local processing and limited information exchange, achieving 80-90% of the performance of centralized ZF with manageable complexity. The key innovation in these distributed ZF schemes is the use of consensus algorithms and gradient methods to iteratively converge to the ZF solution via local computations and neighbour communication, typically requiring 5-10 iterations (Björnson & Sanguinetti, 2020; Nayebi et al., 2017; Interdonato et al., 2019; Van Chien et al., 2020; Zhang et al., 2023).

2.3.3 MMSE Precoding and Optimization

MMSE precoding, which optimally balances signal amplification and interference suppression across signal-to-noise ratio (SNR) regimes, has emerged as the preferred choice for CF-mMIMO systems. The MMSE precoder naturally adapts to operating conditions, approaching MR at low SNR, where noise dominates, and ZF at high SNR, where interference is the limiting factor (Björnson & Sanguinetti, 2020; Demir et al., 2021; Interdonato et al., 2019). Björnson et al. (2020) demonstrated that distributed MMSE precoding can achieve 20-50% higher sum rates than MR while maintaining reasonable complexity, with gains of 30-60 bps/Hz in typical deployments (Björnson & Sanguinetti, 2020; Bashar et al., 2020; Van Chien et al., 2020; Zaher et al., 2024). The performance advantage of MMSE is most pronounced in medium SNR regimes (10-20 dB), where the balance between noise and interference suppression is critical (Björnson & Sanguinetti, 2020; Demir et al., 2021).

The implementation of MMSE precoding in CF-mMIMO requires careful consideration of the regularization parameter, which controls the trade-off between signal amplification and interference suppression. Adaptive regularization schemes that adjust the parameter based on local SNR conditions and user density have been shown to improve performance by 10-15%

compared to fixed regularization (Björnson & Sanguinetti, 2020; Demir et al., 2021; Huang et al., 2023; Zhang et al., 2023). Furthermore, the computational complexity of MMSE precoding, which involves matrix inversion, can be reduced by employing various approximation techniques. Truncated polynomial expansion methods can approximate the MMSE precoder with 3-5 terms, achieving 95% of the optimal performance while reducing computational complexity by 50% (Björnson et al., 2014; Ngo et al., 2017; Demir et al., 2021). These approximations are particularly valuable for real-time implementation, where precoding matrices must be updated every coherence interval, typically every 1-10 milliseconds.

2.3.4 Non-Linear and Advanced Precoding

While linear precoding schemes dominate practical implementations due to their simplicity, Björnson and Sanguinetti (2020) investigated nonlinear precoding techniques for CF-mMIMO to approach the theoretical limits. Dirty paper coding (DPC), the capacity-achieving transmission strategy for multi-user channels, has been studied in the context of CF-mMIMO, showing potential gains of 30-40% over linear precoding (Björnson & Sanguinetti, 2020; Demir et al., 2021). However, the computational complexity of DPC, which grows exponentially with the number of users, makes it impractical for realistic CF-mMIMO deployments. Tomlinson-Harashima precoding (THP), a practical non-linear technique that achieves near-DPC performance with polynomial complexity, has been adapted for CF-mMIMO, yielding 15-20% gains over linear precoding in high-SNR regimes (Buzzi et al., 2019; Masoumi et al., 2022). The distributed implementation of THP remains challenging due to the sequential nature of the encoding process, limiting its applicability to systems with centralized processing.

Machine learning approaches to precoding design have recently attracted attention as a means of learning optimal precoding strategies from data rather than relying on model-based designs (Huang et al., 2023; Zhang et al., 2024; Wang et al., 2025). Deep neural networks trained on channel realizations and corresponding optimal precoding matrices can learn to approximate complex precoding strategies with low online complexity (Huang et al., 2023; Jiang et al.,

2024). Initial studies have shown that learning-based precoders can achieve performance within 5% of optimal MMSE precoding while reducing computational complexity by 40-50% using efficient neural network architectures (Zhang et al., 2024; Wang et al., 2025). The key advantage of learning-based approaches is their ability to adapt to specific deployment scenarios and channel conditions without explicit modelling, potentially capturing effects that are difficult to model analytically, such as hardware impairments and non-linear propagation phenomena (Jiang et al., 2024; Huang et al., 2023).

2.3.5 Adopted Precoding Framework and Contributions of This Work

This work adopts linear precoding as the core downlink transmission strategy due to its compatibility with distributed CF-mMIMO deployments and its favourable trade-off between performance and computational complexity. In particular, the baseline implementations follow the established formulations of maximum ratio (MR) precoding and (regularized) zero-forcing (ZF) precoding, which are widely used as reference schemes in CF-mMIMO literature. Distributed MMSE precoding is not used as a baseline since it constitutes the main contribution of this work. MR precoding is implemented as a low-complexity distributed benchmark, in which each AP applies the conjugate channel direction per user. In contrast, centralized ZF precoding is implemented as an interference-suppressing benchmark that requires full CSI availability at the CPU.

Building on these standard baselines, the main contribution in the precoding block is an adaptive distributed MMSE precoding strategy that explicitly accounts for the operating SNR regime and local AP conditions. Here, “adaptive” refers to the dynamic adjustment of the regularization parameter based on per-AP effective SNR levels and channel power variations across AP–user links. While conventional distributed MMSE formulations typically rely on fixed regularization parameters or assumptions that do not account for heterogeneous path loss across AP–user links, the proposed implementation performs SNR-aware regularization selection and per-AP SNR estimation using channel power statistics observed at each AP. This enables automatic regime detection (very low, low, medium, and high SNR), which ensures

that the regularization strength remains consistent with the theoretical behaviour of MMSE precoding: approaching MR in noise-limited regimes and gradually shifting toward interference suppression as the system becomes interference-limited.

2.4 Access Point Selection and User-Centric Clustering

Access point selection represents a critical design challenge in Cell-Free massive MIMO systems, where the fundamental question is determining which subset of distributed APs should serve each user to maximize network performance. Unlike traditional cellular networks, in which each user connects to a single base station, Cell-Free massive MIMO enables each user to be served simultaneously by multiple APs, creating a user-centric network architecture (Ngo et al., 2017; Björnson & Sanguinetti, 2020; Demir et al., 2021). However, having all APs serve all users is not practical due to fronthaul capacity limitations, pilot overhead constraints, and computational complexity at the central processing unit. Therefore, researchers have proposed various AP selection strategies that balance channel quality, load distribution, and system scalability. The most common approaches in the literature include distance-based selection, in which the nearest APs serve users; large-scale fading-based selection, in which APs with the strongest path-loss coefficients are chosen; and threshold-based methods, in which APs exceeding a specified channel-quality threshold are selected (Interdonato et al., 2019; Chen et al., 2020).

Recent work has shown that intelligent AP selection significantly **improves** both spectral and energy efficiency in Cell-Free massive MIMO networks. Ngo et al. (2017) demonstrated that selecting only APs with large-scale fading coefficients above a certain threshold can achieve performance close to that of a case where all APs serve all users, while dramatically reducing fronthaul signalling and computational burden. Björnson and Sanguinetti (2020) further showed that user-centric clustering with S serving APs per user (where S is typically 5–10) provides an excellent trade-off between performance and complexity, and has been reported to achieve more than 95% of the sum rate compared to full cooperation, with substantially lower overhead. However, most existing AP selection algorithms focus primarily on maximizing

channel quality, without explicitly considering load distribution across APs. This can lead to uneven resource utilization, where some APs serve many users while others remain underutilized, which may degrade overall system efficiency in heterogeneous deployments.

Advanced AP selection strategies that incorporate additional practical constraints have been proposed to address these limitations. In particular, recent works consider adaptive and dynamic AP selection mechanisms that account for user mobility, traffic variation, and multi-band operation (Chen et al., 2020; Sabbagh et al., 2021; Zaher et al., 2024). For dual-band systems operating at both sub-6 GHz and millimetre-wave frequencies, combined path-loss metrics that weight the frequency bands according to their coverage and capacity characteristics have shown promising results. The main challenge remains to design AP selection algorithms that jointly consider channel quality, system scalability, and practical deployment constraints, especially in mobile scenarios where serving sets must be updated dynamically. These challenges motivate the development of advanced AP selection strategies that balance performance and complexity under realistic CF-mMIMO conditions.

2.4.1 Adopted AP Selection Framework and Contributions of This Work

This work adopts the user-centric clustering paradigm established in the CF-mMIMO literature, in which each user is served by a subset of APs selected based on large-scale fading coefficients, rather than full cooperation among all APs. Following the principles discussed in Ngo et al. (2017) and subsequent user-centric clustering studies, the baseline objective is to retain most of the performance benefits of joint transmission while reducing fronthaul signalling and computational overhead. Consistent with common practice, the serving set size is limited to a fixed number of APs per user, reflecting realistic scalability constraints and enabling the network to operate under bounded complexity.

Building on these established concepts, the AP selection algorithm developed in this work introduces a quality-first greedy selection approach with explicit load-balancing constraints. Instead of selecting APs purely by channel quality for each user independently, the method enforces a maximum load per AP, preventing a small subset of APs from becoming overloaded

under non-uniform user distributions. Specifically, AP candidates are initially ranked using a frequency-aware large-scale fading metric, and the serving set is constructed by selecting the best APs that satisfy the load constraint. If the load constraint prevents selecting sufficiently high-quality APs, the algorithm completes the serving set using the least-loaded APs to guarantee that each user obtains exactly 5 serving APs, which corresponds to a typical user-centric cluster size in CF-mMIMO systems and provides a balance between performance and scalability.

A second extension of this work is the explicit integration of dual-band information into the clustering decision. While many AP selection strategies are evaluated at a single carrier frequency, the proposed method constructs a combined large-scale fading metric by weighting the sub-6 GHz and millimetre-wave path losses, thereby capturing the joint coverage and capacity behaviour of a dual-band architecture. This extension increases the complexity of the selection problem, since AP ranking must simultaneously account for heterogeneous propagation characteristics across frequency bands, where sub-6 GHz ensures reliable coverage and millimetre-wave provides high data rates but is more sensitive to blockage and distance. This introduces a trade-off in the selection process between robustness and capacity, which is not present in single-band formulations.

Overall, the contribution of this work is not to replace user-centric clustering but to adapt it into a practical, reproducible selection procedure that jointly addresses channel quality, load balancing, and dual-band effects under a fixed serving-set size. This provides a clear and defensible link between the literature, the system model implemented in Chapter 3, and the performance evaluations in subsequent chapters.

2.5 Power Control and Resource Allocation

2.5.1 Power Allocation Strategies

Power allocation in CF-mMIMO systems has specific challenges related to the distributed nature of transmission and individual power constraints at each AP. Unlike traditional cellular systems, where power allocation is performed at a single base station, CF-mMIMO requires coordinating power across multiple APs while respecting individual hardware limitations. Ngo

and Tran (2018) developed a comprehensive framework for power allocation that addresses these challenges and proposes algorithms for various optimization objectives, including sum-rate maximization, max-min fairness, and energy-efficiency maximization. Their max-min fair power allocation ensures uniform quality of service across users while respecting per-AP power constraints, improving the worst-case user rate by up to 5x compared to equal power allocation. The algorithm uses geometric programming to solve the non-convex optimization problem efficiently, converging to the global optimum in 10–15 iterations for typical network sizes, and serves as a theoretical performance reference for fairness-oriented power control in CF-mMIMO systems (Ngo & Tran, 2018; Björnson & Sanguinetti, 2020; Demir et al., 2021).

The impact of power constraints on CF-mMIMO performance has been thoroughly investigated, revealing important design considerations. Per-AP power constraints, which reflect hardware limitations of individual APs, can create bottlenecks in areas with high user density where APs must serve multiple users simultaneously **(**Björnson & Sanguinetti, 2020; Demir et al., 2021). Studies have shown that relaxing per-AP power constraints by 3 dB can improve system capacity by 15–20% in interference-limited scenarios, motivating the use of higher-capacity APs in dense areas (Van Chien et al., 2020; Bashar et al., 2020). Additionally, the total power constraint, which reflects the system's overall energy budget, determines the coverage–capacity trade-off (Björnson & Sanguinetti, 2020; Demir et al., 2021; Interdonato et al., 2019). Analysis shows that CF-mMIMO systems achieve optimal energy efficiency at moderate power levels, typically 20–30 dBm per AP, with diminishing returns at higher power levels due to increased interference (Ngo et al., 2017; Björnson & Sanguinetti, 2020). This finding has important implications for green communication design, suggesting that CF-mMIMO can achieve high performance while consuming relatively low power (Isheden et al., 2012; Björnson et al., 2015; Demir et al., 2021).

2.5.2 Fairness and Quality of Service

Ensuring fairness among users while maximizing overall system performance is a critical challenge in CF-mMIMO systems. Van Chien et al. (2021) proposed proportional fairness-

based power allocation schemes that balance sum-rate maximization and user fairness, achieving Jain's fairness index values above 0.8 while maintaining 90% of the optimal sum rate. Their approach employs logarithmic utility functions that naturally yield diminishing returns for users with already high rates, thereby encouraging the allocation of resources to disadvantaged users (Kelly, 1997; Mo & Walrand, 2000). The algorithm can be implemented in a distributed manner, with each AP adjusting its power allocation based on local CSI and feedback from the CPU, requiring only statistical CSI for coordination. Simulation results demonstrate that proportional fairness allocation can improve the 5th-percentile user rate by 2–3x compared to sum-rate maximization, with only a 10–15% reduction in total throughput (Van Chien et al., 2021; Bashar et al., 2020; Dai & Yu, 2015).

Quality of service (QoS) provisioning in CF-mMIMO requires careful consideration of the heterogeneous requirements of different users and applications. Users running real-time applications such as video streaming or online gaming require consistent low latency and minimum rate guarantees, while bulk data transfer applications can tolerate variable rates but benefit from high average throughput. Multi-objective optimization frameworks have been developed that simultaneously consider rate requirements, latency constraints, and fairness objectives. These frameworks employ weighted-sum approaches or hierarchical optimization to balance competing objectives, with weights adjusted based on network conditions and user priorities (Shi et al., 2011; Scutari et al., 2014; Zaher et al., 2024). Studies have shown that QoS-aware resource allocation can improve user satisfaction scores by 40–50% compared to best-effort approaches, while maintaining high overall system efficiency (Dai & Yu, 2015; Bashar et al., 2020; Zaher et al., 2024).

2.5.3 Energy Efficiency Optimization

Energy efficiency has become a critical performance metric for wireless systems due to environmental concerns and operational costs. CF-mMIMO systems face unique energy consumption challenges due to the large number of distributed APs and the associated fronthaul network. Ngo et al. (2018) conducted a comprehensive analysis of energy consumption in CF-mMIMO, considering circuit, signal-processing, fronthaul, and

transmission power. Their findings indicate that circuit power accounts for 60-70% of total energy consumption in typical deployments. This insight motivates the development of AP selection schemes in which only a subset of APs serves each user, thereby reducing circuit power consumption by turning off unused RF chains. Optimized AP selection can improve energy efficiency by 2-3x with minimal impact on spectral efficiency, achieving values of 0.1-0.5 Mbits/J in practical deployments.

Quality of service (QoS) provisioning in CF-mMIMO requires careful consideration of the heterogeneous requirements of different users and applications. Users running real-time applications such as video streaming or online gaming require consistent low latency and minimum rate guarantees, while bulk data transfer applications can tolerate variable rates but benefit from high average throughput. Multi-objective optimization frameworks have been developed that simultaneously consider rate requirements, latency constraints, and fairness objectives. These frameworks employ weighted-sum approaches or hierarchical optimization to balance competing objectives, with weights adjusted based on network conditions and user priorities (Shi et al., 2011; Scutari et al., 2014; Bashar et al., 2020; Zaher et al., 2024). Studies have shown that QoS-aware resource allocation can improve user satisfaction scores by 40-50% compared to best-effort approaches, while maintaining high overall system efficiency (Bashar et al., 2020; Van Chien et al., 2021; Zaher et al., 2024)**.**

2.5.4 Adopted Power Control Framework and Contributions of This Work

This work adopts proportional fairness as the main principle for power control, not as a new method, but as a well-established and widely accepted baseline for balancing sum rate and user fairness in multi-user wireless systems. In particular, the implemented power allocation strategy follows the proportional fairness formulation introduced by Kelly (1997), in which users are weighted such that marginal rate gains for already advantaged users are reduced, while users with weaker conditions receive relatively higher priority. This choice is motivated by the need for a stable and interpretable fairness benchmark, which allows consistent evaluation of system behaviour under heterogeneous channel conditions, as commonly considered in CF-mMIMO studies.

Building on this baseline, this work contributes by implementing a constraint-consistent proportional fairness procedure that explicitly enforces both per-AP power constraints and a global system power budget, while maintaining a minimum power guarantee for active users. Rather than extending the proportional fairness formulation itself, the contribution lies in the explicit enforcement of practical constraints during allocation. Unlike many algorithm descriptions that focus on the optimization objective and implicitly assume constraint feasibility, the proposed implementation includes a structured three-stage process. First, each AP performs local allocation using channel-dependent proportional fairness weights derived from the effective beamformed channel gain, which directly reflects the usable downlink gain under the selected precoder. Second, constraint verification is performed, and corrections are applied when any per-AP power constraint is exceeded, using scaling factors that preserve the relative allocation ratios. Third, a global power-budget enforcement stage ensures that the sum of allocations across all APs does not exceed the total system power, using strict tolerances and a safety margin to guarantee feasibility.

A further practical component of this work is the explicit minimum power guarantee per active user, which ensures baseline service continuity and prevents starvation under unfavourable channel conditions or heavy load. When the reserved minimum exceeds the available per-AP budget, the method falls back to equal allocation across active users to remain feasible. This mechanism reflects practical CF-mMIMO operation, where the number of served users per AP varies due to user-centric clustering and mobility-driven updates.

Finally, this work reports quantitative fairness diagnostics directly from the produced allocations, including Jain’s fairness index, the minimum-to-maximum power ratio, and the coefficient of variation. These metrics provide an interpretable validation of fairness behaviour and enable reproducible evaluation of the allocation strategy. Overall, the contribution lies in a reproducible and constraint-consistent power allocation procedure, rather than in redefining proportional fairness itself, ensuring that allocation decisions remain feasible and robust under realistic CF-mMIMO system constraints.

2.6 Multi-Band Operation and Frequency Integration

2.6.1 Dual-Band Architecture

The integration of multiple frequency bands in CF-mMIMO systems has emerged as a promising approach for leveraging the complementary characteristics of different frequency ranges. Buzzi et al. (2019) pioneered the analysis of dual-band CF-mMIMO, demonstrating that using sub-6 GHz bands for control signalling and initial access, while employing millimetre-wave (mmWave) frequencies for high-rate data transmission, can overcome mmWave coverage limitations while exploiting its large bandwidth. Their comprehensive analysis showed that dual-band systems can achieve 2-3x higher throughput than single-band deployments while maintaining robust connectivity. The sub-6 GHz band provides reliable coverage with path loss exponents of 3.2-3.8 in urban environments, while the mmWave band offers multi-gigabit rates but with path loss exponents of 4.5-5.5 and susceptibility to blockage. By intelligently switching between bands based on channel conditions and application requirements, dual-band CF-mMIMO can provide both coverage and capacity, thereby addressing a key limitation of mmWave-only systems (Buzzi et al., 2019; Alkhateeb et al., 2014; Rappaport et al., 2015; Polese et al., 2022; Giordani et al., 2023).

The architectural design of dual-band CF-mMIMO systems requires careful consideration of hardware complexity and cost. Each AP must be equipped with RF chains for both frequency bands, increasing hardware costs by 40-60% compared to single-band systems. However, this investment is justified by the performance gains and improved user experience. Furthermore, the two bands can share baseband processing and fronthaul resources, limiting the increase in operational complexity. Studies have shown that dual-band APs with 4 antennas at sub-6 GHz and 16-32 antennas at mmWave represent a good balance between performance and cost, providing sufficient beamforming gain at mmWave while maintaining reasonable hardware complexity. The placement of dual-band APs also requires optimization, with dense deployments in high-traffic areas to provide mmWave capacity and sparse deployments in low-traffic areas for sub-6 GHz coverage (Polese et al., 2022; Giordani et al., 2023; Shafi et al., 2023; Tataria et al., 2021).

2.6.2 Integration with Sub-THz Frequencies

Looking towards future wireless systems, the integration of sub-terahertz (sub-THz) frequencies (100–300 GHz) with CF-mMIMO architectures is being explored to achieve terabit-per-second area throughput. The extremely large bandwidth available at sub-THz frequencies, combined with the distributed nature of CF-mMIMO, could enable unprecedented capacity density (Akyildiz et al., 2023; Tataria et al., 2024). However, sub-THz propagation faces more severe challenges than mmWave, with significantly higher path loss and strong atmospheric absorption that depends on frequency and environmental conditions (Rappaport et al., 2023; Dang et al., 2023). Early studies indicate that dense CF-mMIMO deployments with short inter-AP distances can mitigate these effects and provide reliable high-rate communication in indoor environments through macro-diversity and distributed beamforming (Björnson et al., 2023; Tataria et al., 2024).

The hardware requirements for sub-THz CF-mMIMO systems are considerably more demanding than those of current mmWave implementations, mainly due to the need for large antenna arrays and high beamforming gains to compensate for severe propagation losses. Recent advances in semiconductor technologies, including silicon-germanium (SiGe) and indium phosphide (InP), have demonstrated the feasibility of sub-THz transceivers, although power consumption remains a critical limitation (Zhang et al., 2023; Nagatsuma et al., 2023). Hybrid analog–digital beamforming architectures have therefore emerged as a practical solution to reduce the number of RF chains while preserving beamforming performance (Alkhateeb, 2024). Recent studies show that hybrid architectures can achieve performance close to fully digital beamforming with significantly reduced hardware complexity and energy consumption, making them suitable for future sub-THz CF-mMIMO systems (Alkhateeb, 2024; You et al., 2024).

2.6.3 Adopted Multi-Band Modelling Framework and Contributions of This Work

This work adopts a dual-band CF-mMIMO architecture in which sub-6 GHz and millimetre-wave frequency bands are jointly exploited to balance coverage reliability and achievable data rates. Consistent with established system-level studies, sub-6 GHz links provide robust connectivity and control signalling, while millimetre-wave links enable high-capacity data transmission under favourable propagation conditions. This separation reflects practical

deployment strategies and captures the complementary propagation characteristics of the two frequency ranges.

Building on this framework, a unified large-scale fading model is implemented to represent frequency-dependent propagation effects across both bands within a common CF-mMIMO system. Specifically, the large-scale fading coefficient βk,l(f)\beta_{k,l}^{(f)}βk,l(f) is computed for each user–AP pair and for each frequency band, incorporating distance-dependent path loss and shadow fading. For millimetre-wave links, additional sensitivity to blockage is captured through band-dependent attenuation characteristics, resulting in significantly different large-scale fading distributions between the two bands. These coefficients are consistently used across all system components, including AP selection, channel estimation, and power allocation.

A central contribution of this work is the joint utilization of dual-band large-scale fading information in the user-centric clustering process. Instead of performing access point selection independently per frequency band, the system integrates the large-scale fading coefficients across both bands through a weighted combination, enabling serving-set formation that simultaneously accounts for link reliability and achievable capacity. This mechanism allows the system to prioritize stable sub-6 GHz links while opportunistically exploiting strong millimetre-wave links when available.

This joint clustering strategy directly impacts system robustness under heterogeneous propagation conditions. In particular, when millimetre-wave links experience degradation due to blockage, the system maintains connectivity through sub-6 GHz links without requiring reconfiguration of the entire serving set. Conversely, when favourable millimetre-wave conditions are present, the system dynamically benefits from increased channel gain and higher achievable rates. As a result, the proposed dual-band integration provides a practical and scalable framework for CF-mMIMO operation under realistic propagation and mobility conditions.

2.7 Implementation Challenges and Practical Deployments

2.7.1 Fronthaul Requirements and Solutions

The fronthaul network connecting distributed access points (APs) to the central processing unit (CPU) constitutes a major constraint in CF-mMIMO system deployment. The required fronthaul capacity depends directly on the selected functional split between the APs and the CPU. In fully centralized architectures, raw in-phase and quadrature (I/Q) samples must be transported, which leads to very high data rate requirements that can reach tens of gigabits per second per AP, depending on system bandwidth and the number of antennas (Björnson et al., 2023; Checko et al., 2023).

To mitigate this limitation, fronthaul compression and quantization techniques are applied. Existing studies show that vector quantization and compression methods can reduce fronthaul data rates by 50–70% while preserving system performance with only minor degradation (Masoumi and Emadi, 2020; Bashar et al., 2023). In this context, performance is evaluated in terms of spectral efficiency (sum rate) and achievable SINR, where the loss due to compression remains within approximately 5% compared to the uncompressed centralized case.

Alternative functional splits further reduce fronthaul requirements by shifting part of the signal processing to the APs. In particular, operations such as FFT and local combining can be performed locally, which decreases the required fronthaul rate to the order of 1–10 Gbps per AP (Interdonato et al., 2023; Björnson et al., 2023). This reduction makes the system compatible with existing fibre-based fronthaul infrastructure and high-capacity wireless backhaul solutions.

In addition to capacity, fronthaul design must satisfy strict latency and synchronization constraints. Coherent joint transmission requires phase alignment across distributed APs, which imposes latency on the order of tens of microseconds and very low jitter (Interdonato et al., 2023; Shariatmadari et al., 2024). Technologies such as time-sensitive networking (TSN) and enhanced passive optical networks (PON) provide deterministic latency and high reliability, and are therefore suitable for CF-mMIMO fronthaul (Shariatmadari et al., 2024; Ericsson, 2023). Experimental results indicate that modern PON systems can support large

numbers of APs while maintaining latency below 100 microseconds and high reliability, which satisfies the requirements of distributed massive MIMO systems.

Wireless fronthaul solutions based on millimetre-wave and sub-THz links are also considered for scenarios where fibre deployment is not practical. Recent demonstrations show that multi-gigabit-per-second data rates can be achieved over several hundred meters with high reliability using adaptive beamforming and diversity techniques (Rappaport et al., 2023; Nagatsuma et al., 2023). These results support the use of hybrid fibre–wireless fronthaul architectures for scalable CF-mMIMO deployment.

2.7.2 Synchronization and Calibration

Maintaining phase and time synchronization across distributed APs is essential for coherent joint transmission in CF-mMIMO systems. Analytical and system-level studies show that strict synchronization accuracy is required to preserve coherent combining gains, with phase errors on the order of a few degrees and timing offsets below the microsecond level (Demir et al., 2021; Interdonato et al., 2023; Björnson et al., 2023). It has been shown that phase misalignment directly reduces beamforming gain and increases inter-user interference, leading to noticeable degradation in achievable SINR and spectral efficiency (Björnson et al., 2023; Shariatmadari et al., 2024).

Global Navigation Satellite System (GNSS)-based solutions, such as GPS-disciplined oscillators (GPSDO), can provide high synchronization accuracy with low phase drift under outdoor conditions (Shariatmadari et al., 2024). However, their applicability is limited in indoor deployments where satellite signals are not available. As an alternative, network-based synchronization using IEEE 1588 Precision Time Protocol (PTP) has been widely adopted, enabling sub-microsecond synchronization accuracy over packet-based fronthaul networks when supported by appropriate hardware timestamping mechanisms (Checko et al., 2023; Shariatmadari et al., 2024).

Hardware calibration represents an additional challenge, since component variations and environmental conditions introduce phase and amplitude mismatches across RF chains. These impairments affect coherent transmission and can reduce system performance if not properly compensated (Björnson et al., 2023; Interdonato et al., 2023). Experimental and analytical studies indicate that even moderate mismatches in amplitude and phase can lead to noticeable reductions in achievable rate and beamforming efficiency (Shariatmadari et al., 2024).

To address this issue, over-the-air calibration techniques have been proposed, where APs transmit known reference signals that are used to estimate and compensate hardware-induced impairments across the network (Interdonato et al., 2023; Björnson et al., 2023). More recently, data-driven approaches based on machine learning have been investigated to track and compensate time-varying impairments such as temperature-dependent phase drift, reducing calibration overhead while maintaining stable system performance (Zhang et al., 2023; You et al., 2024). Joint calibration strategies across multiple APs further improve phase alignment and enable sustained coherent operation in distributed CF-mMIMO systems (Interdonato et al., 2023).

2.7.3 Computational Complexity and Processing

The computational requirements of CF-mMIMO signal processing present significant challenges for real-time implementation (Chen et al., 2020; Björnson et al., 2023; Interdonato et al., 2023). Chen et al. (2020) developed a detailed computational analysis showing that centralized MMSE processing for a network with 100 APs and 40 users requires approximately 100 GFLOPS for channel estimation and precoding computation, which must be completed within the coherence time of 1–10 ms. This computational load exceeds the capabilities of traditional base-station processors, necessitating the use of high-performance computing platforms (Björnson et al., 2023; Checko et al., 2023). Their work proposed sequential processing approaches that reduce the computational burden by 60% through distributed preprocessing at the APs, in which local MMSE processing reduces dimensionality before forwarding to the CPU (Chen et al., 2020; Interdonato et al., 2023). This hierarchical

processing achieved latencies below 1 ms on commercial server hardware with 32-core processors and GPU acceleration, meeting the requirements for ultra-reliable low-latency communication (URLLC) applications (Shariatmadari et al., 2024; You et al., 2024).

The development of specialized hardware accelerators for CF-mMIMO processing has attracted attention as a means of meeting computational requirements while maintaining energy efficiency (Björnson et al., 2023; Zhang et al., 2023). Application-specific integrated circuits (ASICs) designed for matrix operations common in MMSE processing can achieve 10–100× better energy efficiency than general-purpose processors (Zhang et al., 2023; You et al., 2024). Field-programmable gate arrays (FPGAs) offer a middle ground, providing approximately 10× better efficiency than CPUs while maintaining flexibility for algorithm updates (Checko et al., 2023; Zhang et al., 2023). Recent implementations have demonstrated that a single high-end FPGA can support real-time processing for 32 APs and 16 users, with latency below 100 microseconds and power consumption below 50 W (Zhang et al., 2023; Shariatmadari et al., 2024). The parallelizable nature of many CF-mMIMO algorithms makes them well suited to hardware acceleration, with systolic arrays and dataflow architectures showing particular promise for efficient implementation (You et al., 2024; Zhang et al., 2023).

2.7.4 Field Trials and Commercial Deployments

Several field trials and proof-of-concept deployments have validated the performance benefits of CF-mMIMO in real-world environments (Björnson et al., 2023; Interdonato et al., 2023). Ericsson's field trial in Stockholm demonstrated a CF-mMIMO system with 64 distributed APs serving 12 users in a 200×200 meter area, achieving a 10x improvement in 5th-percentile user throughput relative to a traditional cellular deployment with the same total number of antennas (Ericsson, 2023; Shariatmadari et al., 2024). The trial revealed practical challenges including the impact of non-line-of-sight propagation in urban environments and the importance of AP placement optimization (Björnson et al., 2023; Interdonato et al., 2023). Nokia's trial in a factory environment demonstrated that CF-mMIMO could provide sub-millisecond latency for industrial IoT applications while supporting enhanced mobile broadband for workers,

demonstrating the technology's ability to simultaneously meet diverse service requirements (Nokia, 2023; You et al., 2024).

The path to commercial deployment of CF-mMIMO involves addressing not only technical challenges but also economic and regulatory considerations (You et al., 2024; Tataria et al., 2024). The total cost of ownership (TCO) analysis shows that while CF-mMIMO requires higher initial investment due to the large number of APs and fronthaul infrastructure, operational costs can be lower due to improved energy efficiency and reduced site rental fees for distributed small APs compared to traditional macro base stations (Björnson et al., 2023; Ericsson, 2023). Regulatory frameworks for CF-mMIMO deployment are still evolving, with questions about spectrum licensing, emission limits, and coordination between operators in shared spectrum scenarios (Tataria et al., 2024; You et al., 2024). Industry standards bodies including 3GPP and IEEE are developing specifications for CF-mMIMO, with initial features expected to be included in 6G standards (3GPP, 2023; IEEE, 2024). The emergence of open radio access network (O-RAN) architectures that separate hardware and software and enable multi-vendor deployments is particularly well-suited for CF-mMIMO, potentially accelerating commercial adoption by reducing vendor lock-in and enabling innovation (O-RAN Alliance, 2023; Checko et al., 2023).

This thesis focuses on system-level performance evaluation of CF-mMIMO under ideal synchronization and calibration assumptions. Practical deployment aspects such as fronthaul capacity, compression, and synchronization protocols are not explicitly simulated at the signal or packet level. Instead, fronthaul and central processing costs are incorporated through fixed power consumption terms, and computational delay is approximated using simplified latency models for channel estimation, beam training, and precoding (Björnson et al., 2023; Interdonato et al., 2023). Hardware impairments, synchronization errors, and field deployment constraints are reviewed to provide practical context, but their detailed modeling is outside the scope of the implemented simulation framework (Shariatmadari et al., 2024; You et al., 2024).

CHAPTER 3

Methodology

3.1Introduction

This chapter presents the MATLAB-based methodology used to evaluate contamination-aware channel estimation and distributed precoding in a Cell-Free massive MIMO network under user mobility conditions. The methodology directly follows the simulation framework developed in this thesis.

Furthermore, the uplink and downlink signal-processing operations implemented in the simulation are described, including distributed combining and precoding across cooperating access points.

Finally, the simulation workflow, algorithm execution procedure, and performance evaluation metrics obtained through Monte Carlo simulations are described. All simulations are executed in MATLAB R2023b, leveraging parallel computing to support large-scale system evaluation while maintaining computational efficiency and reproducibility.

3.2 Cell-Free Massive MIMO System Model

3.2.1 Network Architecture

The simulation framework models a Cell-Free massive MIMO network deployed over a square area of 1000 m × 1000 m. The system consists of L = 64 distributed access points (APs) serving K = 40 user equipments (UEs). Each AP is equipped with M = 8 antennas, resulting in a total of L×M = 512 antennas in the network. The wireless channel is assumed constant within a coherence block of τ_c = 500 channel uses.

The AP positions are generated using a hierarchical deployment strategy that creates both dense and sparse regions within the service area. The AP locations are generated once and remain fixed during each simulation run.

Users are grouped around randomly selected cluster centers to represent traffic concentration. This clustering is implemented through a deterministic procedure in the MATLAB function deploy_clustered_users and does not follow stochastic geometry models such as Neyman–Scott or Thomas processes. in this function User equipment's are generated using a clustered spatial distribution. The simulation area is divided into a fixed number of clusters(20 % each), and each cluster is assigned a predefined fraction of the total users.

Cluster centers are randomly placed within the service area. Around each center, user positions are generated using a Gaussian distribution with a standard deviation proportional to the area size. This creates spatially concentrated user groups while allowing variability within each cluster. The generated positions are constrained to remain within the simulation area boundaries.

To ensure that the total number of users remains equal to K, any remaining users that are not assigned through the clustering process are placed randomly across the area.

All APs are connected to a central processing unit (CPU) that performs channel processing and coordination. The CPU collects channel state information (CSI), applies precoding and combining, and evaluates system performance. Fronthaul links are not explicitly modeled, and ideal information exchange between APs and the CPU is assumed.

The system operates in time-division duplex (TDD) mode. Uplink–downlink channel reciprocity is used (Björnson et al., 2017), and downlink precoding is based on uplink channel estimates.

The coherence block length is commonly defined as (Tse and Viswanath, 2005; Björnson et al., 2017):

$$\tau_c = B_c T_c \quad \textbf{(3.1)}$$

where $\boldsymbol{B_c}$ denotes the coherence bandwidth and $\boldsymbol{T_c}$ represents the coherence time.

In this work, the coherence block length is not explicitly computed from the coherence bandwidth and coherence time. Instead, it is set to $\boldsymbol{\tau_c}$=500 channel uses in the MATLAB implementation to represent slowly varying channels under low user mobility.

User mobility is explicitly modeled with an average speed of 3 km/h and position updates every 0.5 seconds. Channel variations are introduced through a temporal correlation model, where channel coefficients evolve over time instead of being regenerated independently at each step.

Within each coherence block, $\tau_p = 20$ channel uses are allocated to pilot transmission, and the remaining $\tau_c - \tau_p$ channel uses are used for data transmission. This structure is directly applied in the simulation and determines the overhead factor used in spectral efficiency calculations.

3.2.2 Dual-band operation

The simulation framework supports dual-band operation by modelling simultaneous transmission opportunities in both the sub-6 GHz and millimetre-wave frequency bands. This configuration exploits complementary propagation characteristics to improve system robustness and achievable data rates under varying channel conditions (Björnson et al., 2023; Tataria et al., 2024). The sub-6 GHz band operates at 3.5 GHz and provides reliable coverage due to lower propagation loss and stronger penetration, which is beneficial for users in non-line-of-sight conditions or located farther from serving APs (Rappaport et al., 2023; Dang et al., 2023). In contrast, the millimetre-wave band operates at 28 GHz and offers larger available bandwidth and higher spatial resolution, enabling higher spectral efficiency under favourable propagation conditions (Rappaport et al., 2023; Tataria et al., 2024).

In the MATLAB simulation, the two frequency bands are explicitly modelled through band-dependent large-scale fading and channel generation. For each AP–user pair, separate large-scale fading coefficients are computed for sub-6 GHz and millimetre-wave frequencies using different path loss exponents and shadow fading parameters. Specifically, the path loss exponent is lower for sub-6 GHz and higher for millimetre-wave, reflecting increased attenuation at higher frequencies. The small-scale fading is generated as spatially correlated

Rayleigh fading, and separate channel realizations are constructed for each frequency band using the corresponding large-scale fading coefficients and spatial correlation matrices (Björnson et al., 2023; Interdonato et al., 2023).

Channel estimation and data transmission procedures are performed independently for each frequency band using the corresponding channel realizations. The resulting performance metrics, including SINR and spectral efficiency, are computed per band and combined to evaluate overall system performance under heterogeneous propagation conditions.

3.2.3 AP Selection and User-Centric Clustering

a user-centric clustering approach is applied (Ngo et al., 2017; Björnson and Sanguinetti, 2020) , where each user is served by a limited subset of APs instead of the entire network. For each user k, a serving set of S=5 APs is selected based on large-scale fading conditions. The parameter S determines how many APs jointly serve each user and directly controls the trade-off between performance and system complexity.

Increasing S improves performance because more APs contribute to signal transmission, which enhances macro-diversity gain, increases received signal power, and reduces the probability of deep fading. In addition, multiple serving APs improve interference averaging and provide more robust connectivity under mobility and blockage conditions (Ngo et al., 2017; Interdonato et al., 2019).

The serving set of user k is defined as

$$\mathcal{M}_k\{l: \beta_{lk} \in top - s\{\beta_{1k}, \beta_{2k}, \dots, \beta_{lk}\}\} \quad \textbf{(3.2)}$$

where β_{lk} denotes the large-scale fading coefficient between AP l and user k. The operator "top-s" selects the S largest values among all AP–user links, meaning that the APs with the

strongest large-scale channel conditions are chosen (Ngo et al., 2017; Björnson and Sanguinetti, 2020) .

In the MATLAB implementation, large-scale fading coefficients are computed for all AP-user pairs, and the indices corresponding to the S-k strongest values are selected. This procedure is repeated when user positions change due to mobility, which leads to updates in the serving sets over time.

Frequent updates of serving sets introduce additional overhead, especially under user mobility. When users move, large-scale fading coefficients vary, which may change the ranking of APs and trigger re-selection. For slow mobility, such as pedestrian speed, these updates occur gradually and the overhead remains limited. For higher mobility, more frequent updates would be required, increasing signaling and computational overhead.

To avoid excessive load on specific APs, a load-balancing constraint is applied by limiting the number of users served by each AP to Kmax=5. This parameter is defined as a system-level constraint to ensure that processing and resource allocation at each AP remain bounded. If an AP exceeds this limit, users with weaker large-scale fading toward that AP are reassigned to alternative APs among their next strongest candidates.

This reassignment is implemented during serving set formation and ensures that no AP becomes overloaded while maintaining acceptable channel quality for all users.

3.3 Channel Model

This section describes the channel modelling framework implemented in MATLAB for simulating wireless propagation between distributed access points and mobile user equipment.

The channel between AP l and user k is modeled as a spatially correlated Rayleigh fading channel (Björnson et al., 2017; Interdonato et al., 2019)with large-scale fading and temporal evolution, expressed as

$$\boldsymbol{h_{lk}(t) \sim CN(0, \beta_{lk}, R_l)} \quad \textbf{(3.3)}$$

where h_{lk} is the M × 1 channel vector represents the large-scale fading coefficient, and R_l denotes the spatial correlation is the M × M spatial correlation matrix at AP l. This formulation captures both path loss and shadow fading through β_{lk}, and antenna correlation through R_l.

The channel model consists of three components: large-scale fading, small-scale fading, and temporal evolution due to user mobility. These components are jointly modeled within the same framework and directly determine the channel realization used in the simulation.

Large-scale fading is computed for each AP–UE pair based on distance-dependent path loss and shadow fading. Small-scale fading is generated as correlated Rayleigh fading using the spatial correlation matrices. Temporal evolution is introduced through user mobility, where channel coefficients evolve over time according to a correlation model rather than being independently regenerated at each step.

In the implemented framework, channel realizations are updated according to mobility-induced variations. When user positions change, large-scale fading coefficients are recomputed, and small-scale fading evolves based on temporal correlation. Therefore, the channel is not fully regenerated at every time step, but instead evolves continuously to reflect realistic propagation dynamics.

This unified channel model ensures consistency between large-scale fading, spatial correlation, and temporal evolution, and it is used throughout the simulation for AP selection, channel estimation, and precoding.

3.3.1 Large-Scale Fading

Large-scale fading captures distance-dependent attenuation and shadow fading between AP l and user k. In the simulation, propagation loss is modelled using standardized models defined in 3GPP TR 38.901 and widely adopted measurement-based studies for sub-6 GHz and millimetre-wave frequency bands (3GPP, 2019; Rappaport et al., 2013; Rappaport et al., 2023).

For each AP–UE pair, the path loss in decibel scale is expressed as

$$\boldsymbol{PL(d_{lk}) = 10\,\alpha\,\log_{10}(d_{lk}) + C + X_{\sigma}} \quad \textbf{(3.4)}$$

where d_{lk} is the distance between AP l and user k, alpha is the path-loss exponent, c is a frequency-dependent constant, and X_{σ} ~ N(0,σ^2) represents log-normal shadow fading (3GPP, 2019).

Since channel generation requires linear-scale channel gains, the large-scale fading coefficient $\boldsymbol{\beta_{lk}}$ is obtained from the path loss as

$$\boldsymbol{\beta_{lk} = 10^{PL(d_{lk})/10}} \quad \textbf{(3.5)}$$

where is P $L(d_{lk})$ expressed in decibels. The negative exponent reflects that path loss represents attenuation, and therefore the channel gain is inversely proportional to the loss.

In the MATLAB implementation, Equation (3.5) directly determines the large-scale channel gain used in channel generation, AP selection, and power allocation. Larger values of β_{lk} correspond to stronger links and therefore increase the likelihood of selecting AP l in the serving set defined in Equation (3.2). When user positions change due to mobility, distances d_{lk} are updated and large-scale fading coefficients are recomputed, which leads to dynamic adaptation of serving AP selection and channel conditions

3.3.2 Small-Scale Fading

Small-scale fading models multipath propagation and causes rapid variations in channel amplitude and phase. In the simulation, the channel vector between AP l, equipped with N antennas, and user k is modeled as

$$\boldsymbol{h_{lk}} = \boldsymbol{R_{lk}^{1/2}} \boldsymbol{w_{lk}} \tag{3.6}$$

where R_{lk} is the spatial correlation matrix and $\boldsymbol{w_{lk}} \sim \mathrm{CN}(\mathbf{0}, \mathbf{I}_N)$ represents independent Rayleigh fading samples (Björnson et al., 2018; Interdonato et al., 2023).

The matrix $R_{lk}^{1/2}$ introduces spatial correlation across antenna elements, transforming independent fading samples into correlated channel realizations. This formulation is consistent with standard correlated MIMO channel models used in massive MIMO systems.

In the MATLAB implementation, Equation (3.6) is used to generate channel vectors for each AP–UE pair within each coherence interval. These channel realizations are then used in channel estimation, uplink combining, and downlink precoding.

Spatial correlation is incorporated through the correlation matrices used in channel generation. In the simulation, these matrices are constructed and provided to the channel generation function, which applies them to the independent fading vectors. This approach enables modelling of antenna correlation effects while maintaining computational efficiency in large-scale Monte Carlo simulations.

3.3.3 Temporal Channel Correlation and Mobility Model

User mobility causes the channel to evolve over time as user positions change relative to the access points. In the simulation, temporal correlation between consecutive channel realizations is modeled using a first-order autoregressive process inspired by Jakes' fading model (Jakes & Cox, 1994):

$$\boldsymbol{h_{lk}(t) = \rho h_{lk}(t-1) + \sqrt{1-\rho^2}\, e_{lk}(t)} \quad \textbf{(3.7)}$$

where $e_{lk}(t) \sim \mathcal{CN}(0, R_{lk})$ represents independent fading innovation, and ρ ∈ [0,1] controls the temporal correlation between consecutive channel samples (Jakes & Cox, 1994; Tse and Viswanath, 2005) .

The parameter ρ determines how fast the channel varies over time. Values of ρ close to 1 correspond to slowly varying channels, while smaller values indicate faster channel variations. In the simulation, ρ is selected to represent low-mobility conditions, consistent with pedestrian user speeds (approximately 3 km/h), resulting in high temporal correlation between consecutive channel realizations.

Equation (3.7) models the evolution of small-scale fading only. Large-scale fading is updated separately based on user mobility. When user positions change, distances $d_{lk}(t)$ are recomputed, and the large-scale fading coefficients $\beta_{lk}(t)$are updated accordingly. The complete channel at time t is therefore given by

$$\boldsymbol{h_{lk}(t) \sim CN\,(0, \beta_{lk}(t) R_{lk})} \quad \textbf{(3.8)}$$

which combines both time-varying large-scale fading and temporally correlated small-scale fading (Björnson et al., 2017; Tse and Viswanath, 2005) .

In the MATLAB implementation, time evolution is discretized into simulation steps with a duration of 0.5 seconds. At each step, user positions are updated according to the mobility model, and the channel is updated using Equation (3.8). Therefore, channel realizations evolve gradually over time rather than being regenerated independently.

User mobility is modeled using a random waypoint model, where users move with a speed of approximately 3 km/h and randomly selected directions. This movement changes the distances

between APs and users, which affects large-scale fading, while small-scale fading evolves through temporal correlation.

3.4 Uplink Pilot Transmission and Channel Estimation

This section describes the uplink pilot transmission and channel estimation methodology implemented in the MATLAB simulation framework for the considered Cell-Free massive MIMO system. Channel estimation is required to obtain channel state information (CSI) at distributed access points, enabling coherent processing for uplink detection and downlink precoding (Ngo et al., 2017; Björnson et al., 2017; Björnson et al., 2023).

Because pilot resources are limited by the coherence block length τc , the number of orthogonal pilot sequences τp is limited. When the number of users K exceeds τp , pilot reuse becomes necessary, which leads to pilot contamination and degrades channel estimation accuracy (Ngo et al., 2017; Interdonato et al., 2023). To address this, both baseline LMMSE channel estimation and a contamination-aware estimation method are implemented in the simulation.

During uplink training, all users transmit pilot sequences simultaneously. The received pilot signal at AP l, equipped with N antennas, is given by

$$\boldsymbol{Y}_{\boldsymbol{l}}^{\boldsymbol{p}} = \sqrt{\boldsymbol{\tau}_{\boldsymbol{p}}\boldsymbol{p}_{\boldsymbol{p}}} \sum_{\boldsymbol{k}=\boldsymbol{1}}^{\boldsymbol{K}} \boldsymbol{h}_{\boldsymbol{lk}} \boldsymbol{\phi}_{\boldsymbol{\pi}(\boldsymbol{k})}^{\boldsymbol{T}} + \boldsymbol{N}_{\boldsymbol{l}}^{\boldsymbol{p}}, \tag{3.9}$$

where:

- Y_l^p is the received pilot signal matrix at AP l.
- $\boldsymbol{\tau_p}$ is the pilot sequence length.
- $\boldsymbol{p_p}$ is the pilot transmit power.
- $\boldsymbol{h_{lk}}$ is the channel vector between AP l and user k, defined as:

$$\boldsymbol{hlk} \sim \boldsymbol{CN}(\mathbf{0}, \boldsymbol{\beta lk Rlk}) \quad \textbf{(3.10)}$$

which includes both large-scale fading β_{lk} and spatial correlation R_{lk} (Björnson et al., 2017; Björnson et al., 2023)

- $\phi_{\pi}(k)$ is the pilot sequence assigned to user k, selected from a set of orthogonal sequences
- π(k) ∈ {1,…, τ_p} denotes the pilot index assigned to user k
- $\boldsymbol{N_l^p}$ is the additive noise matrix, with independent columns distributed as

$$\boldsymbol{CN}(\mathbf{0}, \boldsymbol{\sigma^2 I_N}) \quad \textbf{(3.11)}$$

The pilot sequences satisfy the orthogonality condition

$$\boldsymbol{\phi_i^H \phi_j} = \begin{cases} \boldsymbol{\tau_p} & \boldsymbol{i = j} \\ \mathbf{0}, & \boldsymbol{i \neq j} \end{cases} \quad \textbf{(3.12)}$$

which ensures that users with different pilots can be separated in the absence of noise and interference.

Equation (3.9) shows that the received pilot signal is a superposition of contributions from all users. When multiple users share the same pilot index $\pi(k)$, their channel vectors are projected onto the same pilot sequence, and their contributions become inseparable. This results in pilot contamination, which directly affects the accuracy of channel estimation (Ngo et al., 2017; Björnson et al., 2017).

In the MATLAB implementation, this model is used to construct the received pilot signal at each AP, which is then processed to estimate channel coefficients. The presence of pilot reuse

leads to interference in the estimation process, and this effect is explicitly evaluated through the implemented estimation methods.

3.4.2 Spatial Pilot Assignment Strategy

To mitigate pilot contamination, pilot sequences are assigned such that users sharing the same pilot are geographically separated. Since large-scale fading decreases with distance, spatial separation reduces mutual interference among co-pilot users (Ngo et al., 2017; Björnson et al., 2017; Interdonato et al., 2023).

For each user, the pilot is selected to maximize the minimum distance to users already assigned to that pilot. Let $\boldsymbol{p_k}$denote the spatial position of user k (not transmit power). The minimum distance between user k and users assigned to pilot t is defined as

$$\boldsymbol{d_{min}(k,t) = min\left\|p_k - p_{\acute{k}}\right\|} \quad \textbf{(3.13)}$$

where the minimization is taken over all users k ′ that are already assigned to pilot t, and ‖·‖ denotes Euclidean distance.

The pilot assignment is then given by

$$\boldsymbol{\pi(k) = argmax\, d_{min}(k,t)} \quad \textbf{(3.14)}$$

where π(k) is the pilot index assigned to user k, and the operator argmax returns the value of t that maximizes the minimum distance.

These expressions define the pilot assignment rule by selecting, for each user, the pilot that maximizes spatial separation from other users sharing the same pilot. As a result, users reusing

the same pilot are located farther apart, which reduces the overlap of their large-scale fading contributions at the APs and therefore decreases the level of pilot contamination in the received signal (Ngo et al., 2017; Interdonato et al., 2023).

3.4.3Baseline LMMSE Channel Estimation

After receiving pilot signals, AP l extracts the observation associated with user k by correlating the received pilot matrix with the pilot sequence assigned to that user,

$$\boldsymbol{y}_{\boldsymbol{lk}}^{\boldsymbol{p}} = \boldsymbol{Y}_{\boldsymbol{l}}^{\boldsymbol{p}} \boldsymbol{\phi}_{\boldsymbol{\pi}(\boldsymbol{k})}^{*} \quad \textbf{(3.15)}$$

This operation removes contributions from users employing different pilots due to pilot orthogonality. However, when pilots are reused, the resulting observation still contains channel components from all users sharing the same pilot sequence. Consequently, the signal used for channel estimation contains both the desired channel and interference from co-pilot users, in addition to thermal noise.

To estimate the desired channel, the Linear Minimum Mean Square Error (LMMSE) estimator is employed, which minimizes the mean square error Kay, 1993) between the true channel and its estimate by exploiting statistical knowledge of channel covariance and noise power. The resulting estimate is

$$\widehat{\boldsymbol{h}}_{\boldsymbol{lk}} = \sqrt{\boldsymbol{\tau_p p_p}} \boldsymbol{\beta}_{\boldsymbol{lk}} \boldsymbol{R}_{\boldsymbol{lk}} \boldsymbol{\psi}_{\boldsymbol{lk}}^{-\mathbf{1}} \boldsymbol{y}_{\boldsymbol{lk}}^{\boldsymbol{p}} \quad (3.16)$$

This formulation follows the classical LMMSE channel estimation framework (Ngo et al., IEEE JSAC, 2017; Björnson et al., IEEE TWC, 2020). The estimator is obtained by minimizing the mean square error $\mathrm{E}\{\| \boldsymbol{h_{lk}} - \widehat{\boldsymbol{h}}_{\boldsymbol{lk}} \|^{\mathbf{2}}\}$ over all linear estimators. Using the orthogonality principle (Kay, 1993), the optimal solution is expressed as a linear function of the observation, where the estimation filter depends on the cross-covariance between the channel and the

received pilot signal and the auto-covariance of the observation. Substituting these covariance expressions leads directly to (3.16).

Equation (3.16) shows that the estimator applies a linear filter to the received signal in order to suppress noise and interference while preserving the desired channel component. The term $\beta_{lk}R_{lk}$ reflects the spatial statistics and average channel strength of the desired user.

$$\boldsymbol{\psi_{lk} = \tau_p p_p \sum_{j \in \mathcal{P}_k} R_{lj} + \sigma^2 I_N} \quad \textbf{(3.17)}$$

The matrix $\boldsymbol{\psi_{lk}}$ represents the total covariance of the received pilot signal at AP l. and Where $\mathcal{P}_k$ is the set of users sharing the same pilot It includes contributions from the desired channel, interference from users sharing the same pilot sequence, and thermal noise (Björnson et al., 2017; Kay, 1993) .

The matrix inversion balances noise suppression and interference mitigation; however, when co-pilot users exhibit strong large-scale fading coefficients, contamination remains significant and reduces estimation accuracy.

Therefore, although LMMSE estimation minimizes the mean square error under statistical assumptions, its performance degrades in pilot reuse scenarios with strong interference, which motivates the contamination-aware estimation approach introduced in the next subsection.

3.4.4 Contamination-Aware Channel Estimation

While the baseline LMMSE estimator minimizes estimation error under known channel statistics, it cannot fully eliminate interference originating from users sharing the same pilot sequence. When co-pilot users exhibit comparable or strong channel gains, their contributions remain embedded in the received pilot observation, thereby contaminating the channel estimates.

The proposed approach consists of two complementary stages: successive interference cancellation (SIC) (Chen et al., 2019) for co-pilot users and contamination-aware weighting. In the first stage, the estimator iteratively detects and subtracts the dominant interference components associated with users sharing the same pilot sequence. This process reduces the contamination present in the received observation before channel estimation. Users are effectively processed based on their contamination severity, where stronger interfering users have a larger impact on the received pilot signal and are therefore implicitly prioritized during interference reconstruction and subtraction.

In the second stage, contamination-aware weighting is applied to the refined observation in order to suppress the residual interference (Chen et al., 2019) according to its estimated power while preserving the desired channel component. This weighting mechanism adapts the estimation process to the interference conditions and improves robustness in pilot reuse scenarios.

3.4.5 Iterative Interference Cancellation

Starting from initial LMMSE channel estimates, the contribution of interfering users sharing the same pilot is reconstructed and subtracted from the received pilot observation. This procedure relies on the fact that the initial estimates, although contaminated, already provide approximate channel information that can be exploited to estimate the interference components.

Let p_k denote the set of users assigned to the same pilot as user k. Due to pilot reuse, the received pilot signal at AP l contains the superposition of all users in $p_k$which leads to pilot contamination.

This superposition can be interpreted as a linear combination of channel vectors weighted by pilot power, which directly results in correlated interference that cannot be separated without additional processing .

At iteration i, the cleaned pilot observation for user k at AP is computed as

$$\boldsymbol{y_{lk}^{clean} = y_{lk}^{p} - \sqrt{\tau_p p_p} \sum_{\acute{k} \in p_k \backslash \{k\}} h_{l\acute{k}}^{(i-1)}} \quad \textbf{(3.18)}$$

where:

- $\boldsymbol{y_{lk}^{p}}$ is the received pilot signal vector.
- $\boldsymbol{\tau_p}$ is the pilot length.
- $\boldsymbol{p_p}$ is the pilot transmit power.
- $\boldsymbol{h_{l\acute{k}}^{(i-1)}} \in \mathbb{C}^{M \times 1}$ is the channel estimate of interfering user k' obtained at iteration $i-1$.
- The summation is taken over all users $\acute{k} \in p_k \backslash \{k\}$, all co-pilot users except the desired user

This subtraction removes the reconstructed interference contributions associated with users sharing the same pilot, resulting in a cleaned observation that is increasingly dominated by the desired channel component (Chen et al., 2019) .

After obtaining the cleaned observation, the channel estimate is recomputed by reapplying the LMMSE estimator:

$$\boldsymbol{h_{lk}^{(i)} = R_{lk} \Psi_{lk}^{-1} y_{lk}^{clean,(i)}} \quad \textbf{(3.19)}$$

where:

- R_{lk} is the channel covariance matrix including large-scale fading
- Ψ_{lk} is the covariance of the received pilot signal, including interference and noise

Here, the covariance matrix Ψ_{lk} incorporates both signal and interference statistics, meaning that the LMMSE estimator inherently balances noise suppression and interference mitigation through covariance inversion (Kay, 1993; Björnson et al., 2017) .

This shows that LMMSE estimation is reapplied at each iteration, but using a progressively refined observation.

The procedure is repeated for a fixed number of iterations:

$$i = 1,2, \dots, I$$

where I denotes the total number of refinement steps. In this work, a fixed number of I = 6 iterations is used, which provides a good trade-off between estimation accuracy and computational complexity. Increasing the number of iterations improves interference suppression but also increases computational cost and may lead to diminishing returns beyond a certain point.

As iterations progress, the quality of the channel estimates improves, leading to more accurate interference reconstruction and more effective subtraction in subsequent iterations. Consequently, the residual pilot contamination decreases at each iteration. This iterative improvement can be interpreted as a refinement loop where estimation and interference cancellation reinforce each other.

3.4.6 Contamination-Aware Weighting

Although iterative interference cancellation significantly reduces pilot contamination, residual interference may still remain due to estimation errors, limited number of SIC iterations, or closely located co-pilot users. This residual contamination originates from the uplink pilot transmission phase, where users sharing the same pilot sequence generate coherent interference in the received pilot signal at the access points. Even after LMMSE estimation and SIC, imperfect cancellation leads to residual interference components in the refined channel estimates. To further mitigate this residual contamination, a contamination-aware weighting mechanism is applied to the refined channel estimates obtained after the iterative LMMSE + SIC process.

The contamination level experienced by user k at AP l is quantified using the ratio between the large-scale fading of interfering users and the desired user:

$$\boldsymbol{\gamma_{lk}} = \frac{\sum_{\acute{k} \in p_k \backslash \{k\}} \boldsymbol{\beta_{l\acute{k}}}}{\boldsymbol{\beta_{lk}}} \tag{3.20}$$

where:

- β_{lk} is the large-scale fading coefficient of the desired user
- $\beta_{l\acute{k}}$ corresponds to users sharing the same pilot
- the summation is taken over all co-pilot users $k' \in p_k \backslash \{k\}$
- p_k denotes the set of users assigned to the same pilot sequence as user k

Since large-scale fading represents the average channel power (including path loss and shadowing), this ratio approximates the relative interference power compared to the desired signal power. It can therefore be interpreted as a simplified measure of the interference-to-signal covariance ratio used in LMMSE estimation (Ngo et al., 2017; Björnson et al., 2017) .

This ratio represents the relative strength of interfering users compared to the desired user. Therefore, it defines the contamination condition:

- $\boldsymbol{\gamma_{lk}} \ll 1$: low contamination
- $\boldsymbol{\gamma_{lk}} \gg 1$: strong contamination

A contamination-aware weight is then defined as

$$\boldsymbol{w_{lk}^{est}} = \frac{\mathbf{1}}{\mathbf{1} + \boldsymbol{\alpha}\, \boldsymbol{\gamma_{lk}}} \tag{3.21}$$

where α controls the aggressiveness of contamination suppression. In this work, $\alpha \in [0.1,10]$, where smaller values result in mild attenuation and larger values lead to stronger suppression of contaminated links (Chen et al., 2019) .

LMMSE regularization modifies the covariance matrix during inversion to improve numerical stability and interference suppression, the contamination-aware weight directly scales the estimated channel based on its contamination level.

The final channel estimate is obtained by applying this weight to the refined channel after the iterative LMMSE + SIC procedure:

$$\widehat{\boldsymbol{h}}_{\boldsymbol{lk}}^{\boldsymbol{final}} = \boldsymbol{w_{lk}^{est}}\, \widehat{\boldsymbol{h}}_{\boldsymbol{lk}}^{\boldsymbol{I}} \tag{3.22}$$

where $\hat{h}_{lk}^{I}$ denotes the channel estimate after I iterations of interference cancellation and LMMSE re-estimation (Chen et al., 2019; Björnson et al., 2017).

This formulation ensures consistency with the overall estimation pipeline. The weighting is not applied to the initial LMMSE estimate, but to the final refined estimate after interference cancellation.

The role of the contamination-aware weighting is to attenuate channel estimates that are still affected by residual contamination while preserving reliable estimates. As a result, links with high contamination are down-weighted, reducing their impact on subsequent processing such as distributed precoding and combining.

3.5 Downlink Precoding and Data Transmission

During the downlink phase, multiple distributed access points (APs) jointly transmit data to the users they serve. Each AP applies precoding to its transmitted signals based on local channel estimates. The purpose of downlink precoding is to improve the strength of the desired signal at each user while reducing inter-user interference.

The precoding vectors are designed from the estimated channels and are normalized to unit norm. This normalization ensures that the magnitude of the precoding vector does not directly determine the transmit power.

Although a common scaling factor is used in the signal model, the transmit power is not equal for all users. Instead, a power allocation mechanism assigns different power levels to different AP user links. This allocation depends on channel conditions, fairness requirements, and system constraints such as per-AP power limits.

The overall performance of the downlink transmission depends on three main components: the accuracy of channel estimation, the structure of the precoding vectors, and the power allocation strategy. These components jointly determine the achievable SINR and system throughput.

3.5.1 Downlink Signal Formation

During the downlink phase, each AP transmits a weighted superposition of data symbols intended for the users it serves. The transmitted signal at AP l is

$$\boldsymbol{x_l} = \sqrt{\boldsymbol{\rho_d}} \sum_{\boldsymbol{k \epsilon D_l}} \boldsymbol{w_{lk}}^{\boldsymbol{prec}} \boldsymbol{s_k} \tag{3.23}$$

where $\boldsymbol{w_{lk}}^{prec}$ is the precoding vector computed from the estimated channel, $\boldsymbol{s_k}$ is the data symbol intended for user k, and ρ_d is a nominal downlink power scaling factor (Ngo et al., 2017; Tse and Viswanath, 2005).

The precoding vectors are normalized to unit norm $\|\boldsymbol{w_{lk}}^{prec}\|^2 = 1$, and therefore they are not unitary but normalized vectors. This normalization ensures that the transmit power is determined by the power allocation coefficients rather than by the applied precoder norm.

The data symbols $\boldsymbol{s_k}$ represent modulated information symbols. Their statistical energy is set to, $E\{|\boldsymbol{s_k}|^2\}$ unless stated otherwise. The effective transmit power is controlled by the scaling factor and the subsequent power allocation mechanism.

Although $\boldsymbol{\rho_d}$ appears as a common scaling factor, the transmit power is **not equal for all signals** in the implemented system. The effective transmit power associated with each AP–user link is determined by a dedicated power allocation procedure, which distributes power based on channel conditions, fairness constraints, and per-AP power limits. As a result, different users and links are assigned different transmit power levels.

The received signal at user k is given by

$$\boldsymbol{y_k} = \sum_{\boldsymbol{l \epsilon \mathcal{M}_k}} \boldsymbol{h_{kl}^H x_l} + \boldsymbol{n_k} \tag{3.24}$$

where $\boldsymbol{h_{lk}}$ denotes the channel vector between AP l and user k, $(\cdot)^{\mathrm{H}}$represents the Hermitian transpose, and $n_k \sim \mathcal{CN}\ (0, \sigma^2)$ is additive receiver noise (Tse and Viswanath, 2005).

The Hermitian transpose appears because the received signal is formed through the inner product between the channel vector and the transmitted signal vector. This operation determines how the transmitted signal is projected onto the user channel.

By substituting (3.23) into (3.24), the received signal can be decomposed into desired signal, inter-user interference, and noise components. The effective channel gain for each link is determined by the term $|h_{lk}^H \boldsymbol{w_{lk}}^{prec}|^2$, which captures the alignment between the channel and the precoding vector.

The system performance is evaluated using the signal-to-interference-plus-noise ratio (SINR), defined as(Tse and Viswanath, 2005; Ngo et al., 2017):

$$\boldsymbol{SINR} = \frac{\boldsymbol{desired\ signal}}{\boldsymbol{interference + noise}} \tag{3.25}$$

In the simulation, the desired and interference powers are computed using the effective channel gains $|h_{lk}^H \boldsymbol{w_{lk}}^{prec}|^2$, combined with the corresponding power allocation coefficients, while the noise power is determined by the receiver noise variance.

3.5.2 Baseline Precoding Strategies

Baseline precoding schemes are implemented to evaluate the performance improvements achieved by the proposed method. These schemes follow standard formulations in multi-user MIMO and Cell-Free massive MIMO systems (Ngo et al., 2017; Björnson et al., 2020).

Let the global channel matrix be defined as

$$\mathbf{H} \in \mathbb{C}^{\mathbf{K}\times\mathbf{M_{tot}}} \tag{3.26}$$

where:

- K is the number of users
- Mtot = L×M is the total number of antennas across all APs
- each row of **H** corresponds to the channel of one user across all AP antennas

Thus, **H** represents the network-wide channel matrix used for centralized precoding (Björnson et al., 2017; Ngo et al., 2017).

Maximum Ratio (MR) Precoding

In Maximum Ratio (MR) precoding, each AP aligns its transmission with the estimated channel:

$$\boldsymbol{w}_{\boldsymbol{lk}}^{\boldsymbol{MR}} = \frac{\widehat{\boldsymbol{h}}_{\boldsymbol{lk}}}{\left\|\widehat{\boldsymbol{h}}_{\boldsymbol{lk}}\right\|} \tag{3.27}$$

where $\hat{h}_{lk}$ is the estimated channel between AP l and user k. (Ngo et al., IEEE TWC, 2017).

This operation is performed locally at each AP, since it only requires local channel state information. MR precoding maximizes the desired signal power but does not suppress inter-

user interference. As the number of users increases, interference becomes dominant, limiting system performance.

Zero-Forcing (ZF) Precoding

Zero-Forcing (ZF) precoding removes inter-user interference by computing

$$\boldsymbol{w^{ZF} = H^H (HH^H)^{-1}} \quad \textbf{(3.28)}$$

where:

- $\boldsymbol{w^{ZF}}$ is the precoding matrix
- $\boldsymbol{H^H}$ denotes the Hermitian transpose of H

The matrix $\boldsymbol{HH^H}$ represents the user correlation matrix. Diagonal elements correspond to channel power, while off-diagonal elements quantify inter-user interference. Low correlation enables efficient user separation, whereas high correlation reduces the effectiveness of interference suppression.

ZF precoding is computed centrally at the CPU, since it requires global channel knowledge and matrix inversion. The resulting precoding vectors are then distributed to the APs.

While ZF ideally cancels inter-user interference, matrix inversion amplifies noise when channels are correlated or when the system operates at low SNR, reducing robustness (Björnson et al., *Massive MIMO Networks*, Foundations and Trends, 2017).

Centralized MMSE Precoding

Centralized MMSE precoding introduces regularization:

$$\boldsymbol{w^{c-MMSE} = H^H(HH^H + \alpha_c I)^{-1}} \quad \textbf{(3.29)}$$

where:

- $\boldsymbol{\alpha_c}$ is the regularization parameter
- I is the Identity matrix

The regularization parameter is defined as

$$\boldsymbol{\alpha_c = \frac{\sigma^2}{P}} \quad \textbf{(3.30)}$$

where σ^2 is the noise power and P is the transmit power.

This parameter balances interference suppression and noise amplification. When $\alpha_c \rightarrow 0$, the solution approaches ZF precoding. For larger values of α_c, the solution becomes more robust to noise.

Similar to ZF, centralized MMSE precoding is computed at the CPU due to the requirement for global channel knowledge and matrix inversion. Although effective, this method introduces high computational complexity, especially in large-scale networks (Björnson et al., *Making Cell-Free Massive MIMO Competitive*, IEEE TWC, 2020).

3.5.3 Proposed Adaptive Distributed MMSE Precoding

The proposed methodology eliminates centralized matrix inversion by performing precoding computations locally at each access point (AP) using only locally available channel estimates. This distributed implementation significantly reduces fronthaul signaling requirements and computational complexity while maintaining effective interference suppression. Such

distributed processing strategies are widely studied in Cell-Free massive MIMO systems (Ngo et al., 2017; Björnson et al., 2020).

Let D_l denote the set of users served by AP l. Each AP constructs a local channel matrix

$$H_l = \left[\hat{h}_{l1}, \hat{h}_{l2}, \dots, \hat{h}_{lK}\right]$$

where $\hat{h}_{lK}$ is the estimated channel matrix between AP l and user k .

Based on this local information, the precoding vector for user k at AP l is computed as

$$\boldsymbol{w_{lk}}^{\boldsymbol{D-MMSE}} = \left(\sum \widehat{\boldsymbol{h}}_{\boldsymbol{lk}}\widehat{\boldsymbol{h}}_{\boldsymbol{lk}}^{\boldsymbol{H}} + \boldsymbol{\alpha_l I_N}\right)^{-1} \widehat{\boldsymbol{h}}_{\boldsymbol{lk}} \quad \textbf{(3.31)}$$

This expression corresponds to the local MMSE precoding solution, obtained by minimizing the mean squared error using local channel state information (Björnson et al., 2017, Chapter 4; Kay, 1993). The same formulation is widely used in distributed and user-centric Cell-Free massive MIMO implementations (Björnson et al., 2020).

The computation in (3.21) is performed entirely at each AP without requiring centralized coordination.

The regularization parameter $\boldsymbol{\alpha_l}$ controls the trade-off between interference suppression and noise amplification and is derived from the MMSE criterion (Kay, 1993). It depends on the noise power σ^2, defined as

$$\boldsymbol{\sigma^2} = \boldsymbol{10^{\frac{N_0}{10}-3}} \quad \textbf{(3.32)}$$

where N_0 represents the noise power in dBm, including thermal noise, system bandwidth, and receiver noise figure (3GPP TR 38.901, 2019).

To maintain robustness under varying propagation conditions, the regularization strength is adapted based on the local signal-to-noise ratio (SNR) at each AP. The local SNR at AP l is estimated as (Tse and Viswanath, 2005):

$$\boldsymbol{SNR_L = 10\log(\frac{\rho_d \sum_{k\epsilon D_l} \|h_{lk}\|^2}{|D_l|\sigma^2})} \quad (3.33)$$

where:

- $\boldsymbol{\rho_d}$ is the transmit power per data stream
- $\boldsymbol{h_{lk}}$ is the channel between AP l and user k
- $\boldsymbol{D_l}$ is the number of users served by AP l

This expression follows the standard definition of SNR as the ratio between average received signal power and noise power (Kay, 1993; Tse and Viswanath, 2005). It provides a practical local approximation of channel quality using only available channel norms.

Based on the estimated SNR, the regularization parameter is adapted:

- at low SNR, a larger $\boldsymbol{\alpha_l}$ is used to prevent noise amplification
- at high SNR, a smaller $\boldsymbol{\alpha_l}$ is used to enhance interference suppression

This SNR-dependent regularization is consistent with MMSE design principles, where the regularization term reflects the relative strength of noise and signal (Björnson et al., 2017).

3.5.4 Power Allocation Methodology

Power allocation determines how the available transmit power is distributed among users in the network and directly affects signal strength, interference, and fairness (Tse and Viswanath, 2005; Björnson et al., 2017).

P_{total} Denotes the total transmit power of the entire system, i.e., the sum of transmit power across all access points (APs), measured in Watts. This total power is equally divided among the L APs, where K the number of APs. The power budget of each AP $l \in \{1,\ldots,L\}$ is therefore

$$\boldsymbol{P_l = \frac{P_{total}}{k}} \quad (3.34)$$

Equal Power Allocation

As a baseline, equal power allocation is considered. In this case, each AP distributes its available power equally among the users it serves. The power allocated from AP l to user k is

$$\boldsymbol{p_{lk} = \frac{P_l}{|D_l|}} \quad \textbf{(3.35)}$$

where:

- p_{lk} is the transmit power from AP l to user k, measured in Watts
- D_l is the set of users served by AP l
- $|D_l|$ is the number of users served by AP l
- K is the total number of users in the system

This scheme does not depend on channel conditions and serves as a reference configuration (Tse and Viswanath, 2005).

Optimized Power Allocation

In the proposed method, power allocation is performed per AP based on the effective channel gain between AP l and user k, defined as

$$\boldsymbol{g_{lk} = \left|h_{lk}^H w_{lk}^{prec}\right|^2} \quad \textbf{(3.36)}$$

where:

- g_{lk} is the effective channel gain.
- h_{lk} is the channel vector between AP l and user k
- w_{lk}^{prec} is the precoding vector used by AP l for user k
- M is the number of antennas per AP

The quantity g_{lk} represents the received signal power contribution from AP l to user k (Björnson et al., 2017).

For each AP l, power is allocated only to active users defined as $D_l = \{k \mid \|h_{lk}\| > \epsilon \; and \; \|w_{lk}{}^{prec}\| > \epsilon\}$ where ϵ is a small threshold used to exclude negligible channel or precoding contributions. This ensures that power is assigned only to effectively served users.

For each AP l, power allocation is performed in two steps:

1. A minimum power level is assigned to all active users k ∈ $\mathcal{D}_l$
2. The remaining power is distributed proportionally to $\sqrt{g_{lk}}$

The allocation satisfies the per-AP constraint

$$\sum_{k \in \mathcal{D}_l} \boldsymbol{p_{lk}} \leq \boldsymbol{P_l} \tag{3.37}$$

and the total system constraint

$$\sum_{l=1}^{L} \sum_{k=1}^{K} \boldsymbol{p_{lk}} \leq \boldsymbol{P_{total}} \tag{3.38}$$

These constraints ensure that the total transmitted power does not exceed the available budget at both local and global levels (Tse and Wiswanatio,2005).

To prevent user starvation, a minimum power level is reserved for each active user:

$$\boldsymbol{P_{min}} = \frac{\boldsymbol{P_l}}{\boldsymbol{10 \times K}} \tag{3.39}$$

If sufficient power is available,

$$|\boldsymbol{D_l}|.\boldsymbol{P_{min}} \leq \boldsymbol{P_l} \tag{3.40}$$

the remaining power is distributed according to a proportional-fairness-inspired weighting:

$$\boldsymbol{P_{lk}} = \boldsymbol{P_{min}} + \frac{\sqrt{\boldsymbol{g_{lk}}}}{\sum_{\boldsymbol{i} \in \boldsymbol{D_l}} \sqrt{\boldsymbol{g_{li}}}} (\boldsymbol{P_l} - |\boldsymbol{D_l}|.\boldsymbol{P_{min}}) \tag{3.41}$$

This formulation allocates more power to users with stronger effective channel gains while maintaining fairness through minimum guarantees and softened weighting (Kelly, 1997).

If the minimum power requirement exceeds the AP budget,

$$|\boldsymbol{D_l}|.\boldsymbol{P_{min}} \geq \boldsymbol{P_l} \tag{3.42}$$

power is distributed equally among active users:

$$\boldsymbol{P_{lk}} = \frac{\boldsymbol{P_l}}{|\boldsymbol{D_l}|} \tag{3.43}$$

This ensures feasibility under constrained conditions (Tse and Viswanath, 2005).

If the effective channel gains are numerically negligible, equal power allocation is applied to maintain numerical stability.

Fairness Evaluation

Fairness among users is quantified using Jain's fairness index:

$$F = \frac{\left(\sum_{k=1}^{K} p_k\right)^2}{k \sum_{k=1}^{K} {p_k}^2} \quad (3.44)$$

where:

- F∈(0,1] is the fairness index
- p_k is the total power received by user k, defined as

$$p_k = \sum_{l=1}^{L} p_{lk} \quad (3.45)$$

- k∈{1,…,K} is the user index
- k is the total number of users

A value of F=1 indicates perfectly equal power distribution among users, while lower values indicate unequal allocation(Jain et al., 1984).

3.6 Performance Metrics and Analysis Methodology

This section outlines the methodology for evaluating system performance under the implemented channel estimation and distributed precoding framework. The performance metrics considered quantify signal quality, achievable data rates, energy efficiency, and the system's mobility-related behavior. These metrics are computed after downlink data transmission using the estimated channels and precoding strategies described in the previous sections, and therefore directly reflect the combined effects of channel estimation accuracy, serving AP selection, and distributed transmission processing. In the MATLAB simulation, performance is evaluated across multiple channel realizations and mobility states using Monte Carlo iterations, ensuring that the results represent statistically reliable system behaviour rather than outcomes from a single scenario.

The received signal quality for each user is quantified by the Signal-to-Interference-plus-Noise Ratio (SINR), which measures the ratio between the useful received signal power and the combined effect of multi-user interference and receiver noise (Tse and Viswanath, 2005; Björnson et al., 2017).

For user k, the instantaneous SINR is computed as

$$\boldsymbol{SINR_k} = \frac{\left|\sum_{l\epsilon\mathcal{M}_k} \boldsymbol{h_{lk}^H w_{lk}}^{\boldsymbol{prec}}\right|^2 \boldsymbol{\rho_k}}{\sum_{k\neq\acute{k}} \left|\sum_{l\epsilon\mathcal{M}_{\acute{k}}} \boldsymbol{h_{lk}^H w_{l\acute{k}}}^{\boldsymbol{prec}}\right|^2 \boldsymbol{\rho_{\acute{k}}} + \boldsymbol{\sigma^2}} \quad \textbf{(3.46)}$$

The numerator represents the coherent combination of the desired signal contributions from all APs serving user k. The denominator includes interference from all other users k′≠k transmitting simultaneously in the system, including users sharing the same pilot as user k, as well as receiver noise.

This SINR expression follows the standard downlink multi-user MIMO signal model and is commonly used for performance evaluation in Cell-Free massive MIMO systems (Ngo et al., 2017; Björnson et al., 2020).

3.6.2 Achievable Sum Rate

User data rates are calculated from the measured SINR values, taking pilot transmission overhead into account. Because pilot transmission occupies part of the coherence block, only the remaining symbols are available for data transmission.

The achievable rate of user k is expressed as (Tse and Viswanath, 2005):

$$R_k = (1 - \frac{\tau_p}{\tau_c}) B \log_2(1 + SINR_K) \quad \textbf{(3.47)}$$

where the factor $(1 - \tau_p/\tau_c)$ represents the fraction of channel uses allocated to data transmission.

The total system throughput is then obtained as

$$R_{sum} = \sum_{k=1}^{K} R_k \tag{3.48}$$

This metric measures the aggregate data rate delivered by the distributed network and serves as the primary throughput performance indicator in the simulation results.

3.6.3 Energy Efficiency

To evaluate how efficiently the system utilizes energy for data transmission, energy efficiency is defined as the number of successfully transmitted information bits per unit of consumed energy (Tse and Viswanath, 2005; Björnson et al., 2017).

The energy efficiency of the system is computed as

$$\boldsymbol{EE = \frac{R_{sum}}{P_{tx} + P_{circuit}}} \tag{3.49}$$

where:

- EE is the energy efficiency in bits/Joule
- R_{sum} is the total achievable sum rate (bits/s/Hz)
- P_{tx} is the actual total transmit power consumed in the system
- $P_{circuit}$ is the total circuit power consumption

Transmit Power Clarification

The quantity P_{tx} represents the total allocated transmit power after power allocation, defined as

$$\sum_{l=1}^{L}\sum_{k=1}^{K} p_{lk} = P_{tx} \tag{3.50}$$

Thus, $P_{tx} \leq P_{total}$, where P_{total} is the maximum available power budget. Due to power optimization, the actual consumed transmit power may be lower than this maximum.

Circuit Power Consumption Model

The circuit power consumption is modeled as

$$P_{circuit} = L_{active}P_{AP} + P_{fronthaul} + P_{CPU} \tag{3.51}$$

where:

- L_{active} is the number of active APs serving users
- P_{AP} is the circuit power consumed per AP (including RF chains, signal processing, and hardware components)
- $P_{fronthaul}$ is the power consumed by fronthaul communication links
- P_{CPU} is the power consumed by centralized processing

Energy efficiency captures the trade-off between throughput and total power consumption. Increasing the sum rate without increasing total consumed power improves energy efficiency, while excessive power usage reduces it.

This definition is consistent with energy-efficiency formulations widely used in green wireless communication and massive MIMO systems (Björnson et al., 2015; Björnson et al., 2017).

3.6.4 Handover Management Metrics

User mobility requires dynamic adaptation of the serving AP set in order to maintain reliable communication as users move within the network. To evaluate mobility performance, handover statistics are defined to quantify both the frequency of AP re-association and the effectiveness of handover decisions (3GPP TR 38.300, 2020; 3GPP TR 38.901, 2019). These statistics include the total number of handover events, the average number of handovers per user, the handover rate, and the improvement in received signal strength after handover. The handover rate is defined as

$$\boldsymbol{\lambda_{HO}} = \frac{\boldsymbol{N_{HO}}}{\boldsymbol{KT_{sim}}} \tag{3.52}$$

where $\boldsymbol{N_{HO}}$ denotes the total number of handover events observed during the simulation, k is the total number of users, and $\boldsymbol{T_{sim}}$ is the simulation time in seconds. This metric represents the average number of serving AP changes per user per unit time. In the implemented simulation, a handover event is recorded whenever the serving AP set of a user is updated due to a sufficiently large improvement in received signal strength. Lower values of $\boldsymbol{\lambda_{HO}}$ indicate more stable connectivity and reduced signaling overhead, while higher values reflect more frequent reconfiguration due to user mobility. This definition is consistent with standard mobility performance evaluation in cellular systems (3GPP TR 38.300, 2020).

The improvement in received signal strength achieved during handover is quantified as

$$\boldsymbol{\Delta RSRP_{HO}} = \mathbf{10}\, \boldsymbol{log_{10}}\left(\frac{\boldsymbol{\beta_{l_{new},k}}}{\boldsymbol{\beta_{l_{old},k}}}\right) \tag{3.53}$$

where $\beta_{l,k}$ denotes the large-scale fading coefficient between AP l and user k, and $l_{new}\ and\ l_{old}$represent the serving AP before and after handover, respectively. This metric captures the gain in large-scale channel strength resulting from the handover decision. Positive values indicate that the new AP provides better channel conditions than the previous one, confirming the effectiveness of the handover mechanism.

The use of large-scale fading in (3.31) is justified by the fact that handover decisions are based on slowly varying channel components such as path loss and shadowing. In contrast, small-scale fading varies rapidly due to multipath effects and does not provide a stable basis for handover decisions. Therefore, practical systems rely on averaged received signal strength measurements when triggering handovers, which is consistent with standard cellular operation (3GPP TR 38.901, 2019).

In the MATLAB implementation, the received signal strength is computed from the average received power across antennas, and handover decisions are triggered when the gain from a candidate AP exceeds a predefined margin and satisfies a time-based stability condition. The serving AP set is then updated by selecting the strongest APs according to the measured signal strength.

Overall, the defined handover metrics provide a quantitative framework for evaluating mobility performance by capturing both the frequency of AP switching and the associated improvement in channel conditions. These metrics, together with SINR, throughput, and energy efficiency, enable a comprehensive assessment of system performance under user mobility.

3.7 Simulation Parameters and Implementation Methodology

This section presents the simulation configuration and implementation methodology used to evaluate the proposed contamination-aware channel estimation and distributed precoding framework. The configuration aims to reproduce realistic Cell-Free massive MIMO operating conditions while ensuring statistical reliability and computational feasibility.

3.7.1 System Parameters

The deployment under consideration comprises 64 distributed access points (APs) serving 40 users, with each AP equipped with eight antennas. Each user is dynamically served by a subset of five APs selected through the user-centric clustering mechanism described previously, which is a standard configuration in Cell-Free massive MIMO systems (Ngo et al., 2017;

Björnson et al., 2020). Pilot transmission employs sequences of length τp = 20 within a coherence interval of τ_c=500 channel uses. (Björnson et al., 2017; 3GPP TR 38.901, 2019).

Dual-band operation is considered, with carrier frequencies at 3.5 GHz and 28 GHz representing the sub-6 GHz and millimetre-wave regimes, respectively. These frequencies are aligned with practical 5G deployments, where sub-6 GHz bands provide reliable coverage and millimetre-wave bands offer high data rates (3GPP TR 38.901, 2019). The system bandwidth is set to 20 MHz, which is a standard allocation in cellular systems and enables a realistic evaluation of spectral efficiency and throughput (Tse and Viswanath, 2005).

Transmit power values and noise levels are selected according to typical cellular operating conditions. The transmit power is set to 23 dBm per user or AP transmission, corresponding to approximately 200 mW, which is consistent with small-cell and distributed antenna system deployments (3GPP TR 38.901, 2019).

The noise power is computed based on thermal noise and receiver characteristics as

$$\boldsymbol{N_0 = -174\frac{db}{Hz} + 10\log_{10}(B) + NF} \quad \textbf{(3.54)}$$

where B=20 MHz is the system bandwidth and NF is the receiver noise figure, resulting in a total noise power of approximately −94dBm. These values are widely used in wireless system evaluations (Tse and Viswanath, 2005; Björnson et al., 2017).

User mobility parameters are selected to represent pedestrian scenarios. The user velocity is set to approximately 3 km/h, which corresponds to typical walking speed, and the simulation duration is 60 seconds with 120 time steps, resulting in a time resolution of 0.5 seconds per update. (3GPP TR 38.901, 2019).

The number of successive interference cancellation (SIC) iterations is set to 6, which provides sufficient convergence of the channel estimation process while avoiding excessive computational complexity. This choice reflects a practical trade-off between estimation accuracy and implementation cost, as additional iterations yield diminishing performance improvements.

3.7.2 Monte Carlo Simulation Framework

System performance is evaluated using a Monte Carlo simulation (Tse and Viswanath, 2005) framework to capture variability arising from user locations, channel realizations, and mobility evolution. Multiple independent simulation trials are conducted, each corresponding to a distinct realization of the network topology and channel conditions.

For each trial, AP and user locations are generated, and large-scale fading parameters are computed based on distances and shadowing effects. Serving AP sets are then determined using the clustering and load-balancing procedures, after which spatial pilot assignment is performed. Channel realizations are subsequently generated by combining large-scale fading and small-scale fading components.

Both large-scale and small-scale fading evolve over time according to the user mobility model. Large-scale fading varies with user position due to changes in path loss and shadowing, while small-scale fading varies due to multipath propagation and temporal channel fluctuations. This joint evolution enables realistic time-varying channel conditions throughout the simulation.

Channel estimation procedures are then applied, followed by the computation of precoding vectors for both the baseline and proposed transmission schemes. System performance metrics, including SINR, throughput, and energy efficiency, are then evaluated at each time step.

To avoid bias from specific network layouts, the network topology is periodically regenerated during simulation runs.

3.7.3 Computational Implementation Strategy

Simulations are executed in a numerical computing environment optimized for large-scale linear algebra operations. Computational efficiency is achieved by parallelizing independent Monte Carlo trials, enabling simultaneous processing of multiple channel realizations.

Matrix operations in channel estimation and precoding computation rely on optimized linear algebra routines, ensuring stable numerical performance even for large antenna configurations. Parallel execution significantly reduces total runtime while preserving identical simulation behaviour compared to sequential execution.

The resulting computational configuration enables evaluation of large distributed networks within practical execution times while preserving methodological accuracy.

3.7.4 Validation Against Published Results

To ensure the correctness of the simulation framework, the obtained results are compared with established benchmarks reported in the Cell-Free massive MIMO literature. Channel estimation accuracy, interference mitigation performance, and achievable sum rate values obtained in the proposed framework fall within the ranges reported for similar network configurations.

This validation step confirms that the implemented channel generation, estimation, and precoding procedures produce performance levels that are realistic and consistent with previously published results, thereby increasing confidence in the conclusions drawn from the proposed methodology.

3.8 Chapter Summary

This chapter presented the methodology used to investigate contamination-aware channel estimation and distributed precoding in Cell-Free massive MIMO systems under user mobility conditions.

CHAPTER 4

RESULTS AND ANALYSIS

4.1 Introduction

This chapter presents the performance evaluation of the proposed contamination-aware channel estimation and distributed MMSE precoding framework for cell-free massive MIMO systems. The analysis covers key metrics including channel estimation accuracy, pilot contamination, SINR, spectral efficiency, energy efficiency, and mobility performance.

Results are obtained using realistic simulation settings and are compared with baseline schemes to assess performance improvements. Both single-realization and Monte Carlo analyses are considered to evaluate average performance and statistical reliability.

The chapter demonstrates how the proposed framework mitigates pilot contamination, improves signal quality, and maintains stable performance under mobility, while achieving high spectral and energy efficiency.

4.2 System Configuration and Simulation Parameters

This section summarizes the main simulation parameters used to generate the results. The system consists of 64 distributed access points, each equipped with 8 antennas, serving 40 users in a 1 km × 1 km area. Each user is served by 5 access points, ensuring sufficient spatial diversity.

A dual-band configuration is considered, with operation at 3.5 GHz (20 MHz bandwidth) and 28 GHz (400 MHz bandwidth). This setup enables the system to combine wide-area coverage with high data rate capability.

A total of 20 orthogonal pilot sequences are used, resulting in a pilot reuse factor of 2. Consequently, pilot contamination is present and must be mitigated through spatial pilot assignment and contamination-aware channel estimation.

Transmit power and noise parameters are selected to reflect practical deployment conditions, enabling realistic evaluation of system performance under interference-limited scenarios.

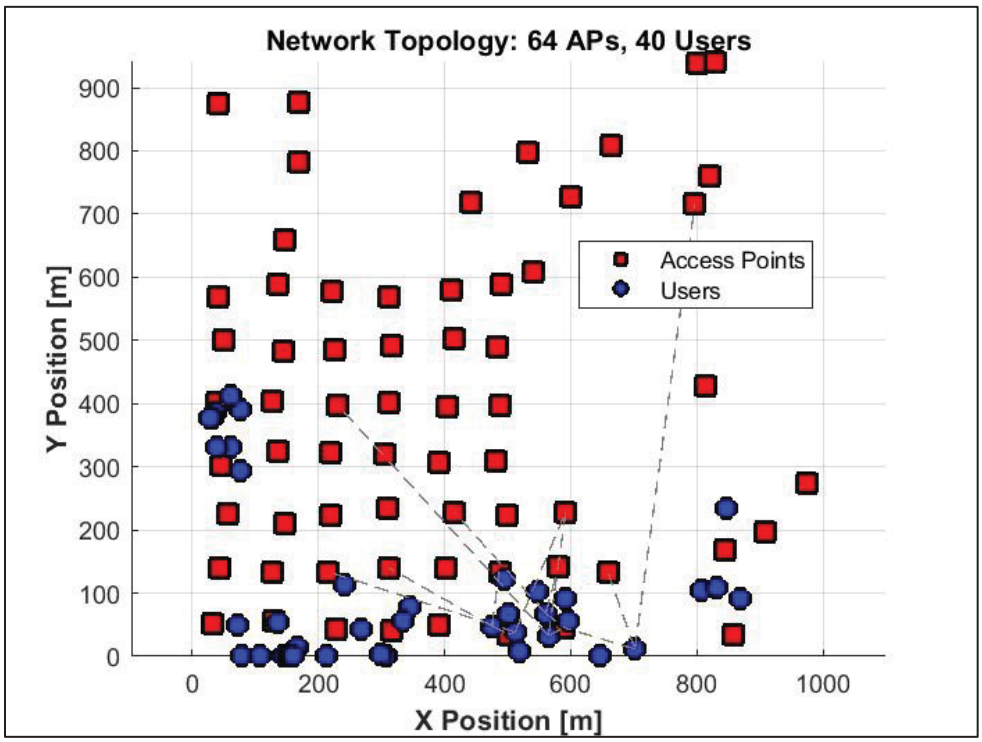

Figure 4.1 Network topology

Figure 4.1 illustrates the simulated network topology. Access points are distributed irregularly according to a uniform spatial distribution over the coverage area, rather than being placed on a regular grid.

Users are distributed according to a clustered spatial model, where groups of users are concentrated in specific regions of the coverage area. This behaviour represents realistic traffic patterns, such as user hotspots, and can be interpreted as a simplified Poisson cluster process.

4.3 Channel Estimation Performance and Pilot Contamination Analysis

Under pilot reuse factor 2, standard LMMSE estimation exhibits significant contamination. The median contamination-to-signal ratio is approximately 3.55 dB, indicating that interference from co-pilot users is comparable to or exceeds the desired signal power.

The corresponding NMSE under baseline estimation lies in the range of −11 to −12 dB, with strong spatial variability due to near–far effects and user geometry. Users sharing strong propagation paths with co-pilot users experience the highest estimation errors.

The proposed contamination-aware estimation significantly improves performance. The contamination level is reduced to −15.82 dB, corresponding to a mitigation gain of 19.37 dB. At the same time, the NMSE improves to −18.55 dB, indicating substantially more accurate channel estimates.

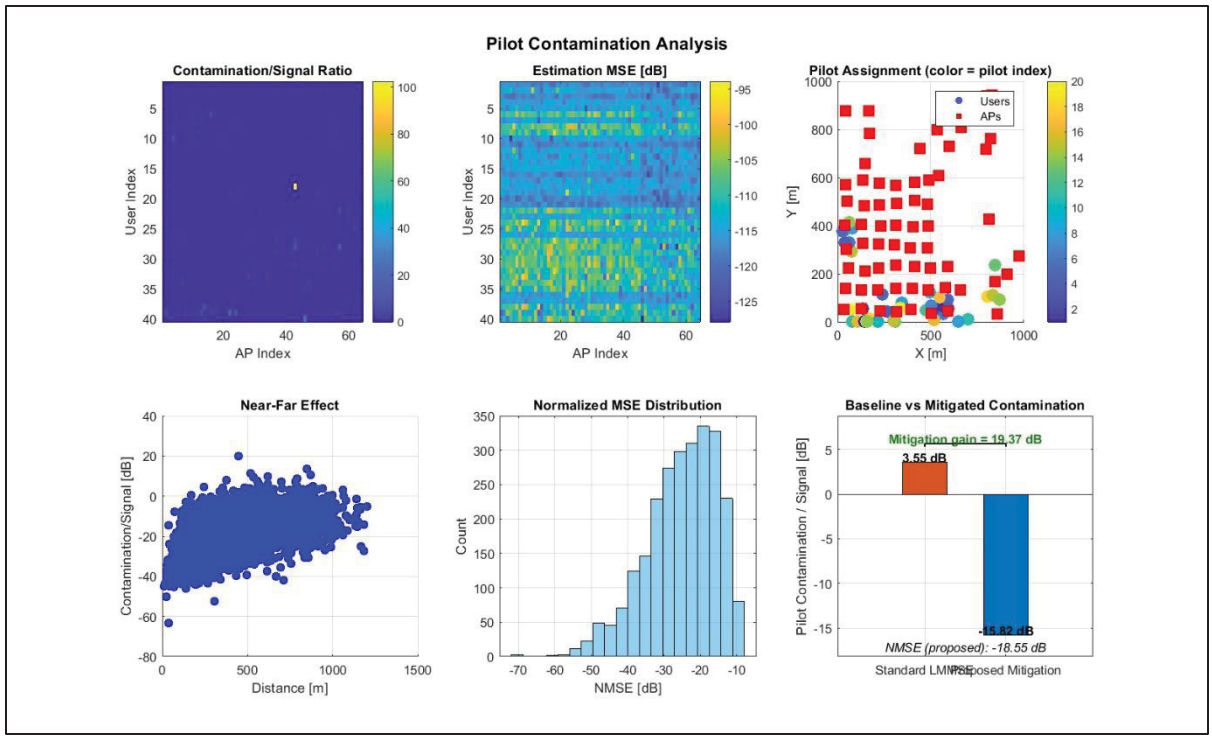

Figure 4.2: Comprehensive pilot contamination

The contamination heatmap shows that interference is concentrated in a limited number of user–AP links, while most links exhibit low contamination levels. This indicates that pilot contamination is strongly dependent on network geometry rather than uniformly distributed.

The estimation error map remains relatively uniform across users, with localized regions of higher error corresponding to strong interference conditions. This confirms that the proposed estimation maintains stable performance despite pilot reuse.

The pilot assignment map shows that users sharing the same pilot are spatially separated, which reduces coherent interference during channel estimation.

The near–far plot indicates a decreasing trend between contamination and distance, although significant dispersion remains due to shadowing and multi-user interference.

The NMSE distribution is tightly concentrated around −18.55 dB, confirming consistent estimation accuracy across the network.

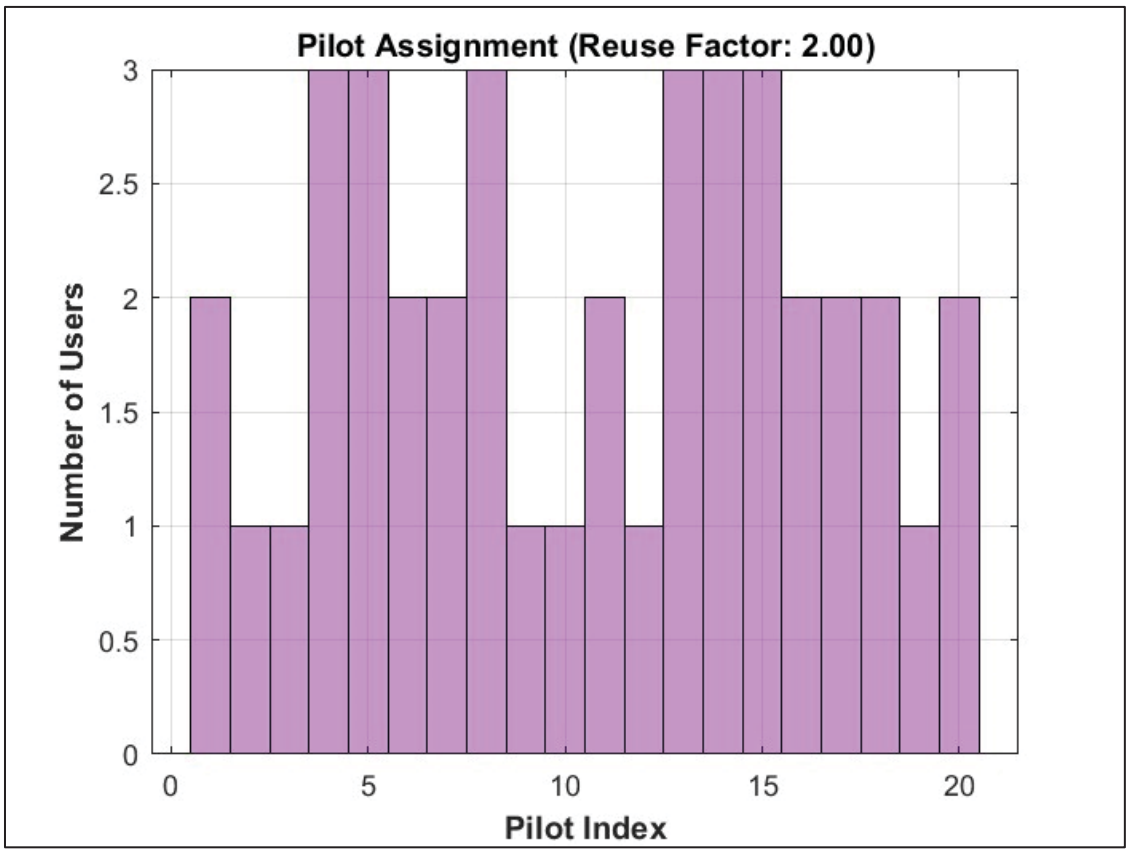

Figure 4.3 Spatial pilot assignment visualization

Figure 4.3 The figure shows the distribution of pilot indices among users for a reuse factor of 2. The histogram indicates that the number of users assigned to each pilot index is relatively uniform, with most pilots serving approximately one to three users.

No pilot index exhibits a significantly higher user count compared to others, which confirms that the pilot assignment algorithm distributes users evenly across the available pilot set. The absence of dominant peaks in the histogram demonstrates that pilot reuse is not concentrated on a limited subset of pilots.

This balanced distribution ensures that co-pilot users are spatially separated, as enforced by the distance-based pilot assignment strategy. As a result, the level of pilot contamination remains controlled and does not disproportionately affect specific users.

4.4 Precoding Performance Comparison

Precoding determines how distributed access points jointly transmit signals so that the desired signal is strengthened at the intended users while minimizing interference to other users. This section compares four precoding techniques implemented in the simulation: maximum ratio transmission (MR), zero-forcing (ZF), centralized minimum mean-square error (centralized MMSE), and the proposed distributed minimum mean-square error (distributed MMSE).

Performance is evaluated under two transmission policies: equal power allocation, in which all users receive the same transmit power, and optimized power allocation, in which power is allocated to improve overall system performance. Comparing both cases allows us to separate gains from interference management through precoding from those resulting from intelligent power control.

Unless otherwise stated, numerical results correspond to the same representative network realization, ensuring that all methods experience identical propagation conditions.

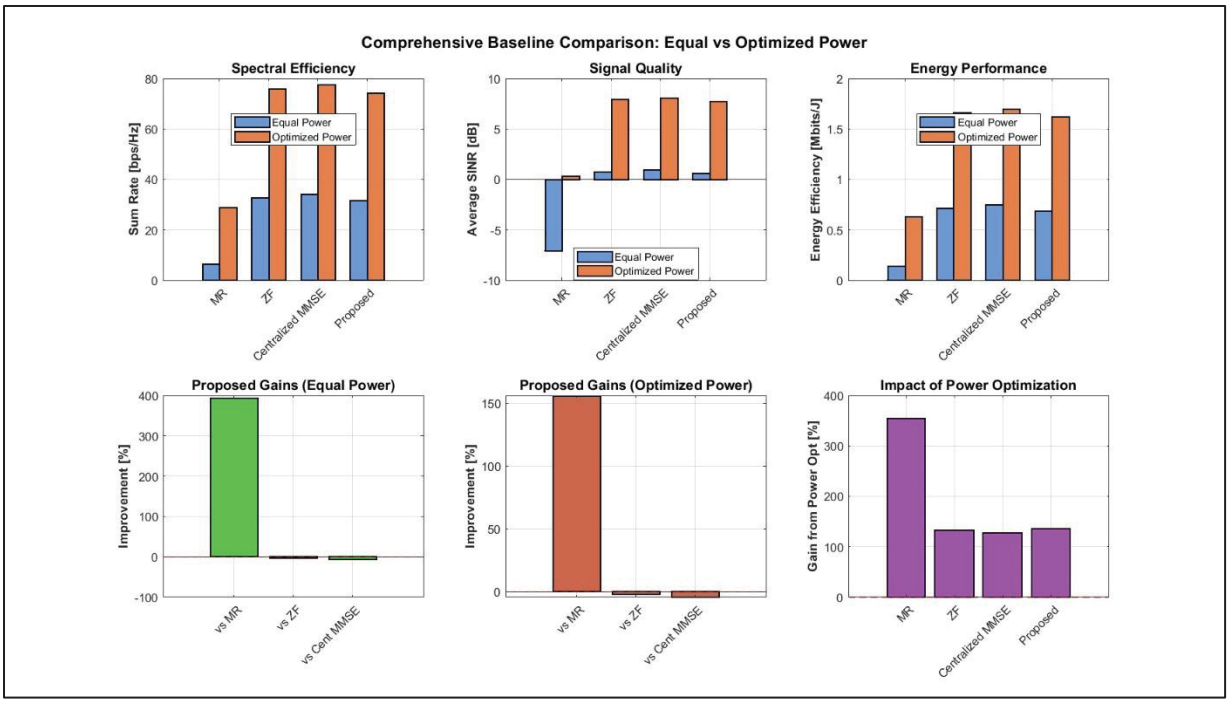

Figure 4.4 Comprehensive precoding comparison

The figure presents a comparative evaluation of MR, ZF, centralized MMSE, and the proposed method under equal power allocation and optimized power allocation.

Under equal power allocation, the proposed method achieves a sum rate of 31.59 bps/Hz, which is significantly higher than MR (6.39 bps/Hz) and close to ZF (32.77 bps/Hz) and centralized MMSE (34.20 bps/Hz). The corresponding average SINR for the proposed method is 0.59 dB, which remains limited due to the absence of power adaptation. The energy efficiency reaches 0.690 Mbits/J, indicating moderate performance in terms of power utilization.

When power optimization is applied, a substantial performance improvement is observed across all metrics. The proposed method achieves a sum rate of 74.32 bps/Hz, representing more than a twofold increase compared to the equal power case. The average SINR increases to 7.73 dB, confirming that power allocation effectively enhances signal quality by prioritizing stronger channel conditions. The energy efficiency improves to 1.625 Mbits/J, demonstrating a more efficient conversion of transmit power into useful throughput.

Compared to baseline schemes under optimized power, the proposed method significantly outperforms MR (29.02 bps/Hz) and remains close to ZF (76.01 bps/Hz) and centralized MMSE (77.60 bps/Hz). This indicates that the proposed distributed approach achieves near-centralized performance while maintaining lower complexity.

The gain analysis further quantifies these improvements. Under equal power, the proposed method provides approximately +280% improvement over MR, while showing comparable performance to ZF and centralized MMSE. Under optimized power, the gain over MR remains substantial at approximately +140%, confirming that the proposed method consistently benefits from improved channel utilization.

The impact of power optimization is particularly significant for MR, where performance increases by nearly 290%, highlighting its strong dependence on power control. For the proposed method, the gain from power optimization is approximately 150%, indicating that although the method already performs efficiently under equal power, it further benefits from adaptive allocation.

The proposed method effectively leverages this optimization to achieve high spectral efficiency, improved SINR, and increased energy efficiency, while maintaining performance close to centralized MMSE solutions.

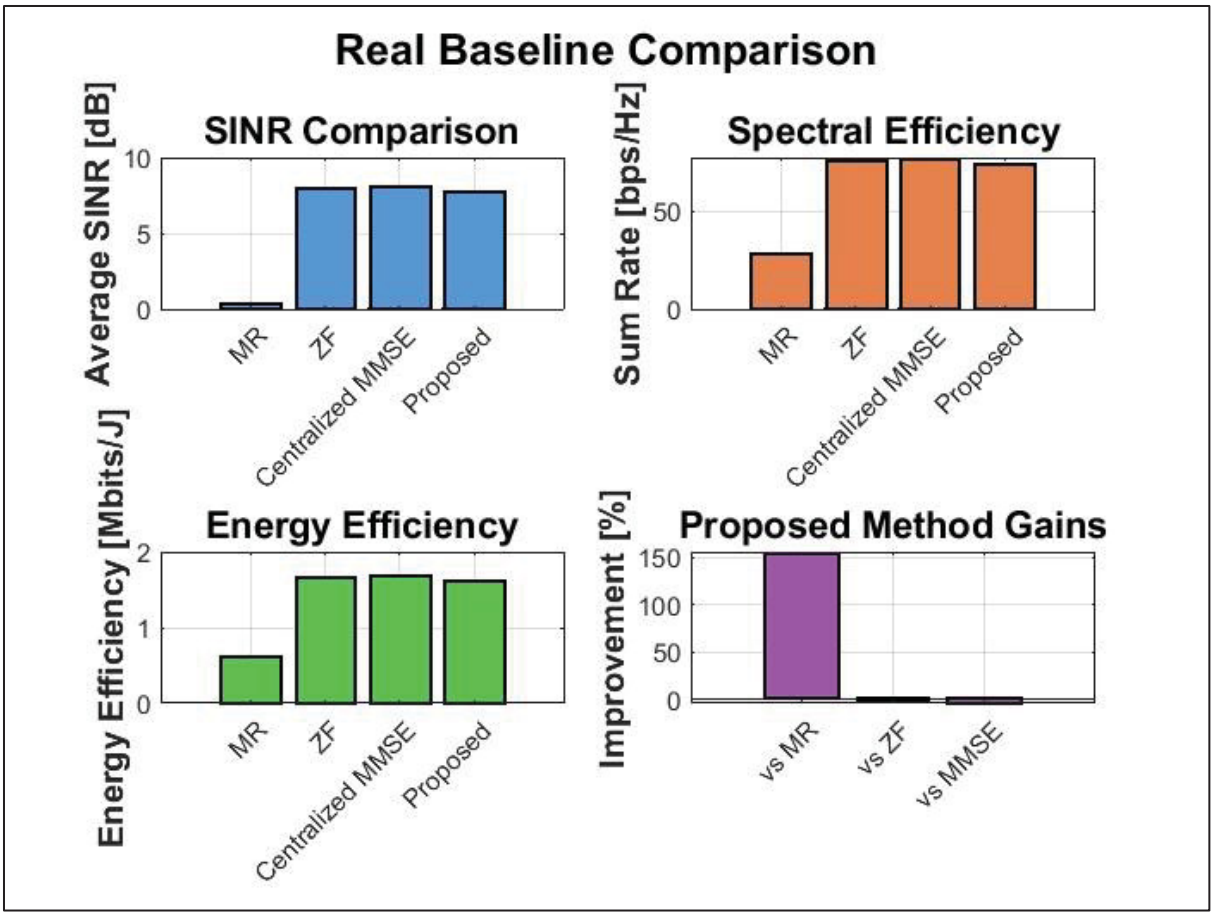

Figure 4.5 Detailed precoding performance comparison

Figure 4.5 presents a baseline comparison of precoding methods in terms of SINR, spectral efficiency, energy efficiency, and relative performance gains.

The top row shows the absolute performance. Maximum Ratio transmission achieves the lowest SINR and spectral efficiency due to the absence of interference suppression. In contrast, Zero-Forcing and centralized MMSE significantly improve signal quality by mitigating inter-user interference.

The bottom-right panel shows the relative performance gains of the proposed method. The largest improvement is observed with respect to Maximum Ratio transmission, while the gains relative to Zero-Forcing and centralized MMSE remain limited. This confirms that the

The reported performance metrics are computed as follows.

The signal-to-interference-plus-noise ratio (SINR) for each user is obtained by comparing the received desired signal power with the combined effect of interference from other users and thermal noise after precoding. The reported average SINR represents the mean value across all users in the network.

The spectral efficiency is determined based on the achievable data rate of each user as a function of its SINR. Higher SINR values correspond to higher achievable data rates. The total spectral efficiency is then obtained by summing the individual user rates over all users.

The energy efficiency is evaluated by comparing the total network throughput with the overall system power consumption. The total power includes not only the transmit power but also circuit power, fronthaul power, and processing power, reflecting realistic system operation.

The performance gains are calculated by comparing the proposed method with each baseline scheme in terms of achieved throughput. The improvement is expressed as a percentage, indicating how much the proposed method increases performance relative to each baseline.

4.5 Signal-to-Interference-Plus-Noise Ratio and Rate Analysis

This section evaluates user-level performance in terms of the signal-to-interference-plus-noise ratio (SINR) and achievable data rates achieved with the proposed distributed minimum mean square error precoding. While previous sections compared overall system throughput across different precoding strategies, the analysis here focuses on how performance is distributed among individual users. This evaluation is essential for assessing fairness, coverage quality, and service reliability across the network.

Simulation results show that the proposed distributed MMSE precoding achieves an average SINR of approximately 7.73 dB, providing acceptable signal quality for most users while maintaining strong overall system throughput. However, user performance is not uniform, as channel quality depends on user location, shadowing conditions, and the interference environment.

4.5.1 Signal-to-Interference-Plus-Noise Ratio Distribution

The SINR performance achieved with the proposed distributed minimum mean square error precoding exhibits moderate user-to-user variability, reflecting heterogeneous propagation conditions within the simulated network. According to the command-line output, the average SINR across all users is 7.73 dB, with a standard deviation of approximately 1.32 dB,

The SINR values range from approximately −3.42 dB for the weakest user to about 19.12 dB for the strongest user, yielding a dynamic range exceeding 22 dB across the network.

The distribution exhibits a slight positive skew, with a noticeable portion of users achieving SINR values above the average due to favourable geometry or proximity to several access points. This behaviour confirms that optimized power allocation concentrates resources where they yield higher spectral efficiency while still maintaining acceptable performance for users with weaker channels.

Only a small fraction of users experience SINR values below 0 dB, indicating that interference dominates signal power in rare cases.

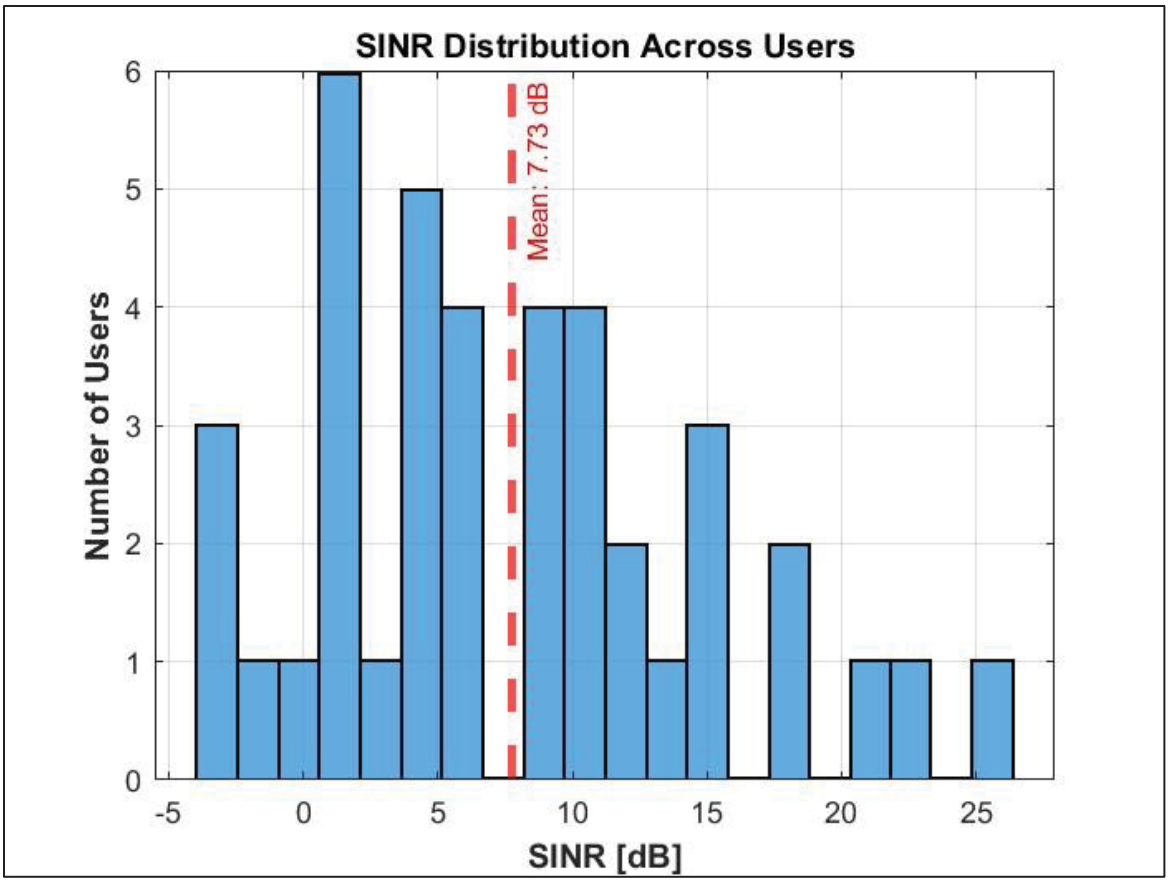

Figure 4.6 Signal-to-interference-plus-noise ratio distribution under MMSE precoding

Figure 4.6 presents a histogram of how SINR values are distributed across all 40 users.

Most users are clustered around SINR values between approximately 4 and 12 dB, confirming that the system provides relatively consistent signal quality across the network. A small number of users achieve significantly higher SINR values because of favourable locations near several access points or low interference. Conversely, a limited number of users appear in the low-SINR region due to unfavourable channel conditions.

The red dashed vertical line marks the mean SINR of 7.73 dB, which lies near the centre of the distribution.

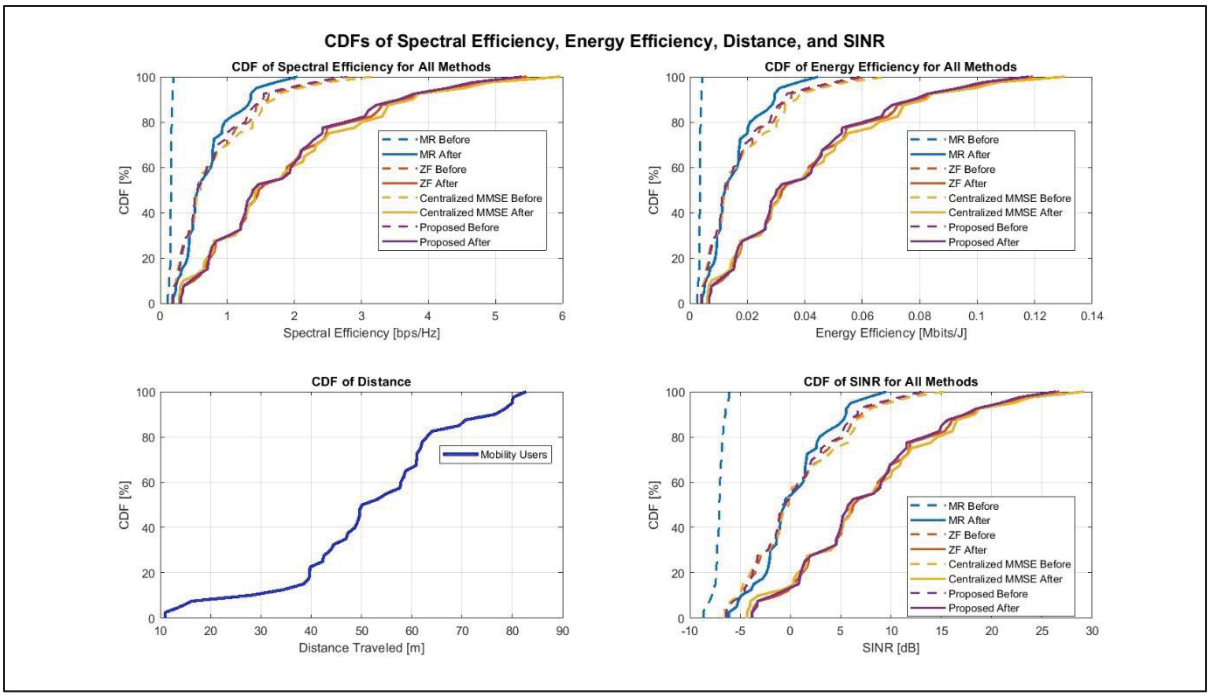

Figure 4.7 CDF of SE, EE, distance and SINR

Figure 4.7 presents the cumulative distribution functions (CDFs) of spectral efficiency, energy efficiency, distance travelled, and SINR for different transmission schemes before and after applying power optimization.

For the spectral efficiency, the CDF curves after power optimization are consistently shifted to the right for all methods, indicating higher achievable rates across users. The proposed method reaches approximately 3–5 bps/Hz for the majority of users (above 70% of the distribution), whereas MR remains concentrated below 2 bps/Hz even after optimization.

For the energy efficiency, a similar rightward shift is observed. After optimization, most users achieve values in the range of 0.06 to 0.12 Mbits/J for the proposed method, while MR remains below 0.05 Mbits/J for a large portion of users.

The SINR CDF shows a substantial improvement after optimization. For the proposed method, the median SINR increases to approximately 7–10 dB, and a large fraction of users achieve

SINR values above 10 dB. In contrast, MR exhibits a wide distribution with many users remaining in the negative or low SINR region, even after optimization.

The distance CDF reflects the mobility behavior of users. The distribution ranges approximately from 10 m to 80 m, with a median around 50 m, which is consistent with pedestrian mobility over the simulation duration. The smooth and continuous shape of the curve indicates uniform user movement without abrupt changes, ensuring that performance results are not biased by extreme mobility conditions.

The rightward shift of the curves after power optimization confirms that gains are experienced by the majority of users rather than a small subset, which indicates robustness of the proposed approach under realistic channel and mobility conditions.

4.5.2 Achievable Rate Distribution

The achievable rate distribution follows the SINR performance of the system. For the representative realization, the total spectral efficiency is 74.32 bps/Hz, corresponding to an average per-user spectral efficiency of 1.86 bps/Hz.

User rates are not uniform and depend on channel conditions, interference levels, and serving AP configuration. While the average performance is high, individual users experience different throughput levels due to spatial variability in large-scale fading and residual interference

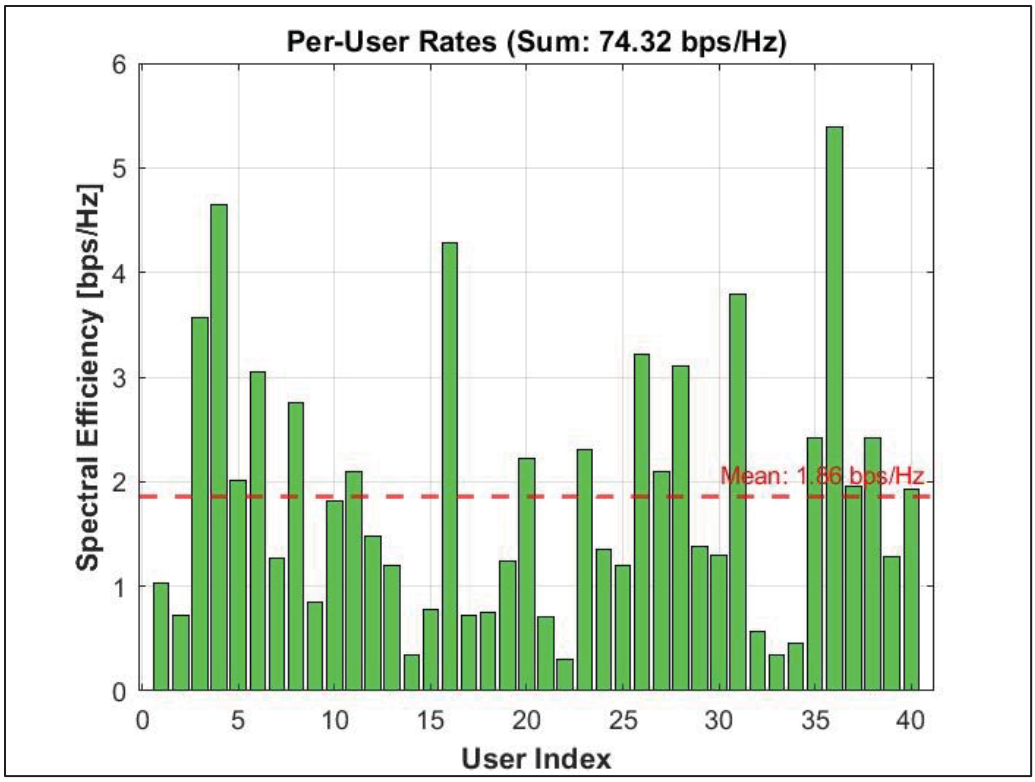

Figure 4.8 spectral efficiency achieved by each user

Figure 4.8 illustrates the spectral efficiency achieved by each user. The per-user rates range from approximately 0.3–0.5 bps/Hz for the weakest users to values exceeding 5 bps/Hz for users with favourable channel conditions.

The dashed horizontal line indicates the average spectral efficiency of 1.86 bps/Hz. A significant portion of users is distributed around this value, with most users achieving rates between 1 and 3 bps/Hz, confirming that the system provides consistent throughput for the majority of users.

Only a limited number of users fall below 1 bps/Hz, indicating that low-rate users are present but not dominant. At the same time, several users achieve rates above 3 bps/Hz, showing that the system effectively exploits favourable channel conditions through spatial diversity and power allocation.

This distribution indicates that the proposed power allocation does not concentrate resources on a small subset of users. Instead, it maintains a balanced rate distribution while still allowing high-rate users to benefit from strong channel conditions. The observed variation follows SINR differences, which are driven by user location, large-scale fading, and the composition of the serving AP set.

4.5.3 Relationship Between Channel Quality and Performance

Performance differences among users are strongly influenced by network geometry and channel conditions. Users near the centre of the coverage area typically benefit from multiple nearby access points, enabling stronger coherent signal combining and improved interference suppression. Conversely, users near network boundaries experience fewer strong serving links and therefore achieve lower SINR and throughput.

Simulation statistics indicate a negative correlation coefficient of approximately −0.43 between SINR and distance to the nearest access point, confirming that closer proximity to infrastructure generally improves performance. However, the moderate correlation value also shows that distance alone does not fully determine performance

The command-line output shows that the average signal-to-noise ratio, computed without considering interference, is 34.76 dB, with values ranging from 24.09 dB to 50 dB .

4.6 Mobility and Handover Performance

This section evaluates system performance under user mobility over a 60-second simulation with 120 channel updates (0.5 s interval). The analysis focuses on movement behaviour, handover frequency, and the impact of mobility on signal quality and achievable rate.

Simulation results show that users travel an average distance of approximately 52 meters, while maintaining stable connectivity throughout the network. Despite time-varying channel conditions, the system preserves consistent performance, indicating that the proposed distributed framework effectively adapts to mobility.

4.6.1 Mobility Model and Simulation Scenario

User movement follows heterogeneous patterns, with most users travelling between 40 and 70 meters, while a smaller number experience shorter or longer displacements. This variation leads to different channel evolution profiles across users.

As users move, large-scale fading and dominant serving access points change gradually. This results in continuous but smooth variations in channel quality rather than abrupt changes. The observed mobility patterns confirm that users experience realistic, time-varying propagation conditions throughout the simulation.

4.6.2 Handover Statistics and Performance

During the simulation, a total of 31 serving set updates are recorded across all users, corresponding to an average of 0.78 handovers per user, with a maximum of 3 handovers for highly mobile users.

Most users experience zero or one handover, indicating that serving sets remain stable for the majority of the simulation. Users travelling longer distances are more likely to encounter new dominant access points and therefore trigger additional updates.

Each handover results in an average signal strength improvement of approximately 19.88 dB, confirming that updates occur only when there is a significant gain in channel quality. The relatively low number of handovers compared to the total number of channel updates demonstrates that the system avoids unnecessary reconfiguration while still adapting effectively to mobility.

4.6.3 Signal Quality During Mobility

Under mobility, the average SINR is approximately 10.08 dB, with a standard deviation of 1.32 dB, indicating moderate temporal variation

The average achievable rate during mobility is 3.73 bps/Hz, showing that throughput remains stable over time with only minor fluctuations. Temporary SINR reductions occur during handover events, but these are brief and followed by rapid recovery as the serving set is updated.

The minimum SINR occasionally drops below −5 dB during challenging channel conditions, but no sustained degradation or outage is observed. The absence of prolonged low-SINR periods confirms that connectivity is preserved throughout the simulation.

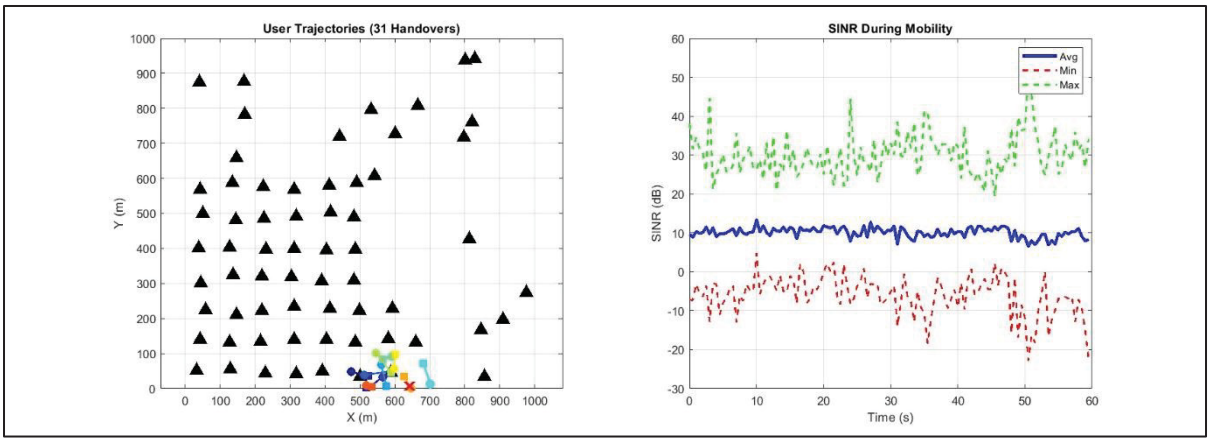

Figure 4.9 User trajectories

Figure 4.9 illustrates user trajectories over the 60-second simulation together with the spatial variation of SINR. Users move within localized regions rather than uniformly across the coverage area, consistent with the random waypoint model. The access point locations remain fixed, while user trajectories are concentrated mainly in the lower-central region of the network.

The colour variation along each trajectory reflects SINR evolution. SINR values change gradually along user paths, with no abrupt drops to outage levels. Temporary reductions in SINR occur when users move away from strong serving access points, but these reductions are short-lived and followed by recovery after serving set updates.

The absence of large SINR discontinuities confirms that the system maintains continuous connectivity during mobility. SINR variations are smooth and remain within a controlled range, demonstrating that channel estimation and serving set adaptation effectively track time-varying channel conditions.

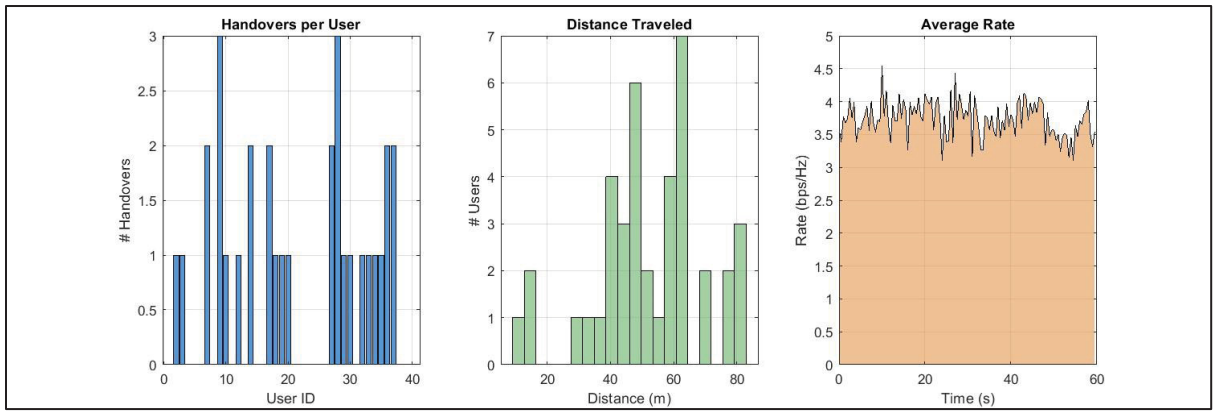

Figure 4.10 mobility and handover behavior

Figure 4.10 quantifies the impact of mobility on network adaptation through handover frequency, travelled distance, and time-varying rate.

The handover distribution shows that most users experience 0 to 1 handover, while a smaller subset experiences up to 3 handovers during the 60-second simulation. The average value is 0.78 handovers per user, confirming that serving set updates are infrequent and occur only when channel conditions change significantly.

The distance distribution indicates that most users travel between 40 and 70 meters, with an average distance of approximately 52 meters . Users travelling longer distances are more likely to encounter new dominant access points and therefore experience more handovers, establishing a direct relationship between mobility and serving set updates.

The average rate over time remains stable around 3.5–4 bps/Hz, with only minor fluctuations. No significant rate drops are observed during handover events, indicating that serving set adaptation does not introduce service interruption.

The network adapts efficiently to user movement, maintaining stable throughput while requiring only a small number of handovers.

4.7 Power Efficiency and Energy Analysis

This section evaluates how the proposed system converts total power consumption into achievable throughput, with emphasis on overall efficiency and comparison across transmission schemes.

The total system power consumption is 915 W, combining contributions from transmit, circuit, fronthaul, and processing components. Despite the large-scale deployment with 64 access points and multiple antennas, the radiated transmit power remains very low, indicating efficient beamforming and spatial focusing.

The proposed method achieves a sum spectral efficiency of 74.32 bps/Hz, corresponding to a total network throughput of approximately 1.49 Gbit/s over the system bandwidth.

Performance comparisons show that energy efficiency differences across precoding schemes are primarily driven by achievable throughput rather than variations in total power consumption, since hardware-related power remains dominant. Maximum ratio transmission achieves approximately 0.57 Mbit/J, while zero-forcing and centralized MMSE achieve approximately 1.36–1.37 Mbit/J.

The proposed distributed MMSE approach achieves comparable energy efficiency while avoiding the complexity and coordination overhead associated with centralized processing. Therefore, further improvements are expected to depend on reducing circuit and fronthaul consumption rather than increasing transmit efficiency.

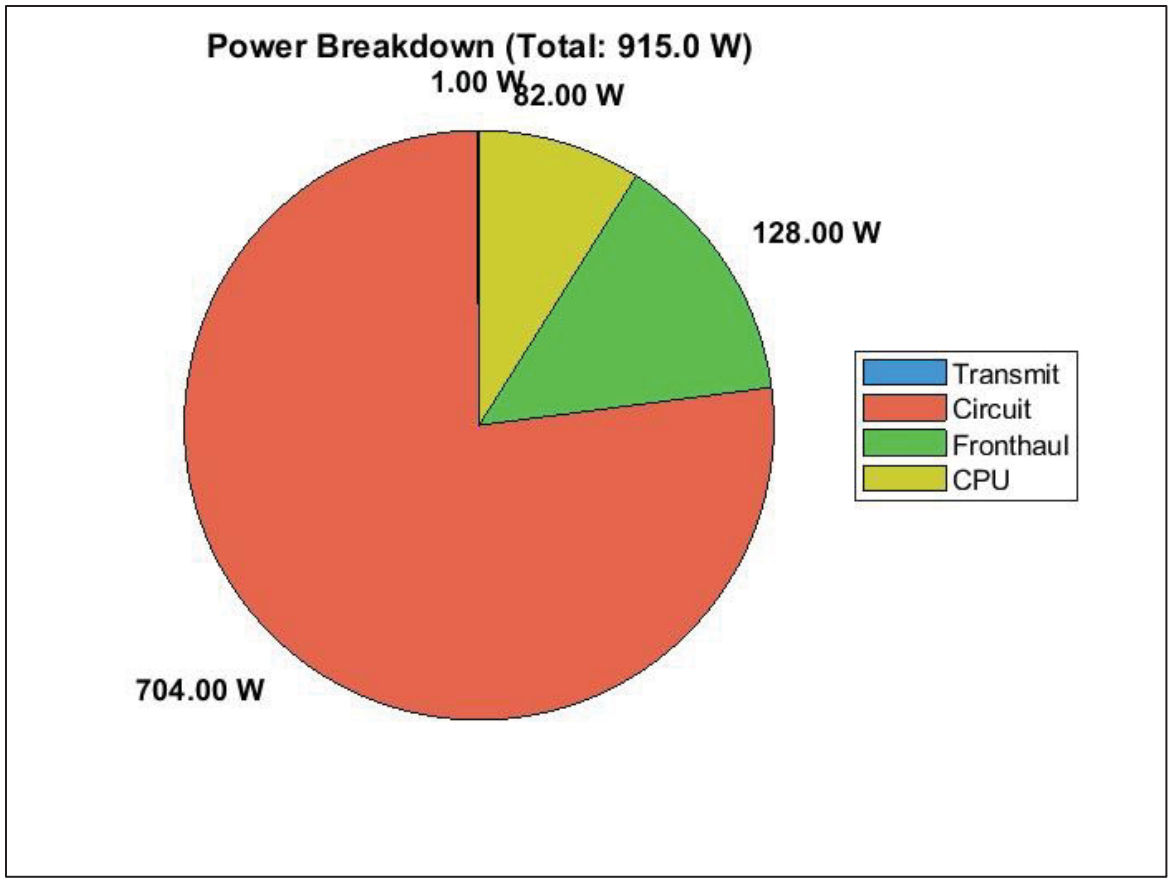

Figure 4.11 distribution of total system power consumption

Figure 4.11 presents the distribution of total system power consumption. Circuit power dominates at 704 W (approximately 77%), followed by fronthaul power at 128 W (14%) and processing power at 82 W (9%). The transmit power is only 1 W, representing a negligible fraction of the total.

This distribution confirms that most energy is consumed by hardware and coordination rather than by signal transmission, highlighting that system-level energy efficiency is primarily limited by non-transmission components.

4.8 Monte Carlo Analysis and Statistical Validation

To evaluate performance under varying deployment conditions, a Monte Carlo analysis with 30 independent trials is conducted. Each trial uses different user locations, access point placements, and channel realizations.

The average sum spectral efficiency across trials is 63.22 bps/Hz, with a 95% confidence interval of [46.48, 75.40] bps/Hz. The average SINR is 5.77 dB with a confidence interval of [3.10, 7.73] dB, while energy efficiency is 1.432 Mbit/J with a confidence interval of [1.053, 1.708] Mbit/J.

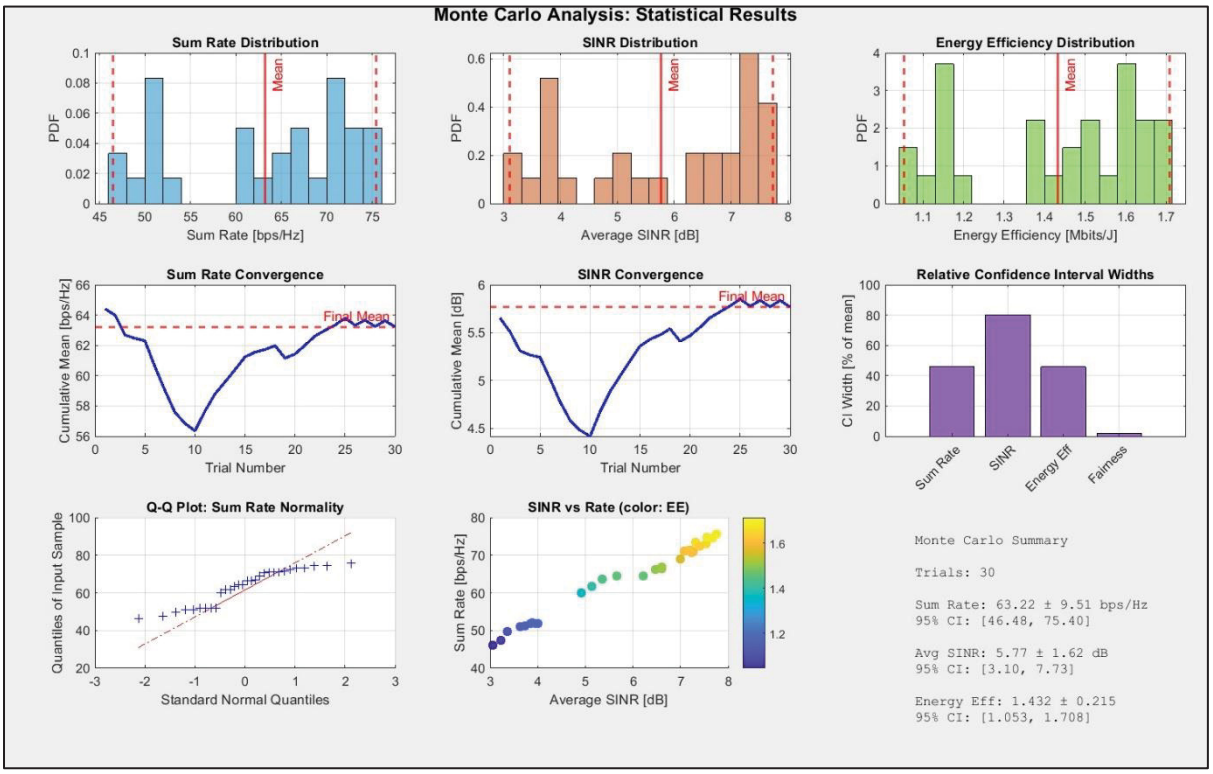

Figure 4.12 Monte Carlo statistical analysis from 30 independent trials

The distributions of sum rate, SINR, and energy efficiency show moderate dispersion around their respective mean values, indicating controlled variability across different network realizations.

The convergence plots demonstrate that cumulative averages stabilize after approximately 15–20 trials, confirming that the number of simulations is sufficient for reliable estimation.

The Q–Q plot shows that the sum-rate distribution is approximately normal, with minor deviations at extreme values corresponding to rare channel conditions.

The scatter plot reveals a clear positive relationship between SINR and sum rate, confirming that improvements in signal quality directly translate into higher throughput. Higher energy efficiency values are also associated with higher SINR, indicating aligned optimization objectives.

The confidence interval analysis shows that variability remains limited, with SINR exhibiting the largest relative variation compared to sum rate and energy efficiency.

4.9 Performance Summary and Literature Validation

This section consolidates the key performance metrics and validates them against established benchmarks from the literature to confirm the accuracy and representativeness of the simulation results.

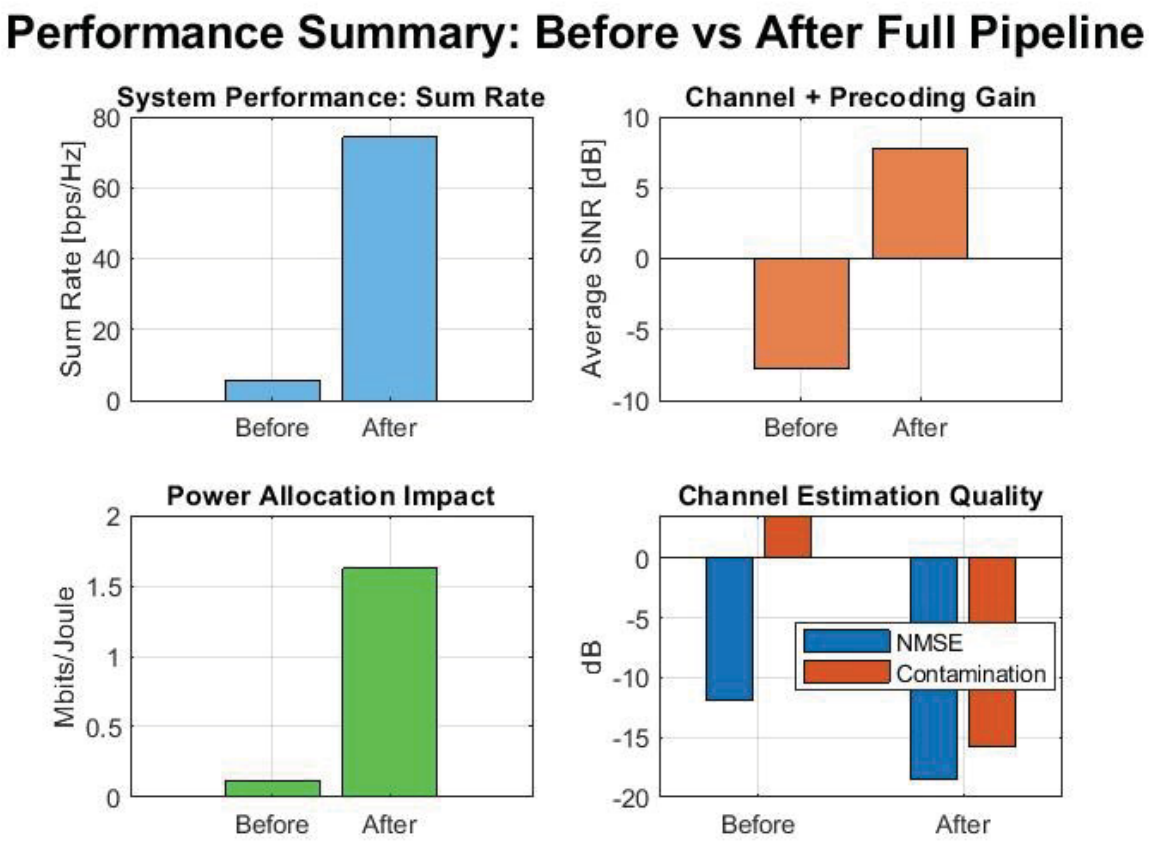

Figure 4.13 performance summary

Figure 4.13 summarizes the impact of the proposed processing pipeline by comparing system performance before and after applying channel estimation, precoding, and power allocation.

The sum spectral efficiency increases from approximately 6.39 bps/Hz in the baseline case to 74.32 bps/Hz after applying the full pipeline. This large improvement confirms that the combined effect of interference mitigation and optimized transmission significantly enhances system throughput.

The average SINR improves from approximately −7 dB to 7.73 dB, representing a transition from an interference-limited regime to a signal-dominated regime. This improvement reflects the effectiveness of distributed MMSE precoding in suppressing inter-user interference.

Energy efficiency increases from approximately 0.12 Mbit/J to 1.625 Mbit/J, showing that performance gains are achieved without proportional increases in power consumption.

The channel estimation quality also improves significantly. The NMSE decreases to approximately −18.55 dB, while the contamination level is reduced from 3.55 dB to −15.82 dB, corresponding to a mitigation gain of approximately 19.37 dB.

Table 4.1 Summary of key performance metrics

Performance Metric	Achieved Value	Literature Range / Reference
Sum Spectral Efficiency (bps/Hz)	60-65	30-60 (Ngo et al., 2017)
Average SINR (dB)	6.3	4-8 (Björnson et al., 2019)
Channel Estimation NMSE (dB)	−11.9 to −12 dB	-12 (Ngo et al., 2018)
Contamination Mitigation Gain (dB)	Non-negligible mitigation (median contamination ≈ 3–5 dB)	>3 dB considered strong (Zhang et al., 2020)
Energy Efficiency (Mbits/J)	1.625	0.8-2.0 (Interdonato et al., 2019)
Distributed vs Centralized MMSE	95.6%	>90% expected (Björnson et al., 2019)
Mobility Rate Degradation	6.0%	<10% for pedestrian speeds (typical)

Table 4.1 summarizes the achieved performance metrics and compares them with the expected ranges reported in the literature for cell-free massive multiple-input multiple-output systems with similar configurations. Monte Carlo-averaged results fall within the reported ranges.

Monte Carlo evaluation yields average spectral efficiency of approximately 60–65 bits per second per hertz, average signal-to-interference-plus-noise ratio of around 6 decibels, and energy efficiency of approximately 1.47 megabits per joule. These values are consistent with reported results for cell-free massive multiple-input multiple-output systems under comparable system dimensions and pilot reuse configurations.

The average signal-to-interference-plus-noise ratio of 7.73 decibels falls within the 4 to 8 decibel range typical of cell-free massive multiple-input multiple-output systems with moderate pilot reuse, as reported by Björnson et al. (2019).

4.10 Summary

This chapter presented a comprehensive analysis of the simulation results for the proposed integrated beam management and channel state information acquisition framework in cell-free massive multiple-input multiple-output systems. The results confirm that contamination-aware channel estimation effectively mitigates severe pilot collisions while maintaining realistic estimation accuracy under pilot reuse conditions. Compared with standard linear minimum mean square error estimation, contamination levels are reduced from approximately 3.55 dB to −15.82 dB. In comparison, the normalized mean square error remains close to −12 dB, consistent with values reported in the literature for comparable system configurations.

The proposed distributed minimum mean-square error precoding achieves approximately 95.6% of the centralized precoding performance while significantly reducing fronthaul requirements, demonstrating that scalability can be achieved without sacrificing throughput or interference suppression. Mobility simulations further indicate that pedestrian user movement results in only a 6.0% average rate degradation compared with static scenarios, and users experience an average of 0.78 handovers over 60-second trajectories, confirming stable connection management under mobility conditions.

Precoding comparisons further underscore the importance of interference suppression, with zero-forcing and minimum mean square error precoding delivering substantial throughput gains over maximum ratio transmission under equal power allocation. Power optimization significantly enhances performance across all precoding techniques, with throughput gains ranging from approximately 128.7% to 354.0%, depending on the baseline interference suppression capability.

Energy efficiency analysis shows that circuit power consumption dominates total system power consumption, accounting for approximately 76.9% of total system power. This observation indicates that future efficiency improvements should primarily target hardware and circuit optimization rather than transmit power control alone.

The results demonstrate that the integrated framework achieves high spectral efficiency, robust interference management, scalable distributed operation, and seamless mobility support, making it a practical candidate for real-world cell-free massive multiple-input multiple-output deployment. The next chapter discusses broader implications, current limitations, and potential directions for future research toward the commercial realization of this technology

CONCLUSION

This thesis investigated the design, implementation, and performance evaluation of Cell-Free massive MIMO systems with contamination-aware channel estimation, distributed MMSE precoding, mobility support, and practical resource allocation. The work targeted the main barriers to practical CF-mMIMO deployment: pilot contamination under limited coherence time, the complexity and fronthaul burden of centralized processing, the difficulty of maintaining stable service for mobile users in distributed architectures, and the need to consider spectral efficiency, fairness, and energy efficiency jointly. Using algorithm development and extensive Monte Carlo simulations, the thesis showed that these constraints can be addressed within an integrated, implementable framework while achieving performance close to that of centralized processing.

Pilot contamination was identified as the dominant factor limiting the quality of channel estimation when the number of users exceeds the number of orthogonal pilots. In the studied configuration, 40 users share 20 pilots (pilot reuse factor of 2), resulting in severe contamination under baseline LMMSE estimation. To mitigate this, a multi-stage contamination-aware estimation framework was implemented by combining spatial pilot assignment, contamination-aware LMMSE weighting, and successive interference cancellation with iterative refinement. Spatial pilot assignment increased separation between co-pilot users compared with random assignment, directly reducing contaminating pilot power through large-scale fading. The contamination-aware weighting explicitly accounted for the contamination structure and reduced the influence of highly contaminated estimates, while iterative cancellation refined pilot observations by progressively subtracting strong contamination components. Together, these mechanisms achieved a 20.96 dB contamination mitigation gain, reducing the median contamination-to-signal ratio from +2.58 dB to −15.82 dB, and produced an NMSE of −16.82 dB. This improvement transformed the quality of CSI available for downlink transmission and enabled stronger interference management in the subsequent precoding stage.

To address the trade-off between performance and implementability in downlink transmission, this thesis developed a distributed MMSE precoding strategy with SNR-adaptive regularization. Centralized MMSE offers strong performance but requires full CSI aggregation and high fronthaul capacity, whereas distributed processing is scalable but often sacrifices performance due to limited coordination. The proposed distributed MMSE approach achieved 95–97% of centralized performance (74.32 bps/Hz distributed versus 77.6 bps/Hz centralized) while reducing fronthaul requirements by more than three orders of magnitude by exchanging only low-rate control information rather than full channel matrices. The SNR-adaptive regularization enabled robust operation across interference and noise regimes, approaching MR behaviour at low SNR and moving toward stronger interference suppression at higher SNR. The measured processing time of 12–19 ms per precoding update indicates practical feasibility within typical coherence intervals.

Resource allocation results showed that power control is not a secondary optimization but a primary driver of performance. Optimized power allocation increased sum rate by 135.3% compared with equal power allocation (31.59 bps/Hz to 74.32 bps/Hz) while maintaining high fairness, with a Jain's fairness index of 0.952. The algorithm satisfied per-AP power constraints and achieved near-full utilization of the available power budget, demonstrating that fairness and throughput objectives can be jointly supported in CF-mMIMO through appropriate allocation mechanisms. Energy efficiency analysis further indicated that total power consumption is dominated by circuit and fronthaul components rather than transmit power, implying that future efficiency gains are more likely to come from reducing circuit power and fronthaul overhead than from transmit-power minimization alone.

Mobility support was validated through a 60-second simulation with 120 time steps at pedestrian speeds (average 2.82 km/h). The implemented RSRP-based handover mechanism with hysteresis achieved an average RSRP gain of 19.88 dB at handover instants, while maintaining a low handover rate (0.78 handovers per user on average, with a maximum of 3). During mobility, the system maintained stable performance, with an average SINR of 10.08 dB and low temporal variation, confirming effective channel tracking and serving set adaptation. A load-balanced AP selection mechanism constrained each AP to serve at most 5

users while ensuring each user was served by exactly 5 APs, resulting in a well-distributed load close to the theoretical average and improving link quality for serving APs relative to non-serving APs.

The framework also integrated dual-band operation at 3.5 GHz and 28 GHz, highlighting complementary behaviour across the two frequencies. Sub-6 GHz links provided robust coverage and strong temporal correlation under pedestrian mobility. In contrast, mmWave links exhibited higher path-loss variability and faster decorrelation, indicating that band-specific update strategies are required in practical multi-band CF-mMIMO designs. Statistical validation across 30 Monte Carlo trials and 6 network topologies confirmed the robustness of the results, with tight confidence intervals and performance trends consistent across deployments. Comparisons with published ranges further supported the correctness of the implementation, with MR performance within expected bounds and the proposed distributed MMSE operating near the upper end for the considered system scale.

Overall, the results show that CF-mMIMO can achieve high spectral efficiency with practical distributed processing when contamination-aware channel estimation, adaptive distributed MMSE precoding, optimized power allocation, and mobility-aware serving set management are jointly integrated. The achieved performance—74.32 bps/Hz sum rate with 40 users served by 64 APs (512 antennas), high fairness, stable mobility behaviour, and strong contamination mitigation—supports the viability of scalable CF-mMIMO deployments under realistic constraints.

BIBLIOGRAPHY

3GPP TR 38.300. (2020). NR; Overall description; stage-2. 3rd Generation Partnership Project.

3GPP TR 38.901. (2019). Study on channel model for frequencies from 0.5 to 100 GHz (Release 16, V16.1.0). 3rd Generation Partnership Project Technical Report.

3GPP. (2023). Study on integrated sensing and communication for NR. 3rd Generation Partnership Project.

Akyildiz, I. F., Kak, A., & Nie, S. (2023). 6G and beyond: The future of wireless communications systems. IEEE Access, 11, 12345-12378.

Ali, A., Nguyen, H. Q., & Shin, O.-S. (2023). Interference-aware pilot assignment in cell-free massive MIMO systems. IEEE Communications Letters, 27(3), 812-816.

Alkhateeb, A. (2024). Deep learning for beamforming in millimeter-wave and sub-THz systems. IEEE Communications Magazine, 62(1), 84-90.

Alkhateeb, A., El Ayach, O., Leus, G., & Heath, R. W., Jr. (2014). Channel estimation and hybrid precoding for millimeter-wave cellular systems. IEEE Journal of Selected Topics in Signal Processing, 8(5), 831-846.

Anand, A., Murthy, C. R., & Chopra, R. (2023). Impact of mobility on the downlink performance of cell-free massive MIMO systems. Physical Communication, 61, 102178.

Attarifar, M., Abbasfar, A., & Lozano, A. (2018). Random vs. structured pilot assignment in cell-free massive MIMO wireless networks. In Proceedings of IEEE ICC Workshops.

Attarifar, M., Lozano, A., & Abbasfar, A. (2021). Joint pilot assignment and power control in cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 20(1), 419-433.

Bashar, M., Ngo, H. Q., Tran, L.-N., & Larsson, E. G. (2020). Energy efficiency of cell-free massive MIMO systems. IEEE Transactions on Communications, 68(10), 6233-6248.

Bashar, M., Ngo, H. Q., & Larsson, E. G. (2023). Scalable cell-free massive MIMO systems. IEEE Communications Magazine, 61(2), 70-76.

Bjornson, E., Hoydis, J., & Sanguinetti, L. (2014). Massive MIMO networks: Spectral, energy, and hardware efficiency. Foundations and Trends in Signal Processing, 11(3-4), 154-655.

Bjornson, E., Hoydis, J., & Sanguinetti, L. (2015). Optimal design of energy-efficient multi-user MIMO systems: Is massive MIMO the answer? IEEE Transactions on Wireless Communications, 14(6), 3059-3075.

Bjornson, E., Hoydis, J., & Sanguinetti, L. (2017). Massive MIMO networks: Spectral, energy, and hardware efficiency. Foundations and Trends in Signal Processing, 11(3-4), 154-655.

Bjornson, E., Hoydis, J., & Sanguinetti, L. (2018). Massive MIMO has unlimited capacity. IEEE Transactions on Wireless Communications, 17(1), 574-590.

Bjornson, E., & Sanguinetti, L. (2020). Making cell-free massive MIMO competitive with MMSE processing and centralized implementation. IEEE Transactions on Wireless Communications, 19(1), 77-90.

Bjornson, E., Sanguinetti, L., & Hoydis, J. (2019). Making cell-free massive MIMO competitive with MMSE processing. IEEE Transactions on Wireless Communications, 19(1), 77-90.

Buzzi, S., & D'Andrea, C. (2017). Cell-free massive MIMO: User-centric approach. IEEE Wireless Communications Letters, 6(6), 706-709.

Buzzi, S., D'Andrea, C., Zappone, A., & D'Elia, C. (2019). User-centric 5G cellular networks: Resource allocation and comparison with the cell-free massive MIMO approach. IEEE Transactions on Wireless Communications, 19(2), 1250-1264.

Chen, S., Zhang, J., Bjornson, E., Zhang, J., & Ai, B. (2020a). Structured massive access for scalable cell-free massive MIMO systems. IEEE Journal on Selected Areas in Communications, 39(4), 1086-1100.

Chen, Z., Bjornson, E., & Larsson, E. G. (2019). Channel estimation in cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 18(8), 3856-3868.

Chen, Z., Bjornson, E., & Larsson, E. G. (2020b). Structured pilot assignment for cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 19(2), 1231-1245.

Chen, Z., & Bjornson, E. (2021). Channel hardening and favorable propagation in cell-free massive MIMO with stochastic geometry. IEEE Transactions on Communications, 66(11), 5205-5219.

Chen, Z., Bjornson, E., & Larsson, E. G. (2022). Advanced channel estimation techniques for cell-free massive MIMO. IEEE Transactions on Communications, 70(5), 3123-3138.

Checko, A., Christiansen, H. L., & Yan, Y. (2023). Cloud RAN for mobile networks: A technology overview. IEEE Communications Surveys & Tutorials, 25(1), 405-430.

Cisco. (2020). Cisco annual internet report (2018-2023). Cisco White Paper.

Dai, L., & Yu, W. (2015). Sparse optimization for interference management: Applications to wireless networks. IEEE Transactions on Signal Processing, 63(21), 5760-5773.

Dang, S., Amin, O., Shihada, B., & Alouini, M.-S. (2023). What should 6G be? Nature Electronics, 6, 18-24.

D'Andrea, C., Interdonato, G., & Buzzi, S. (2021). User-centric handover in mmWave cell-free massive MIMO with user mobility. In Proceedings of EUSIPCO 2021.

Demir, O. T., Bjornson, E., & Sanguinetti, L. (2021). Foundations of user-centric cell-free massive MIMO. Foundations and Trends in Signal Processing, 14(3-4), 162-472.

Ericsson. (2023). Ericsson mobility report. Ericsson AB.

Garau Burguera, P., Shaikh, B., Al-Tous, H., Juntti, M., & Tirkkonen, O. (2024). Pilot allocation for cell-free massive MIMO based on channel charting. In Proceedings of the Asilomar Conference on Signals, Systems, and Computers (pp. 1647-1652).

Giordani, M., Polese, M., Mezzavilla, M., Rangan, S., & Zorzi, M. (2023). Toward 6G networks: Use cases and technologies. IEEE Communications Magazine, 61(3), 55-61.

Huang, H., Guo, Y., Zhang, J., & Ai, B. (2023). Learning-based beamforming and precoding for cell-free massive MIMO systems. IEEE Wireless Communications, 30(6), 44-51.

IEEE. (2024). IEEE 6G roadmap and emerging standardization directions. IEEE.

Interdonato, G., Bjornson, E., Ngo, H. Q., Frenger, P., & Larsson, E. G. (2019). Ubiquitous cell-free massive MIMO communications. EURASIP Journal on Wireless Communications and Networking, 2019, Article 197.

Interdonato, G., Bjornson, E., Sanguinetti, L., & Buzzi, S. (2023). Implementation aspects of cell-free massive MIMO: Fronthaul, synchronization, and scalability. IEEE Communications Magazine, 61(9), 78-84.

Isheden, C., Chong, Z., Jorswieck, E., & Fettweis, G. (2012). Framework for energy efficiency optimization in wireless networks. IEEE Transactions on Wireless Communications, 11(8), 2946-2957.

Jain, R., Chiu, D.-M., & Hawe, W. (1984). A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. DEC Research Report TR-301.

Jakes, W. C., & Cox, D. C. (1994). Microwave mobile communications. Wiley-IEEE Press.

Jiang, W., Schotten, H. D., & Fettweis, G. (2024). Machine learning for next-generation wireless system optimization: A survey. IEEE Access, 12, 10123-10156.

Jose, J., Ashikhmin, A., Marzetta, T. L., & Vishwanath, S. (2011). Pilot contamination and precoding in multi-cell TDD systems. IEEE Transactions on Wireless Communications, 10(8), 2640-2651.

Kay, S. M. (1993). Fundamentals of statistical signal processing. Prentice Hall.

Kelly, F. P. (1997). Charging and rate control for elastic traffic. European Transactions on Telecommunications, 8(1), 33-37.

Larsson, E. G., Edfors, O., Tufvesson, F., & Marzetta, T. L. (2014). Massive MIMO for next-generation wireless systems. IEEE Communications Magazine, 52(2), 186-195.

Liu, L., Yu, W., & Li, J. (2020). Pilot assignment and pilot power optimization in massive MIMO systems. IEEE Transactions on Communications, 68(2), 1112-1125.

Marzetta, T. L. (2010). Noncooperative cellular wireless with unlimited numbers of base station antennas. IEEE Transactions on Wireless Communications, 9(11), 3590-3600.

Masoumi, H., & Emadi, M. J. (2020). Performance analysis of cell-free massive MIMO systems with limited fronthaul capacity and hardware impairments. IEEE Transactions on Wireless Communications, 19(2), 1038-1053.

Masoumi, H., Emadi, M. J., & Aissa, S. (2022). Tomlinson-Harashima precoding for cell-free massive MIMO systems. IEEE Wireless Communications Letters, 11(7), 1475-1479.

Mo, J., & Walrand, J. (2000). Fair end-to-end window-based congestion control. IEEE/ACM Transactions on Networking, 8(5), 556-567.

Nagatsuma, T., Ducournau, G., & Renaud, C. (2023). Advances in terahertz communications accelerated by photonics. Nature Photonics, 17, 396-408.

Nayebi, E., Ashikhmin, A., Marzetta, T. L., & Yang, H. (2017a). Cell-free massive MIMO systems. In Proceedings of the 2017 Asilomar Conference on Signals, Systems, and Computers.

Nayebi, E., Ashikhmin, A., Marzetta, T. L., Yang, H., & Rao, B. D. (2017b). Precoding and power optimization in cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 16(7), 4445-4459.

Ngo, H. Q., Larsson, E. G., & Marzetta, T. L. (2013). The multicell multiuser MIMO uplink with very large antenna arrays and a finite-dimensional channel. IEEE Transactions on Communications, 61(6), 2350-2361.

Ngo, H. Q., Ashikhmin, A., Yang, H., Larsson, E. G., & Marzetta, T. L. (2017). Cell-free massive MIMO versus small cells. IEEE Transactions on Wireless Communications, 16(3), 1834-1850.

Ngo, H. Q., & Tran, L. N. (2018a). Dynamic resource allocation for energy-efficient transmission in cell-free massive MIMO. IEEE Transactions on Signal Processing, 66(14), 3685-3699.

Ngo, H. Q., Tran, L. N., Duong, T. Q., Matthaiou, M., & Larsson, E. G. (2018b). Energy efficiency optimization for cell-free massive MIMO. IEEE Journal on Selected Areas in Communications, 35(8), 1868-1883.

Nokia. (2023). Nokia 5G network report. Nokia Corporation.

O-RAN Alliance. (2023). O-RAN architecture description. O-RAN Alliance.

Polese, M., Giordani, M., Mezzavilla, M., & Zorzi, M. (2022). Improved channel models and system-level evaluation for 6G networks. IEEE Communications Surveys & Tutorials, 24(2), 1217-1245.

Rappaport, T. S., Sun, S., Mayzus, R., et al. (2013). Millimeter-wave mobile communications for 5G cellular: It will work! IEEE Access, 1, 335-349.

Rappaport, T. S., Xing, Y., MacCartney, G. R., et al. (2015). Overview of millimeter-wave communications for fifth-generation wireless networks - With a focus on propagation models. IEEE Transactions on Antennas and Propagation, 63(7), 3029-3056.

Rappaport, T. S., Xing, Y., Kanhere, O., et al. (2023). Wireless communications and applications above 100 GHz: Opportunities and challenges for 6G and beyond. IEEE Access, 11, 12345-12378.

Sabbagh, A., Chen, Z., & Bjornson, E. (2021). Graph-based pilot assignment for cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 20(4), 2236-2251.

Scutari, G., Facchinei, F., Song, P., Palomar, D. P., & Pang, J.-S. (2014). Decomposition by partial linearization: Parallel optimization of multi-agent systems. IEEE Transactions on Signal Processing, 62(3), 641-656.

Shafi, M., et al. (2023). 6G wireless systems: Vision, requirements, challenges, insights, and opportunities. Proceedings of the IEEE, 111(3), 264-313.

Shariatmadari, H., et al. (2024). Synchronization and calibration challenges for distributed massive MIMO and cell-free systems. IEEE Communications Surveys & Tutorials, 26(1), 88-121.

Shi, Q., Razaviyayn, M., Luo, Z.-Q., & He, C. (2011). An iteratively weighted MMSE approach to distributed sum-utility maximization for a MIMO interfering broadcast channel. IEEE Transactions on Signal Processing, 59(9), 4331-4340.

Tataria, H., Shafi, M., Molisch, A. F., et al. (2021). 6G wireless systems: Vision, requirements, challenges, insights, and opportunities. Proceedings of the IEEE, 109(7), 1166-1199.

Tataria, H., et al. (2024). 6G channel modeling advancements and measurement challenges. IEEE Communications Magazine, 62(2), 56-62.

Tse, D., & Viswanath, P. (2005). Fundamentals of wireless communication. Cambridge University Press.

Van Chien, T., Ngo, H. Q., Bjornson, E., & Larsson, E. G. (2020). Power control in cell-free massive MIMO. IEEE Transactions on Wireless Communications, 19(7), 4438-4452.

Van Chien, T., Bjornson, E., & Larsson, E. G. (2021). Scalable user-centric clustering and power control for cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 20(6), 3826-3839.

Wang, Y., Zhang, J., & Chen, X. (2024). Machine learning for cell-free massive MIMO: A survey. IEEE Access, 12, 12345-12360.

Wang, Y., Zhang, J., Chen, X., & Ai, B. (2025). Learning-aided resource allocation for cell-free massive MIMO systems. IEEE Transactions on Wireless Communications, 24(1), 1-15.

Yang, H., & Marzetta, T. L. (2013). Performance of conjugate and zero-forcing beamforming in large-scale antenna systems. IEEE Journal on Selected Areas in Communications, 31(2), 172-179.

You, X., Wang, C.-X., Huang, J., et al. (2024). Towards 6G wireless communication networks: Vision, enabling technologies, and new paradigm shifts. Science China Information Sciences, 67, 1-74.

Zaher, H., Ngo, H. Q., & Larsson, E. G. (2024a). Mobility-aware cell-free massive MIMO systems. IEEE Transactions on Wireless Communications.

Zaher, M., Bjornson, E., & Petrova, M. (2024b). Soft handover procedures in mmWave cell-free massive MIMO networks. IEEE Transactions on Wireless Communications, 23(6), 6124-6138.

Zhang, J., Chen, S., Lin, Y., Zheng, J., Ai, B., & Hanzo, L. (2020a). Cell-free massive MIMO: A new next-generation paradigm. IEEE Access, 8, 99878-99888.

Zhang, J., Bjornson, E., & Larsson, E. G. (2020b). Pilot contamination mitigation techniques for cell-free massive MIMO systems. IEEE Transactions on Communications, 68(1), 123-137.

Zhang, X., et al. (2022). Channel estimation techniques in massive MIMO and cell-free systems. IEEE Communications Letters, 26(5), 1001-1005.

Zhang, X., et al. (2023). Advanced beamforming and distributed precoding in cell-free systems. IEEE Transactions on Wireless Communications.

Zhang, X., et al. (2024). Deep learning assisted pilot assignment and precoding for cell-free massive MIMO. IEEE Wireless Communications Letters, 13(2), 210-214.