Resource Management in Networks with Intelligent Metasurfaces

A Dissertation

submitted to the designated
by the Assembly
of the Department of Computer Science and Engineering
Examination Committee

by

Alexandros Papadopoulos

in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

University of Ioannina
School of Engineering
Ioannina 2026

Advisory Committee:

- **Christos Liaskos**, Assistant Professor, Department of Computer Science and Engineering, University of Ioannina
- **Evangelos Papapetrou**, Assistant Professor, Department of Computer Science and Engineering , University of Ioannina
- **Konstantinos Votis**, Principal Researcher (Grade A), Information Technologies Institute , Centre of Research and Technology, Hellas

Examining Committee:

- **Christos Liaskos**, Assistant Professor, Department of Computer Science and Engineering, University of Ioannina
- **Evangelos Papapetrou**, Assistant Professor, Department of Computer Science and Engineering, University of Ioannina
- **Konstantinos Votis**, Principal Researcher (Grade A), Information Technologies Institute , Centre of Research and Technology, Hellas
- **Sotiris Ioannidis**, Associate Professor, School of Electrical and Computer Engineering , Technical University of Crete
- **Lysimachos-Pavlos Kontis**, Professor, Department of Computer Science and Engineering, University of Ioannina
- **Vassilios V. Dimakopoulos** , Professor, Department of Computer Science and Engineering, University of Ioannina
- **Panagiotis Diamantoulakis**, Assistant Professor, Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki

DEDICATION

This dissertation is dedicated to the memory of my beloved grandpa, Tasos, who recently passed away.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to all those who contributed, directly or indirectly, to the completion of this dissertation.

I am especially grateful to my supervisor, Professor Christos Liaskos, for his guidance, invaluable scientific support, insightful comments, and continuous encouragement throughout all stages of this research. His contribution has been decisive both in shaping the present work and in my overall scientific development.

I would also like to thank the members of the examination and advisory committee for their time, comments, and interest in the subject of my work.

Finally, I am deeply grateful to my family, especially my wife, Livia, and my daughter, Ismini, as well as to all those who stood by me throughout this effort, for their patience, understanding, and unwavering moral support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Abstract

Alexandros Papadopoulos, Ph.D., Department of Computer Science and Engineering,
School of Engineering, University of Ioannina, Greece, 2026.
Resource Management in Networks with Intelligent Metasurfaces.
Advisor: Christos Liaskos, Assistant Professor.

Reconfigurable Intelligent Surfaces (RISs) are a promising physical-layer technology for Beyond-5G/6G, enabling Programmable Wireless Environments (PWEs) in which the propagation medium becomes software controlled. However, translating RIS potential into deployable network infrastructure requires solving three tightly coupled resource-management challenges: (i) how to compile RIS configuration codebooks with sufficient electromagnetic (EM) fidelity at acceptable cost, (ii) how to share a single RIS among concurrent users under operator policies with millisecond-level control latency, and (iii) how to coordinate multiple RIS units at the network level while maintaining tiered fairness under congestion. This thesis advances RIS operational maturity by developing a coherent, lifecycle-consistent framework that connects manufacturing-time computation of physics-consistent configurations to operating-time selection, sharing, and orchestration.

At the manufacturing (offline) phase, the thesis adopts a codebook paradigm: for each supported macroscopic function, configurations are computed via EM-aware optimization and stored for fast retrieval during operation. A central contribution is a sequence of physics-informed compilation tools that progressively raise fidelity and practicality, explicitly accounting for near-field operation, mutual coupling, reflections, and energy redistribution beyond the focal region. The resulting pipeline demonstrates that extracting and exploiting physical structure (e.g., correlated element responses and staged refinement) can reduce optimization cost by orders of magnitude without compromising EM consistency. The thesis further addresses scalability beyond the "optimize-per-location" paradigm by introducing a sparse 3D sampling

and compression strategy: physics-consistent configurations are generated at representative spatial samples, then compressed and organized into a compact, searchable representation that supports bounded, predictable online inference complexity while maintaining near-optimal signal-to-noise ratio (SNR).

At the operating (online) phase, the thesis formalizes an RIS as a shareable network resource hosted by macroscopic EM functions, and introduces a resource-slice notion that links multi-user efficiency to the deviation between the shared configuration and each user's single-user optimum. Building on this abstraction, it develops a practical sharing mechanism based on multiplexing precompiled codebook entries rather than repeated channel estimation and iterative online optimization. This approach enables concurrent multi-user service with low latency, while supporting operator-defined differentiation (e.g., pricing tiers) and predictable performance trade-offs. The thesis validates the deviation–performance link analytically in both far-field steering and near-field focusing, and corroborates it with emulations and real testbed measurements, showing that multiplexing can match or outperform established sharing approaches (e.g., segmentation and time division) in both efficiency and fairness, while remaining computationally lightweight.

To scale from single-surface control to multi-RIS PWEs, the thesis introduces a system-level framework that combines node-level sharing with priority-aware, congestion-pricing routing on an RIS network graph. This enables end-to-end fairness enforcement across cascaded RIS links, ensuring that tier differentiation remains stable when multiple RIS units and congested paths interact. The framework includes a fairness mechanism that enforces an operator-defined efficiency floor while maintaining proportionality to user priority weights. The resulting network behavior remains robust under increasing user densities and varying RIS counts, with practical runtime and modest memory footprint.

Finally, the thesis shows that codebook multiplexing supports PWE-native services beyond LoS restoration, focusing on proactive covert communications. It introduces an RF-fencing service that partitions the angular or spatial domain into friendly delivery zones and hostile suppression zones, and develops a controller that synthesizes the required response by combining manufacturing-time EM-consistent codebook entries with lightweight online refinement. Comprehensive evaluations across THz and mmWave far-field scenarios demonstrate strong suppression in hostile sectors while preserving near-baseline performance in intended service sectors, and near-field studies demonstrate the creation of highly localized “quiet zones” in indoor and outdoor environments. These results highlight that while densification can often approximate coverage restoration, intentional exposure control—covertness, quiet zones, and proactive suppression—benefits directly from coordinated multi-RIS programmability.

Overall, this thesis establishes a unified perspective on RIS resource management across compilation, sharing, and network orchestration. By grounding control in EM consistency while enforcing policy-driven fairness with bounded online complexity, it provides an architecture-consistent path toward practical PWEs for B5G/6G networks.

ΕΚΤΕΤΑΜΕΝΗ ΠΕΡΙΛΗΨΗ

Αλέξανδρος Παπαδόπουλος, Δ.Δ., Τμήμα Μηχανικών Η/Υ και Πληροφορικής, Πολυτεχνική Σχολή, Πανεπιστήμιο Ιωαννίνων, 2026.
Διαχείριση Πόρων σε Δίκτυα με Ευφυείς Μεταεπιφάνειες.
Επιβλέπων: Χρήστος Λιάσκος, Επίκουρος Καθηγητής.

Οι Επαναδιαμορφώσιμες Ευφυείς Επιφάνειες (Reconfigurable Intelligent Surfaces – RIS) αποτελούν μία από τις πιο υποσχόμενες τεχνολογίες για δίκτυα B5G/6G, επειδή μετατρέπουν το περιβάλλον διάδοσης από «παθητικό» μέσο σε προγραμματιζόμενη υποδομή. Στο πλαίσιο των Programmable Wireless Environments (PWEs), τοίχοι, προσόψεις και επιφάνειες επενδύονται με επαναδιαμορφώσιμες μονάδες ώστε η ηλεκτρομαγνητική (HM) διάδοση να ελέγχεται βέλτιστα με γνώμονα τις ανάγκες του δικτύου: να ενισχύεται προς τους επιθυμητούς δέκτες, να περιορίζεται σε ανεπιθύμητες κατευθύνσεις, να ανακατανέμεται χωρικά, ή να καταστέλλεται σε συγκεκριμένες ζώνες. Παρά την εκτεταμένη βιβλιογραφία, το επίπεδο ωρίμανσης τεχνικών λύσεων ώστε να καταστεί εφικτή η χρήση των RIS μέσα σε ρεαλιστικά δίκτυα είναι αναντίστοιχα χαμηλό με τα οφέλη που μπορεί να έχει. Οι βασικές προκλήσεις μπορούν να στοιχειοθετηθούν σε τρίες άξονες: το πώς παράγονται οι ρυθμίσεις των RISs ώστε να επιτυγχάνεται η βέλτιστη HM απόκριση (codebook compilation), το πώς μοιράζεται μία επιφάνεια από πολλούς χρήστες σε πραγματικό χρόνο, και το πώς συντονίζονται πολλά RIS σε ένα δίκτυο ώστε να εξασφαλίζεται δικαιοσύνη (fairness) υπό συμφόρηση και πολιτικές παρόχου. Η παρούσα διατριβή αντιμετωπίζει αυτά τα ζητήματα με ενιαία λογική κύκλου ζωής, συνδέοντας τη φάση κατασκευής/εκπαίδευσης (offline) με τη φάση λειτουργίας (online).

Στη φάση κατασκευής, η διατριβή υιοθετεί το παράδειγμα του κωδικοβιβλίου: για ένα σύνολο μακροσκοπικών HM λειτουργιών και πιθανών θέσεων του χρήστη, υπολογίζονται εκ των προτέρων ρυθμίσεις της επιφάνειας και αποθηκεύονται ώστε η λειτουργία να βασίζεται σε ταχεία ανάκτηση και ελαφριά προσαρμογή, αντί για βαριά επαναβελτιστοποίηση. Το κρίσιμο σημείο εδώ είναι ότι η βέλτιστη ρύθμιση RIS δεν καθορίζεται μόνο από ιδανικά μοντέλα γεωμετρικής οπτικής, αλλά λαμβάνονται υπόψη ΗΜ φαινόμενα που υπάρχουν σε μία πραγματική διάδοση όπως περίθλαση πεπερασμένου μεγέθους, αμοιβαία σύζευξη στοιχείων, και ανακλάσεις/πολλαπλές διαδρομές του περιβάλλοντος. Η διατριβή αναπτύσσει μία φυσικά ενημερωμένη διαδικασία κατασκευής του κωδικοβιβλίου που ενσωματώνει αυτά τα φαινόμενα και αξιοποιεί τη συμπεριφορά τους, ώστε να επιταχύνει τη σύγκλιση και να εξάγει κατάλληλα συμπεράσματα που μπορούν να χρησιμοποιηθούν. Η βελτιστοποίηση δεν εκτελείται απλά αλλά, κατά την διάρκειά της, εξάγεται πληροφορία για το πόσο επηρεάζει κάθε στοιχείο την επιθυμητή απόκριση και ποια σύνολα στοιχείων εμφανίζουν ισχυρή συσχέτιση ως προς τη μακροσκοπική συμπεριφορά. Παράλληλα, όπως είναι προφανές, δεν μπορεί να πραγματοποιηθεί βελτιστοποίηση για κάθε πιθανή θέση του χρήστη στον τρισδιάστατο χώρο. Η διατριβή δείχνει ότι η λύση είναι να εκτελεστεί φυσικά συνεπής βελτιστοποίηση σε αραιό, αντιπροσωπευτικό δειγματοληπτικό πλέγμα του χώρου, και στη συνέχεια να συμπιεστούν οι ρυθμίσεις αξιοποιώντας την ενδογενή πλεονασματικότητα που προκύπτει όταν διαφορετικές θέσεις οδηγούν σε παρόμοιες ΗΜ αποκρίσεις. Η online επιλογή περιορίζεται σε μικρό, σχετικό υποσύνολο υποψήφιων λύσεων και συνεπώς αποκτά οριοθετημένη, προβλέψιμη πολυπλοκότητα, ενώ η απόδοση παραμένει κοντά στη βέλτιστη.

Στη φάση λειτουργίας, το βασικό βήμα της διατριβής είναι η τυπική θεμελίωση του RIS ως δικτυακού πόρου. Μία μονάδα RIS μπορεί να εξυπηρετεί πολλαπλές μακροσκοπικές ΗΜ λειτουργίες και, όταν αυτές συνυπάρχουν, η αποδοτικότητα της καθεμιάς μειώνεται. Η διατριβή εισάγει την έννοια του «τμήματος πόρου» (resource slice) ως λόγο της αποδοτικότητας υπό συνύπαρξη προς την αποδοτικότητα όταν η λειτουργία εξυπηρετείται μόνη της. Το πρακτικό και ταυτόχρονα βαθιά φυσικό συμπέρασμα είναι ότι η υποβάθμιση στην απόδοση συσχετίζεται ισχυρά με την απόκλιση της κοινής ρύθμισης από την ατομικά βέλτιστη ρύθμιση του χρήστη. Η διατριβή στηρίζει αυτόν τον δεσμό αναλυτικά τόσο σε far-field steering όσο και σε near-field focusing, όπου η φασική απόκλιση μεταφράζεται σε μη-συνεκτικό άθροισμα συνεισφορών και άρα σε μειωμένο κέρδος/πεδίο στον στόχο. Στη συνέχεια, επιβεβαιώνεται εμπειρικά ότι η απόκλιση αυτή μπορεί να λειτουργήσει για τον υπολογισμό μιας γρήγορης εκτίμησης της αποδοτικότητας των χρηστών όταν εξυπηρετούνται από το ίδιο RIS. Αυτό καθιστά εφικτή μια πρακτική στρατηγική πολυχρηστικής εξυπηρέτησης: αντί για τμηματοποίηση του RIS (segmentation) ή κυκλική εξυπηρέτηση μέσω χρόνου (TDMA), η κοινή ρύθμιση προκύπτει από σύνθεση/πολυπλεξία προϋπολογισμένων καταχωρήσεων του κωδικοβιβλίου. Η διατριβή δείχνει ότι αυτή η προσέγγιση προσφέρει σταθερή, χαμηλής καθυστέρησης λειτουργία και συγκρίσιμη ή ανώτερη μέση απόδοση, ενώ μειώνει φαινόμενα ακραίας αδικίας που μπορούν να εμφανιστούν όταν κάποιοι χρήστες «χάνουν» μεγάλο μέρος του RIS λόγω κατανομής σε υπο-πάνελ ή χρονικά παράθυρα. Πέρα από την ισότιμη εξυπηρέτηση, η εργασία ενσωματώνει πολιτικές παρόχου και επίπεδα υπηρεσίας: η δικαιοσύνη δεν αντιμετωπίζεται ως απόλυτη ισότητα, αλλά ως ελεγχόμενη διαφοροποίηση, όπου οι χρήστες υψηλότερης βαθμίδας πρέπει να διατηρούν αναλογικά καλύτερη απόδοση.

Η μετάβαση από ένα μεμονωμένο RIS σε δίκτυα PWEs με πολλά RIS απαιτεί επιπλέον έλεγχο σε επίπεδο δρομολόγησης και συμφόρησης. Η διατριβή μοντελοποιεί το RIS-enabled δίκτυο ως γράφο, όπου κόμβοι είναι οι RIS μονάδες και ακμές συνδέουν ζεύγη με κατάλληλες συνθήκες (π.χ. LoS εντός ορίου). Σε αυτό το πλαίσιο, η επιλογή διαδρομής δεν μπορεί να είναι μόνο γεωμετρική (ελάχιστη απόσταση), διότι οι αποφάσεις πολλών χρηστών δημιουργούν συμφόρηση σε συγκεκριμένες ακμές/κόμβους, και η τοπική κοινή ρύθμιση σε ένα RIS επηρεάζεται από το ποιοι χρήστες είναι συνδεδεμένοι σε αυτό. Η διατριβή εισάγει έναν μηχανισμό επιλογής διαδρομής που συνδυάζει το φυσικό κόστος (απόσταση/ποινή διαδοχικών ανακλάσεων) με όρους συμφόρησης και όρους προτεραιότητας που υλοποιούν την πολιτική παρόχου, οδηγώντας σε κατανομές όπου οι χρήστες υψηλότερης βαθμίδας τείνουν να λαμβάνουν μικρότερες ή λιγότερο συμφορημένες διαδρομές, ενώ οι χαμηλότερες βαθμίδες εκτρέπονται προς φθηνότερους/λιγότερο φορτισμένους πόρους. Παράλληλα, εισάγεται ο κατάλληλος μηχανισμός δικαιοσύνης που επιβάλλει ελάχιστο επίπεδο αποδοτικότητας για όλους, ώστε να αποφεύγονται περιπτώσεις πολύ χαμηλής απόδοσης για χρήστες χαμηλής προτεραιότητας υπό συνθήκες υψηλής πυκνότητας του δικτύου. Η αξιολόγηση σε σενάρια με δυαδικές ρυθμίσεις RIS (περιορισμός ιδιαίτερα απαιτητικός για πολυπλεξία) δείχνει ότι η προσέγγιση διατηρεί υψηλά επίπεδα συνολικής δικαιοσύνης ακόμη και όταν αυξάνεται ο λόγος χρηστών προς RIS μονάδες, ενώ οι χρόνοι εκτέλεσης και η κατανάλωση μνήμης παραμένουν πρακτικοί για λειτουργία σε σχεδόν πραγματικό χρόνο.

Πέρα από το κλασικό σενάριο «αποκατάστασης οπτικής επαφής», η διατριβή αναδεικνύει ότι η πραγματική προστιθέμενη αξία των PWEs εμφανίζεται σε υπηρεσίες όπου μπορούν να προσθώσουν στο δίκτυο δυνατότητες ελέγχου και προστασίας της πληροφορίας, κάτι που δεν επιτυγχάνεται με απλή πυκνοποίηση σταθμών βάσης ή πρόσθετα σημεία πρόσβασης. Σε αυτό το πνεύμα, εισάγεται η έννοια του RF-fencing ως υπηρεσία προληπτικής μυστικότητας/ασφάλειας: το πεδίο διάδοσης διαχωρίζεται σε φιλικές περιοχές μετάδοσης του σήματος και εχθρικές περιοχές καταστολής, με στόχο ταυτόχρονα να διατηρείται ποιοτική σύνδεση στους νόμιμους δέκτες και να ελαχιστοποιείται η ακτινοβολία προς δυνητικούς υποκλοπείς/παρεμβολείς. Η υλοποίηση βασίζεται σε σύνθεση προϋπολογισμένων, ΗΜ-συνεπών καταχωρήσεων κωδικοβιβλίου που αφορούν σενάρια beam steering, και σε ελαφριά online βελτιστοποίηση που διορθώνει μικρές αποκλίσεις χωρίς να απαιτείται βαριά επαναβελτιστοποίηση. Η διατριβή τεκμηριώνει μέσω εκτεταμένων εξομοιώσεων σε mmWave και THz ότι επιτυγχάνεται ισχυρή καταστολή σε ανεπιθύμητες κατευθύνσεις με μικρές απώλειες στις επιθυμητές περιοχές εξυπηρέτησης, ακόμη και όταν οι περιοχές-στόχοι βρίσκονται κοντά γωνιακά ή όταν οι απαιτήσεις καταστολής είναι πολλαπλές. Επιπλέον, σε near-field περιβάλλοντα η υπηρεσία επεκτείνεται στη δημιουργία αυστηρά τοπικών «ήσυχων ζωνών» (quiet zones) σε εσωτερικούς και εξωτερικούς χώρους, όπου η ΗΜ ενέργεια καταστέλλεται εντός μιας καθορισμένης χωρικής περιοχής ενώ το υπόλοιπο πεδίο παραμένει ουσιαστικά ανεπηρέαστο. Αυτό αποδεικνύει ότι η προγραμματιζόμενη διάδοση μπορεί να προσφέρει χωρική «απομόνωση» με έλεγχο, ανοίγοντας δρόμους για υπηρεσίες ασφάλειας και ιδιωτικότητας που είναι εγγενείς στα PWEs.

CHAPTER 1

INTRODUCTION

1.1 Programmable Wireless Environments in the 6G Era

6G systems are expected to support a diverse set of demanding services, including immersive extended reality [1], autonomous mobility [2], integrated sensing [3], wireless power transfer [4], and mission-critical control. Such services demand not only high data rates, but also stringent guarantees on latency, reliability, adaptability, spatial selectivity, and energy efficiency [5]. Yet, despite major advances in endpoint intelligence, wireless performance remains fundamentally constrained by the propagation

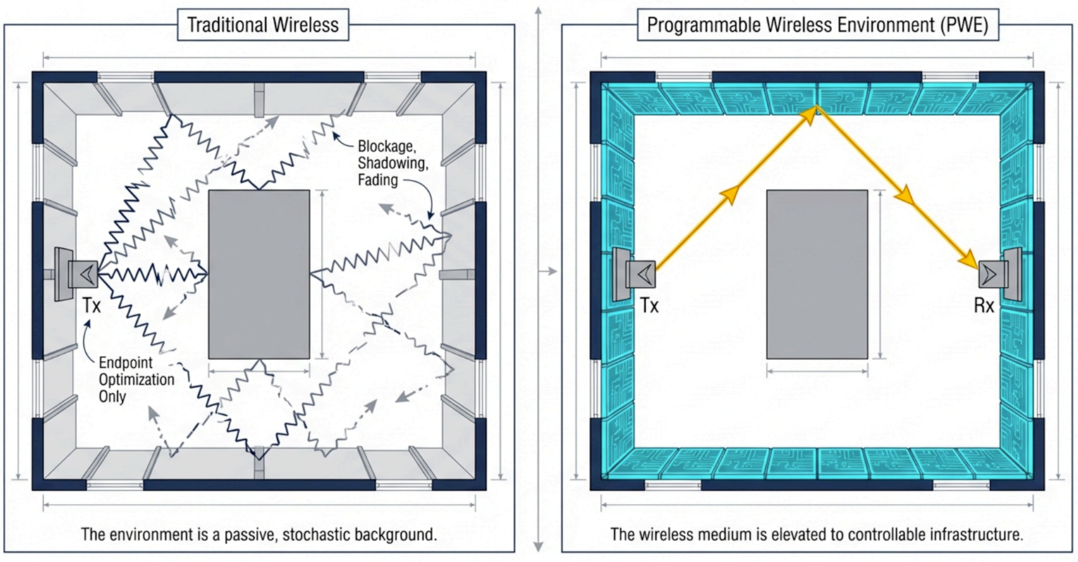

Figure 1.1: Comparison between traditional wireless networks and PWEs.

medium itself. Path-loss, blockage, shadowing, multipath fading, uncontrolled interference, and Doppler effects all arise from the fact that the environment is typically treated as passive and uncontrollable [6].

Historically, wireless-system design has focused on strengthening the transmitter and receiver. Massive Multiple-Input-Multiple-Output (MIMO) [7], beamforming [8], adaptive modulation [9] and coding [10], retransmission schemes [11], and interference-management techniques [12] have all improved performance by compensating for channel impairments at the endpoints. This strategy has been highly successful, but it leaves the propagation medium itself untouched. As a result, the wireless channel is still largely treated as a random object that must be estimated, tracked, and tolerated, rather than deliberately shaped [13].

This limitation becomes even more severe at high frequencies. In the mmWave and THz bands, propagation is more fragile because of stronger free-space attenuation [14], increased blockage sensitivity, molecular absorption [15], and higher vulnerability to mobility-induced effects [16]. Non-line-of-sight communication becomes difficult to sustain, and even relatively small obstacles may severely degrade connectivity [17]. Under such conditions, relying exclusively on endpoint adaptation becomes increasingly restrictive. A complementary paradigm is therefore needed, one that transfers part of the control effort from the endpoints to the environment itself.

A promising answer is the *programmable wireless environment* (PWE) [18]. In a PWE, major environmental objects and surfaces, such as walls, ceilings, facades, partitions, and furniture, are augmented with programmable electromagnetic (EM) structures that actively participate in the communication process. Instead of interacting with waves in an uncontrolled manner, these surfaces become software-actuated propagation resources. The environment is thus elevated from passive background to controllable infrastructure.

As illustrated in Fig. 1.1, propagation in this paradigm is no longer viewed solely as a stochastic channel realization. It becomes a partially software-defined process [19]. EM energy can be redirected along custom air-paths, focused toward intended users, split across directions, attenuated in selected regions, or absorbed where suppression is desired. In this way, a PWE can restore blocked links, improve coverage, reduce interference, enhance spatial selectivity, support wireless power delivery, and enable new security- and sensing-oriented services.

The importance of this paradigm is twofold. First, it expands the design space of wireless systems by introducing the environment itself as an optimization variable. Second, it aligns naturally with the broader 6G vision, in which communication, sensing, computing, and control are increasingly integrated. A programmable environment can therefore support not only link enhancement, but also localization [20], sensing assistance [3], EM exposure shaping [21], and proactive physical-layer protection [22].

1.2 From Metasurfaces to Programmable Propagation

The physical enablers of PWEs are metamaterials and, in particular, metasurfaces. Metamaterials are artificial EM structures whose macroscopic behavior is determined primarily by engineered sub-wavelength geometry rather than by chemical composition [18]. Their planar realization, the metasurface, consists of arrays of sub-wavelength unit cells, often referred to as meta-atoms or simply cells, arranged over a thin substrate [23].

Because of their engineered structure, metasurfaces can realize EM responses that are difficult or impossible to obtain with natural materials [24]. Depending on the design, they can support anomalous reflection [25], wavefront shaping [26], focusing [27], beam steering [28], beam splitting [29], polarization manipulation [30], absorption [31], and dispersion engineering [32]. These functions arise from the collective response of the unit cells and from the way the surface manipulates the induced currents generated by the impinging field [33].

Early metasurfaces were largely static: they were designed for one task under assumed operating conditions and could not adapt after fabrication [34]. The next step was the development of tunable and reconfigurable metasurfaces, whose EM response can be modified through external control [35]. Depending on the implementation, such control may be electrical [36], optical [37], thermal [38], magnetic [39], or mechanical [40].

A particularly important direction within this evolution is that of Software Defined Metamaterials (SDMs) [19]. The core idea is to couple the EM layer with embedded control networks and tunable elements, so that the physical behavior of the material can be altered through software-level actuation. In contrast to conventional rigid structures, SDMs support dynamic reconfiguration, interactivity, and potentially sensing-assisted adaptation [41]. This shifts the role of the material itself: it is no longer merely fabricated to exhibit a desired response, but can be programmed to realize different responses over time.

Several related terms coexist in the literature and are often used interchangeably, although they are not fully equivalent. *SDM* is a broad systems term for controllable metamaterial structures. *Metasurface* denotes the general planar EM realization. *Reconfigurable intelligent surface* (RIS) commonly refers to a more restrictive, and often reflectarray-like, implementation. *Tile* denotes an individually addressable planar unit used to cover larger surfaces within an environment. A *PWE* is then a three-dimensional space whose inner surfaces are partially or fully coated with such programmable tiles, so that EM functions can be assigned to them in software.

This distinction matters. A general metasurface is best understood as a programmable manipulator of induced surface currents, and therefore as a potentially broad field-transformation platform. By contrast, many RIS models used in communication theory approximate the surface as an array of largely isolated phase-shifting elements. This abstraction is analytically convenient and has enabled substantial progress, but it corresponds to a narrower class of hardware and a narrower set of controllable interactions. RIS should therefore be viewed as an important subset of the broader metasurface landscape rather than as an exhaustive physical description of programmable surfaces. This becomes especially relevant when the objective is not only theoretical performance analysis, but also realistic control, implementation, and network integration.

1.3 Hardware Realizability and Control Granularity

At the hardware level, a programmable metasurface typically consists of patterned conductive structures on a dielectric substrate, together with components that enable controllable local impedance [35]. The unit cells may include passive metallic inclusions and active elements such as PIN diodes [42], varactors [43], MEMS switches [44], graphene-based components [45], liquid crystals [46], or other tunable devices. By modifying the electrical state of these components, the local reflection or transmission response of each cell can be adjusted [35].

A central practical issue is the granularity of control. Globally tunable metasurfaces, where a single control signal affects the entire surface, are simpler to realize but inherently limited in functionality [47]. Richer operation requires locally tunable metasurfaces, in which each cell, or a small group of cells, can be controlled independently [48]. This enables spatially varying impedance profiles and thereby supports steering, focusing, splitting, absorption, and polarization manipulation.

The control of active metasurface elements can be continuous or discrete, with binary control being the most common in current prototypes [49]. In practice, however, true multifunctionality requires more than coarse binary switching. Although binary control is sufficient for certain tasks, it strongly restricts the set of achievable impedance states and therefore the range of attainable EM behaviors [49]. More advanced designs thus pursue continuous [50] or quasi-continuous control [51] of both reactive and resistive components. This enables simultaneous phase and amplitude shaping and significantly broadens the functional space of the surface.

Programmable metasurfaces are therefore not merely EM sheets, but cyber-physical devices. As illustrated in Fig. 1.2, behind the EM layer one typically finds bias lines, routing structures, regulators, Radio-Frequency (RF) chokes, assistive circuitry, and a gateway or controller interface [6]. The control substrate must be designed carefully so that the DC and digital networks do not compromise the RF response. EM design and electronic design are thus tightly coupled.

This coupling is evident in recent practical prototypes. Multifunctional microwave implementations have shown that a common hardware platform can support different

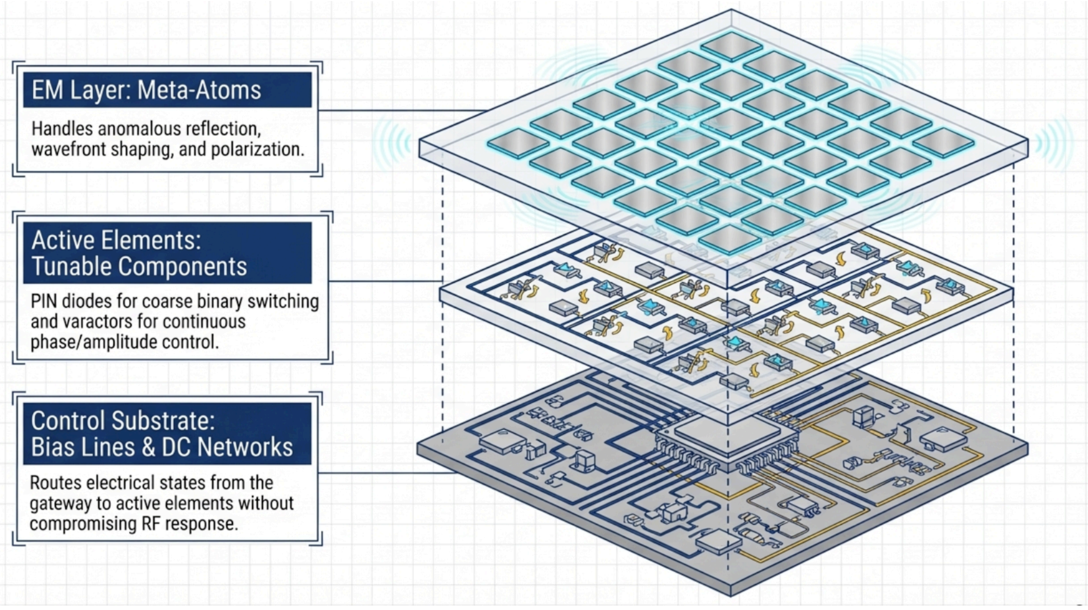

Figure 1.2: Layout of programmable metasurfaces.

EM behaviors through programmable complex loads applied at the meta-atoms [35]. Such results are important because they demonstrate that programmable-surface control is not merely a software abstraction; it is a realizable hardware architecture subject to fabrication constraints, calibration requirements, bandwidth limits, and performance trade-offs [52].

Accordingly, any serious control methodology must remain grounded in hardware realism. If a desired functionality cannot be realized under actual biasing limits, coupling effects, quantization, fabrication tolerances, bandwidth restrictions, or control overhead, then an idealized optimization result has limited practical value. Reducing the gap between elegant abstractions and realistic programmable-surface operation is therefore a recurring motivation throughout this thesis.

1.4 From Microscopic Configurations to Functional Control

A programmable surface operates by controlling how induced surface currents evolve once an incident wave reaches the structure. At the microscopic level, the surface is described by the states of its active components: switch states, bias voltages, tunable impedances, or related control variables assigned across the unit cells. In this thesis, this microscopic assignment is referred to as the *configuration* of the surface [53].

At the macroscopic level, however, what matters to the communication system is not the raw vector of low-level states, but the EM behavior that the surface realizes [54]. Such behavior may correspond to steering an incoming wave toward a desired direction, focusing energy into a target region, partially absorbing radiation, converting polarization, or splitting energy across multiple outgoing paths. In this thesis, this higher-level EM behavior is referred to as the *functionality* of the surface [52].

The distinction between configuration and functionality is fundamental. Much of the RIS literature formulates control directly at the level of individual elements, optimizing phase or amplitude parameters online. While this is useful for benchmarking algorithms and studying theoretical limits, it scales poorly as the surface size grows [55]. A metasurface with hundreds or thousands of controllable cells gives rise to a high-dimensional optimization problem that is difficult to solve online, especially under mobility, imperfect state information, or multi-user operation [56].

A more scalable alternative is to control the surface primarily at the functionality level. In this view, the controller issues a high-level EM request, such as steering, focusing, absorption, or a supported combination thereof, while the mapping to concrete hardware states is handled by a lower layer. This is the key idea behind *codebook-based control* [55].

A codebook is a structured repository that maps desired surface functionalities, together with their relevant parameters, to realizable configurations [57]. A codebook entry stores the hardware-level instructions required to realize a given EM behavior under specified conditions. These conditions may include frequency, angle of incidence, angle of departure, polarization, target region, power split, or other descriptors relevant to the supported function.

This naturally leads to a two-phase workflow, illustrated in Fig. 1.3. In the *offline phase*, the designer or manufacturer characterizes the surface, calibrates the relation between electrical actuation and EM response, and solves the underlying configuration-synthesis problem for a range of supported functionalities [55]. The resulting realizable states are then organized into a codebook. In the *online phase*, the runtime system receives service requests, selects or composes suitable functionalities, queries the codebook, and deploys the corresponding low-level states to the relevant surface units [58].

This offline–online separation is essential for scalability. It shifts the main compu-

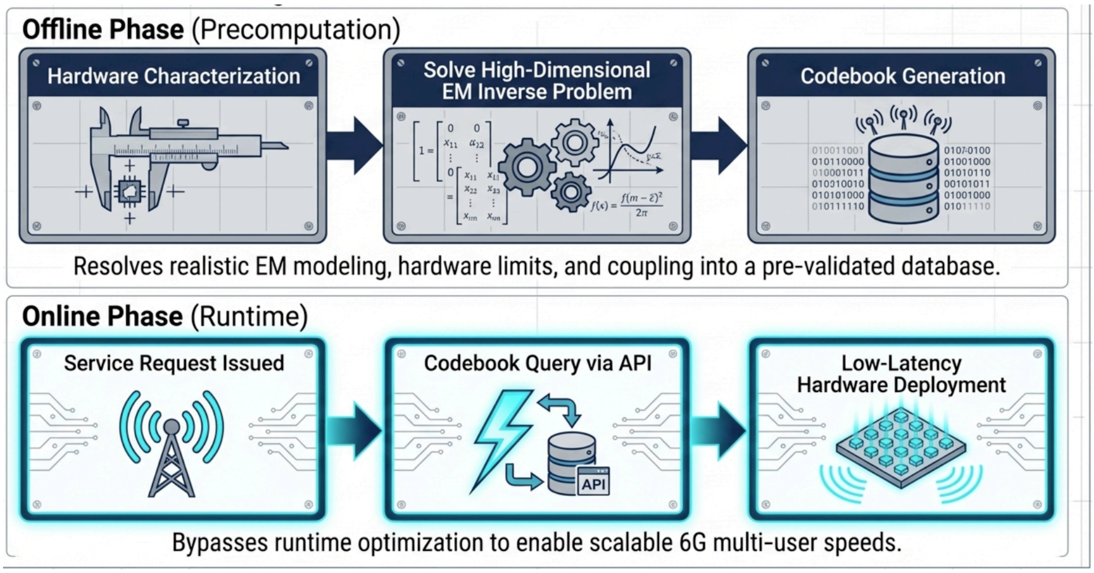

Figure 1.3: The two-phase control workflow of a PWE.

tational burden away from time-critical operation and toward precomputation, calibration, and compilation. The runtime controller no longer needs to solve a full inverse EM problem from scratch for every request; instead, it can retrieve and adapt prevalidated solutions with much lower latency. This becomes especially important in dynamic, multi-user, and service-aware environments, where repeated high-dimensional online optimization would be prohibitive [13].

At the same time, codebook-based control raises its own design questions. How should functionalities be parameterized? How rich should the codebook be? How should multi-function operation be represented? How should model mismatch and hardware inaccuracies be incorporated during compilation? These questions are central rather than secondary, and a major objective of this thesis is to develop methodologies that make codebook-oriented control both physically meaningful and operationally efficient.

1.5 Programmable Surfaces as Network Resources

A decisive conceptual step in the PWE literature is the interpretation of programmable metasurfaces as network components rather than isolated EM devices. Under this perspective, a metasurface or tile can be viewed as a controllable *wave router*. It receives incident energy under certain conditions and applies an EM function that determines

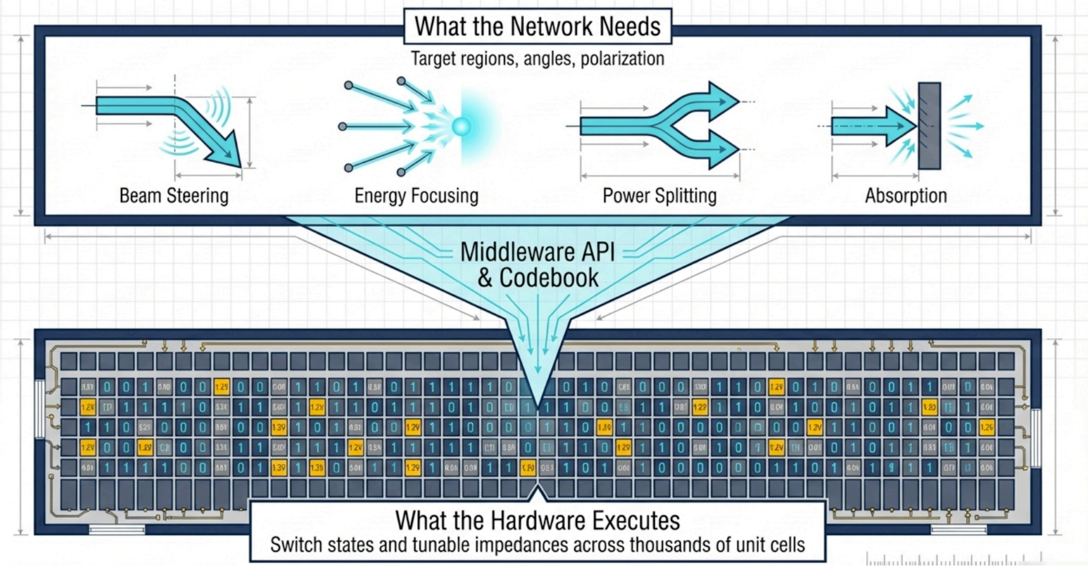

Figure 1.4: Software pipeline for SDM control within the network.

how this energy is redirected, shaped, split, attenuated, or absorbed. This abstraction makes programmable propagation amenable to graph-based and network-oriented reasoning [18].

The software pipeline that supports this abstraction is especially important for making programmable surfaces usable by non-specialists. In the systems view adopted here, the surface exposes an *application programming interface (API)* through which higher-level commands can be issued, such as steering an incoming wave from one direction to another, focusing energy toward a region, absorbing radiation, or splitting power among multiple outgoing directions. The API hides the underlying EM and hardware states. Between the API and the physical surface lies a middleware layer that translates these high-level requests into concrete hardware instructions by consulting the codebook database.

Once multiple programmable surfaces are deployed in a space, the environment can be represented as a graph [18]. Nodes may correspond to surfaces, tiles, access points, gateways, user devices, or other relevant entities. Edges represent feasible EM interactions, such as visibility relations or controllable propagation between nodes. End-to-end communication through the programmable environment can then be modeled as a path-selection problem over this graph, where each path corresponds to an *air-path* realized through a sequence of coordinated EM actions.

This graph view is powerful for several reasons. First, it decomposes the overall control problem into interacting units with explicit relations, rather than treating the

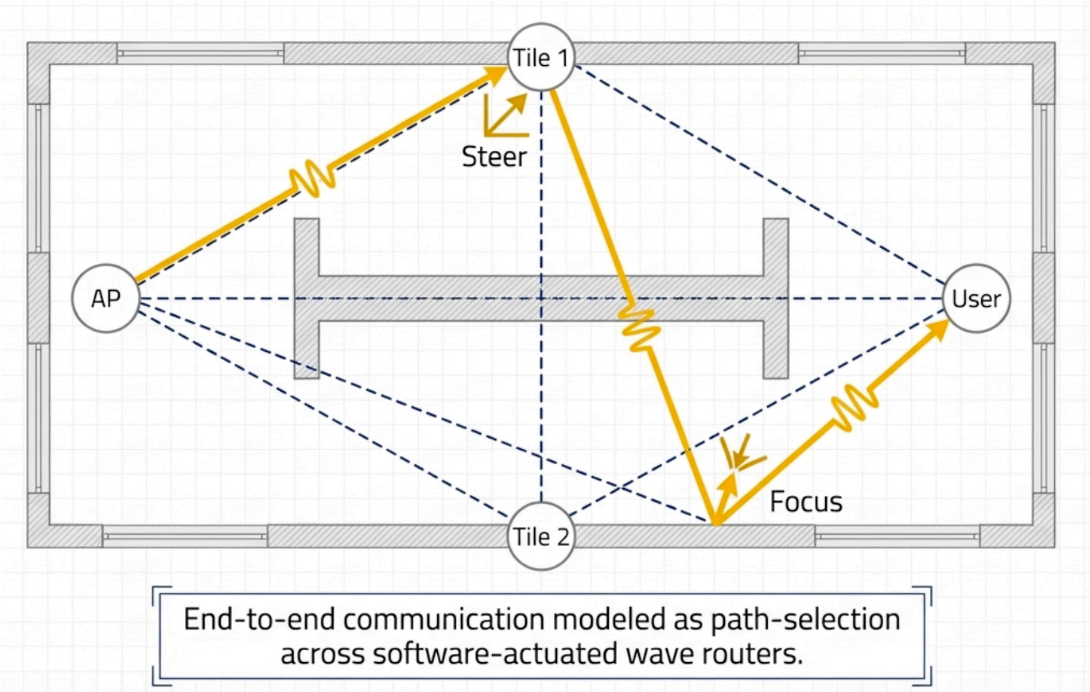

Figure 1.5: Illustration of the graph view in PWEs.

environment as an opaque global optimization space. Second, it aligns naturally with familiar networking notions such as routing, scheduling, resource allocation, policy enforcement, and traffic engineering. Third, it allows communication objectives to be expressed at a level that is closer to end-to-end service provisioning.

This viewpoint also clarifies the difference between the broader PWE perspective and the narrower signal-processing view that often dominates RIS studies. In the broader PWE framework, the architecture spans multiple layers: the physical metasurface layer, the embedded control substrate, the middleware and software-abstraction layer, the orchestration layer, and the application-facing interface. By contrast, much of the RIS literature focuses on the optimization of one or a few surfaces with respect to a communication metric, often assuming that the surrounding control infrastructure already exists. Such works address an important subset of the problem, but not the full system.

This thesis adopts the broader network-centric view. The aim is not only to show that one programmable surface can improve one link, but to understand how realistic surfaces can be characterized, compiled, controlled, shared, and orchestrated as reusable network resources within larger programmable environments.

1.6 Operation, Sharing, and Virtualization in PWEs

A PWE can be described as a layered system, illustrated in Fig. 1.6. At the *physical layer*, programmable surfaces are deployed over relevant environmental objects and interfaced with gateways or embedded controllers that expose connectivity and actuation primitives. At the *control layer*, these surfaces are managed through software abstractions that hide low-level EM details and allow functionalities to be requested in a structured manner. At the *network layer*, the environment is represented and operated as a connected system of controllable propagation resources. Finally, at the *application layer*, the environment exposes services such as connectivity enhancement, privacy-aware routing, wireless power transfer, sensing support, or spatial EM shaping.

A typical operation flow proceeds as follows. First, programmable surfaces are deployed and registered with the infrastructure. Their capabilities are profiled through characterization and offline compilation. Users then enter the environment and explicitly or implicitly express communication and service objectives. A controller gathers relevant context, such as user positions, traffic demands, priorities, and policy constraints. Using the environment graph and the codebook of supported functionalities, it selects suitable air-paths and assigns functionalities to the relevant surfaces. These requests are issued through the API, translated by the middleware into concrete hardware instructions, and deployed to the corresponding surface units. When needed, measurements and feedback can further support local refinement around the initial deployed solution.

This workflow highlights the importance of software interfaces and middleware. The controller should not manipulate raw cell states unless strictly necessary. Instead, it should operate through structured abstractions that expose functionality-level requests while hiding microscopic implementation details. Without such layers, programmable-surface control would remain a specialized hardware task rather than a usable communication-system primitive.

Once programmable surfaces are viewed as infrastructure resources, resource management becomes unavoidable. Surface area, controllable regions, supported functionalities, and control bandwidth are all finite. A customized propagation path consumes surface resources, and multiple users may compete for the same programmable region or control budget. Consequently, programmable surfaces must be shared, scheduled,

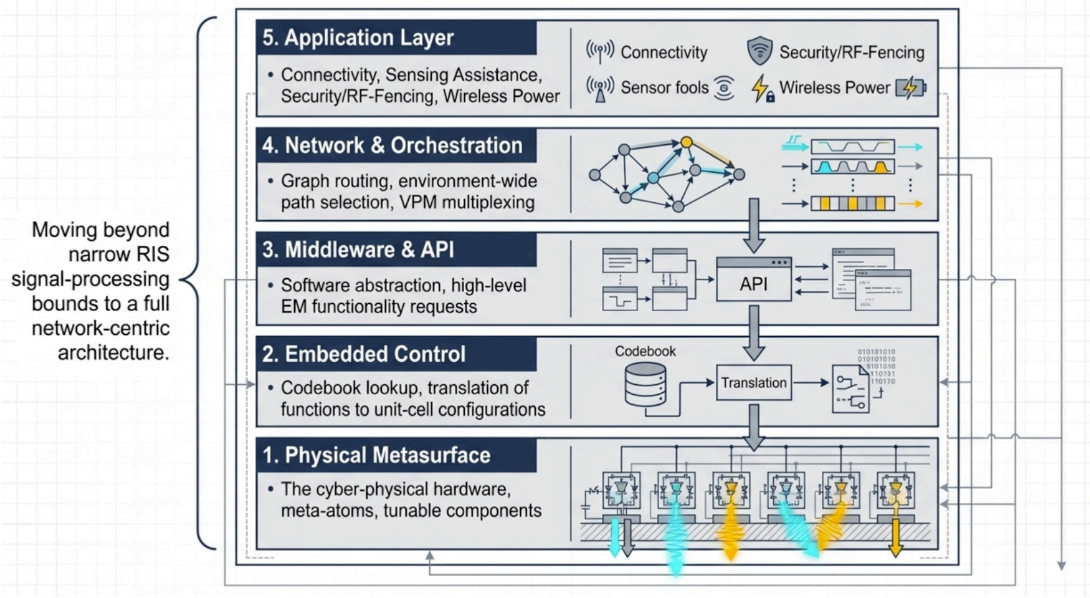

Figure 1.6: Network-centric layered architecture of a PWE.

and possibly sliced among different users, tenants, or services.

This issue becomes especially evident in realistic multi-user operation. A surface is not dedicated to a single link; it must serve heterogeneous demands under potentially conflicting objectives. Some users may require coverage restoration, others secrecy support, others energy delivery, and others low-latency service continuity. Once such heterogeneity is acknowledged, the problem is no longer purely EM. It becomes a resource-allocation problem over programmable propagation capabilities.

This observation naturally leads to virtualization. A physical programmable metasurface can be abstracted into one or more logical resources, each representing a subset of the underlying control and EM capabilities. In this context, a virtual programmable metasurface (VPM) can be understood as a software abstraction that exposes part of the physical surface to a given controller, tenant, or service class. A hypervisor-like entity can mediate between physical resources and virtual instances, enabling controlled sharing, isolation, lifecycle management, and policy enforcement.

This virtualization view is attractive because it connects programmable-surface control with broader telecom softwarization trends. It aligns naturally with concepts from Software-Defined Networking (SDN), Network Functions Virtualization (NFV), and Open Radio Access Network (O-RAN), thereby making metasurface-enabled environments easier to integrate into modern network-management ecosystems. It also provides a principled way to reason about fairness, prioritization, service differentiation, and resource guarantees.

Importantly, metasurface resources may be sliced across several domains. Spatial slicing partitions the surface into regions. Temporal slicing enables time sharing. Frequency slicing allocates different spectral resources. More advanced abstractions may also consider functionality-oriented slicing, where different portions of the EM capability space are assigned to different services. This further reinforces the view that programmable surfaces are not simple on/off reflectors, but rich sharable resources whose abstraction and management require dedicated methodologies.

For this thesis, this perspective is central. A surface that can realize sophisticated EM functions but cannot be efficiently shared, reused, and composed in a live network has limited practical value. Therefore, multiplexing, virtualization, and low-complexity sharing mechanisms form a core part of the programmable-surface problem studied here.

1.7 Beyond Link Improvement: Service-Aware and Security-Aware Propagation

A major strength of PWEs is that they enable services that go beyond received-power maximization. Because the environment can shape where energy goes, and where it does not go, it becomes possible to define spatial propagation policies. Such policies may improve communication quality, but they can also enforce privacy, isolate flows, support wireless power transfer, or suppress activity in selected regions.

This capability connects PWEs naturally to physical-layer security. By steering energy away from unintended listeners, limiting leakage into hostile regions, or creating preferred delivery zones and suppression zones, programmable environments can support proactive covertness and confidentiality mechanisms. Unlike classical security mechanisms that operate on top of an uncontrolled channel, these approaches intervene directly at the level of propagation.

Likewise, programmable environments can support integrated sensing and communication, localization assistance, and environment monitoring. Since the same programmable surfaces can influence and probe the EM field, they may serve not only communication objectives but also broader service objectives related to awareness of the space itself. This is fully consistent with the 6G direction, in which communication, sensing, and control are expected to become tightly intertwined.

The deeper promise of PWEs is therefore not merely stronger links. It is the creation of a programmable spatial infrastructure in which EM behavior can be orchestrated according to application goals, user priorities, and operator policies.

1.8 Motivation and Thesis Objectives

The discussion above reveals a central tension. On the one hand, programmable metasurfaces offer rich EM controllability. On the other hand, this controllability emerges from the microscopic state of many unit cells, under realistic hardware constraints, inside environments that must serve multiple users and services. Bridging these levels is the main motivation of this thesis.

The first challenge is *scalable control*. Online per-element optimization becomes increasingly expensive as the surface size grows and runtime conditions become more dynamic. This motivates offline compilation and codebook-oriented methodologies that preserve physical consistency while reducing online complexity.

The second challenge is *EM realism*. Many communication-oriented abstractions idealize the surface as a simple phase-control device. While analytically useful, such abstractions often omit coupling, amplitude effects, near-field behavior, quantization, and hardware restrictions. There is therefore a need for methods that remain grounded in more realistic EM models and in realistic configuration-synthesis workflows.

The third challenge is *resource sharing*. Surfaces are not dedicated to one user or one objective. They must serve heterogeneous demands under limited resources. This raises questions of multiplexing, scheduling, fairness, virtualization, and policy enforcement that cannot be addressed through isolated single-link optimization alone.

The fourth challenge is *multi-surface coordination*. Once multiple surfaces are deployed in the same environment, the problem becomes explicitly networked. Air-paths must be selected, surface functions must be coordinated, and end-to-end objectives must be achieved across an interconnected programmable space.

The fifth challenge is *service expansion*. Programmable environments make it possible to shape not only communication quality but also spatial exposure, secrecy, and covert delivery. This motivates the study of security-aware and covertness-oriented propagation services that are native to the programmable medium itself.

Taken together, these considerations define the objectives of this thesis: to develop scalable and physics-aware methods for programmable-surface control; to advance offline codebook compilation for efficient operation; to enable low-complexity multi-user sharing and multiplexing; to orchestrate multiple surfaces within a unified network framework; and to extend programmable propagation toward proactive security-oriented services.

1.9 Contributions

This thesis makes the following main contributions:

1. It develops physics-aware tools and methodologies for RIS configuration computation, with emphasis on realistic EM modeling and on the relation between microscopic control states and macroscopic functionalities. Within this context, it also investigates efficient configuration-generation strategies, including meta-heuristic and hybrid optimization methods.
2. It advances manufacturing-phase programmable-surface control through codebook compilation methodologies that precompute, organize, and validate realizable configurations for later reuse, thereby reducing runtime complexity while preserving consistency with realistic propagation behavior.
3. It formalizes programmable surfaces as sharable network resources and introduces low-complexity multiplexing and sharing mechanisms that support multi-user and multi-function operation under heterogeneous demands.
4. It extends control from isolated single-surface settings to environments with multiple RIS units, introducing coordination and path-selection mechanisms for priority-aware end-to-end orchestration across programmable air-paths.
5. It introduces RF-fencing as a RIS-assisted service for proactive covertness and spatially selective EM control, combining precompiled entries with lightweight online refinement to create friendly delivery regions and hostile suppression regions.

Overall, the thesis contributes a unified resource-management perspective for RIS-enabled environments, spanning realistic configuration synthesis, offline codebook generation, efficient online sharing, multi-RIS coordination, and security-oriented programmable propagation.

1.10 Thesis Structure

The remainder of the thesis is organized as follows:

- Chapter 2 reviews the related literature and highlights the main open challenges.
- Chapter 3 presents the tools and methods used for realistic RIS modeling, measurement, and optimization.
- Chapter 4 focuses on codebook compilation and offline configuration generation.
- Chapter 5 studies RIS sharing and multiplexing in multi-user and multi-function scenarios.
- Chapter 6 introduces the proposed framework for multi-RIS coordination.
- Chapter 7 presents RF-fencing as a programmable-environment service for covertness and spatial selectivity.
- Chapter 8 concludes the thesis and outlines future research directions.

CHAPTER 2

RELATED WORK

2.1 Optimization of RIS Configuration & Codebook Compilation

Achieving efficient and physically consistent RIS control requires configuring hundreds or thousands of unit cells under stringent latency constraints, which makes low-complexity runtime methods essential. Currently, much of the early RIS literature relies on online joint active–passive beamforming, e.g., jointly optimizing the transmitter precoder and RIS phases to maximize received power or user SINR under imperfect channel state information (CSI) [59, 60]. While effective in principle, such methods incur repeated channel estimation and iterative per-element adjustments, which are difficult to reconcile with low-latency services and large-scale surfaces.

In this thesis, we propose the control of RIS units via codebook schemes. Codebook compilation occurs primarily during manufacturing: an RIS prototype, together with receivers and transmitters, is placed in an automated positioning and evaluation setup—either physically realized or precisely simulated, as illustrated in Fig. 2.1. For each functionality of interest (i.e., those declared as 'supported' by the manufacturer), an optimization loop determines the RIS configuration. A finite set of these optimum

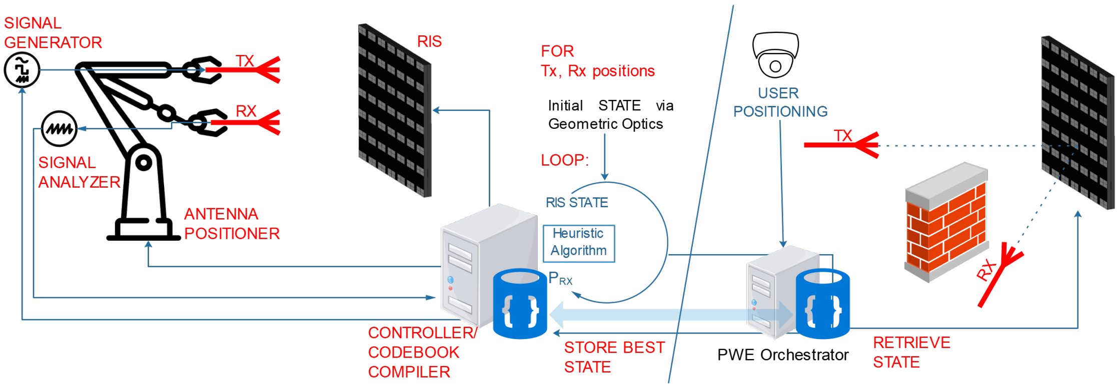

Figure 2.1: Illustration of the codebook compilation workflow during the RIS manufacturing (left), and its use during the RIS operation (right).

reflection patterns is stored in the codebook database for rapid deployment during operation, enabling timely adaptation without continuous channel estimation or iterative online optimization. This is particularly appealing in dynamic and multi-user settings, where frequent reconfiguration is needed to sustain QoS.

The first significant limitation in the research bibliography concerns physics consistency in the codebook compilation itself. Most codebook designs and compilation workflows adopt simplified propagation assumptions, frequently relying on far-field models that overlook near-field effects, where the RIS response is strongly shaped by mutual coupling among unit cells and by specular reflections from the surrounding environment [61, 62, 63]. Recent studies have begun to incorporate near-field modeling in codebook design [64] and coupling-aware formulations [65], yet these effects are commonly treated in isolation or through approximations, leaving open the question of how to compile and organize codebooks under joint near-field propagation, coupling, and environment interactions.

Codebook compilation is inherently a complex optimization problem, making dimensionality reduction crucial for limiting computational cost. Common practices include continuous relaxation followed by discrete quantization of element states [24], and exploiting symmetry, periodicity, and supercell grouping to shrink the search space [33]. Moreover, existing works employ hierarchical beam training to reduce search cost [66], learning-based mappings from measurements or user coordinates to codewords [67, 68], and structural constraints that reduce the effective dimensionality of the RIS control space [69].

In terms of the optimization procedure itself, genetic algorithms coupled with EM solvers are widely used due to their flexibility, despite known efficiency drawbacks [33], while iterative coordinate-wise optimizers are also common but can be slow, fail to converge, or get trapped in local optima [70, 71]. Machine-learning approaches typically rely on extensive simulated or measured data to accelerate parts of the workflow (e.g., early filtering of non-promising candidates), but they act as add-ons rather than replacing the underlying physics-driven optimization [72, 52].

Finally, an optimal configuration cannot be stored for every user location in continuous space, yet enlarging the codebook to better approximate that space makes online search and training overhead prohibitive. Environment-aware protocols further combine short training blocks with statistical CSI to generate candidate channels and select a suitable codeword from an offline-compiled database [73]. Near-field-specific codebooks have also been proposed for extremely large-scale RISs by matching cascaded array steering vectors, with exhaustive or hierarchical procedures for codeword selection [56]. These directions reduce training burden, but they often shift complexity into the selection mechanism (e.g., learned inference or non-trivial probing policies) and may require substantial datasets or protocol overhead. More importantly, they do not directly address how to structure a finite codebook so that near-optimal configurations can be selected scalably and with low online complexity over a continuous user space.

2.2 RIS as a Network Resource

Resource allocation in wireless networks aims to efficiently share finite resources among a time-varying set of users while meeting fairness and constraint objectives (e.g., energy and bandwidth)[74, 75]. In contemporary architectures, this objective is often operationalized through network slicing and virtual switching. Network slicing partitions radio and core resources to satisfy heterogeneous service requirements; recent works employ optimization and learning, including deep RL and AI-driven policies, to handle user mobility and maximize spectrum efficiency under resource scarcity[76, 77, 78]. Virtual switching extends this logic by virtualizing core-network functions and allocating resources across stakeholders (e.g., infrastructure, cloud, and slice owners), typically via iterative optimization and VNF placement strategies that account for deployment costs and bandwidth limitations [79, 80]. Translating these ideas to RIS-assisted networks, the notion of Virtual Programmable Metasurfaces (VPMs) models software-defined entities that encapsulate subsets of RIS elements, improving manageability and providing a natural abstraction for orchestration systems [81].

Within RIS-assisted networks, resource allocation has been studied both through (i) joint optimization of RIS configurations with network parameters and (ii) orchestration of multiple RIS units in PWEs. Learning-based and game-theoretic formulations have been used to optimize transmit power, spectrum allocation, and RIS states under dynamic conditions, targeting performance and energy-efficiency gains [82, 83, 84]. Other works consider multi-RIS coordination via external software services that compute and deploy interaction configurations to meet user demands, and employ multi-stage RL to balance energy efficiency and latency [85, 86]. Complementary contributions address heterogeneous settings where RIS supports traffic offloading, uplink mmWave, NOMA, or MIMO enhancements by jointly optimizing phase shifts and resource variables via iterative and fractional-programming-based methods [87, 88, 89]. While these frameworks demonstrate system-level gains, they typically optimize global metrics and seldom formalize user-specific differentiation aligned with stakeholder policies (e.g., tiering/pricing), nor do they explicitly interface with network slicing/RAN control using policy-native inputs and outputs [90].

A separate line of work addresses multi-user sharing of a single RIS by splitting resources in time or space. In PWEs, RIS multi-tasking has been explored via segmentation into tiles assigned to different tasks or via duty-cycling/TDMA-style sharing, where users are served in dedicated time slots [91, 92]. These approaches incur practical limitations: spatial segmentation reduces effective aperture per user and can degrade performance [93], whereas TDMA requires stringent synchronization and introduces latency due to per-slot hardware state transitions; additionally, users not scheduled in a slot obtain no RIS gain during that interval.

Finally, building PWEs with multiple RIS deployments requires efficient cascaded-RIS path selection, since multi-hop reflective links introduce a routing layer on top of per-surface configuration. For cascaded links under imperfect phase estimation, [94] proposes an SNR-maximizing strategy that balances beamforming gains against phase errors and provides outage/capacity characterizations. Related works decouple routing decisions from subsequent configuration/offloading steps: [95] addresses latency in multi-RIS-assisted mobile edge computing by selecting paths via graph-based shortest-path formulations that trade cascaded gains against reflection losses, while [96] develops analogous weighted-graph approaches in heterogeneous edge-IoT settings, selecting RIS–AP pairs by cascaded gain and routing under gain/loss tradeoffs. These contributions motivate viewing multi-RIS PWEs as coupled problems of (i) selecting the cascaded route and (ii) configuring the participating RIS nodes under system constraints and user demands.

2.3 Physical Layer Security in RIS-empowered Networks

Another property of PWE is the enhancement of privacy-aware networking. In this line, recent work has investigated RIS-assisted covert communication [97, 98]. Early studies showed that RIS configurations can create quiet zones by absorbing or redirecting energy, enabling covertness with low power overhead compared to noise-centric countermeasures [98]. Subsequent contributions largely focused on optimizing classical secrecy metrics under covert constraints, e.g., alternating optimization of beamforming and RIS phases [99] and MIMO extensions based on null-space processing to reduce leakage [100]. However, these formulations mainly treat RIS as an auxiliary knob for secrecy-rate enhancement, rather than as an environmental control primitive that can explicitly sculpt spatial exposure and enforce directional suppression—a core capability envisioned in PWEs.

To leverage RIS-enabled spatial degrees of freedom more directly, [101] proposed a full-duplex covert scheme in which an RIS supports a public downlink while concealing a covert uplink, revealing that multi-user interference can serve as a natural masking mechanism without artificial jamming. Related ideas were adapted to RIS-aided NOMA, where covert messages are embedded within superimposed downlink signaling by exploiting the public user's signal and RIS-induced phase uncertainty [102]. While these works highlight the value of interference and superposition for covertness, they do not align with the PWE architecture, where multiple distributed RIS are orchestrated to jointly structure the propagation environment [103, 104]. This leaves open the design of covert strategies that exploit networked RIS coordination to create spatially programmable covertness over extended areas, rather than optimizing a single localized link.

Recent research has begun to expand RIS-assisted covert communication either by enhancing the architecture of a single surface or by coordinating multiple RIS. On the architectural side, STAR-RIS—enabling simultaneous reflection and transmission—has been studied in multi-antenna and NOMA covert settings, showing that joint STAR-RIS and transmission optimization can improve covertness by shaping the spatial footprint and embedding hidden traffic within legitimate signals [105, 106]. Active RIS architectures with amplification have also been explored, aiming to boost covert link quality under strict detection constraints [107, 108]. However, these approaches do not, by themselves, justify moving away from nearly-passive designs, particularly when comparisons against cooperative multi-RIS deployments are absent. In contrast, [109] considers distributed RIS control via multi-agent reinforcement learning, indicating that multi-RIS coordination can enhance spatial suppression and covert throughput while retaining low-power operation expected in PWEs. Yet the current scope remains limited—typically a single covert link and narrow objectives—without fully addressing how multi-RIS systems can simultaneously support legitimate services and enforce directional suppression as an intrinsic, programmable feature of the environment. Overall, scalable and architecture-consistent covert communication for PWEs, grounded in distributed RIS coordination and explicit spatial exposure control, remains insufficiently developed.

2.4 Main Challenges and Open Issues

As concerns codebook compilation, several open issues remain. First, developing tools and algorithms grounded in accurate EM models or full-wave analysis is essential for realism and direct deployability; beyond realism, such modeling offers insights into RIS behavior that can be exploited to reduce optimization dimensionality and improve codebook-entry quality. Second, an exhaustive benchmarking of available heuristic/metaheuristic methods—and their potential synergies—should be undertaken to move beyond arbitrary optimizer selection. Finally, techniques that avoid running exhaustive optimization at every possible user position during compilation are crucial for realistic RIS use within PWEs.

During operation within the PWE, as with other network components, each RIS must simultaneously serve multiple users with heterogeneous requirements. In networking terms, an RIS is a shareable resource, yet a precise and operational definition of RIS-as-a-resource—including what is allocated, at what granularity (elements/tiles/time/frequency), and under which constraints (hardware switching latency, control overhead, power budget, and inter-user coupling)—remains largely missing. Consequently, translating high-level stakeholder policies (e.g., tiering/pricing) and user QoS demands into actionable RIS control decisions at runtime is still an open problem, particularly under tight latency budgets that preclude heavy online optimization. This challenge becomes even sharper in cascaded-RIS links, where the PWE is naturally modeled as a graph [18] and resource allocation is coupled with route selection; however, without incorporating explicit fairness and differentiation across users and tiers, graph-based formulations capture only part of the required orchestration logic.

Finally, in the context of physical-layer security within PWEs, most RIS-assisted covert strategies still regard the RIS as a tuning knob for covertness enhancement rather than a software-defined mechanism for shaping spatial exposure via programmable propagation control. Many solutions also rely on iterative optimization or learning-based inference, which is often too demanding for real-time security actions and typically validated only numerically, without tight coupling to physics-consistent tools that would improve fidelity and deployability. A codebook-driven RIS operation offers an implementation-aligned alternative: optimized configurations are precompiled offline and retrieved at negligible runtime cost, while codebook-entry multiplexing combines simple precompiled patterns into complex responses on demand.

CHAPTER 3

MODELING, MEASUREMENT METHODS & OPTIMIZATION TOOLS

The development of services and tools for RIS deployment in realistic environments requires accurate modeling of RIS-enabled wireless propagation. This chapter presents the main tools and methods used in this thesis for the simulation of RIS-assisted PWEs and the validation of the proposed solutions. First, physics-based EM frameworks are presented. In addition, the Finite-Difference Time-Domain (FDTD) method is employed, widely recognized for its accuracy. The RIS modeling approach and the tool used for this part are described. The evaluation is supported by measurements on a real testbed, selected results of which are included. Finally, for completeness, we present the optimization algorithms used at different stages of the RIS lifecycle.

3.1 Physics-based EM Modeling of RIS-enabled Environments

Physics-based EM modeling consists of three different approaches. First, the RIS behavior is modeled in the far field, assuming plane-wave propagation from an angle of arrival (*AoA*) to an angle of departure (*AoD*). Next, near-field modeling is presented for confined environments: the transmitter is placed at a specific position in 3D space with spherical-wave transmission to another point. Finally, the near-field model is enhanced with additional aspects—mutual coupling and secondary reflections—for more accurate results.

3.1.1 Far-Field

The far-field regime is appropriate for simulating RIS behavior when the scattered field is observed at distances much larger than both the RIS aperture and the operating wavelength. Under these conditions, the EM behavior is dominated by angular distributions [110], so precise control of propagation directions is essential for spatial selectivity. To model these angular characteristics, the Huygens-Fresnel principle (HFP) is adopted [111], treating the RIS as an ensemble of secondary point sources that re-radiate the incident wavefront in all directions. The RIS is represented as a two-dimensional array with $N_{\text{el}} = N_r \times N_c$ reconfigurable unit cells. Each unit cell acts as an isolated scatterer that applies a controllable phase shift to the incident wavefront, with the phase profile

$$\mathbf{\Phi} = [\Phi_n] \in [0, 2\pi]^{N_r \times N_c}, \;\; n = 1, \ldots, N_{el} \tag{3.1}$$

where Φ_n represents the reflection phase applied by the n_{th} unit. Accordingly, the scattered field produced immediately after interaction with the RIS can be expressed as

$$E_{\text{scat},n} = E_{\text{inc},n} e^{-j\Phi_n}, \tag{3.2}$$

where $E_{\text{inc},n}$ denotes the incident field at the n_{th} unit cell. Per HFP, each unit emits a spherical wave, and the superposition of these waves sets the field in any desired direction, enabling dynamic shaping of the spatial energy distribution.

To specify the geometry, assume the RIS is illuminated by a uniform plane wave of wavelength λ arriving from the far-field direction (θ_i, φ_i). The elements lie on a rectangular grid with uniform spacings Δ_x and Δ_y along the horizontal and vertical axes, respectively. Thus, the position of each unit cell is given by

$$x_n = \Delta_x(x - x_c), \quad y_n = \Delta_y(y - n_c), \tag{3.3}$$

where (x_c, y_c) indicates the coordinates of the array center, serving as the reference point for calculating phase differences across the surface. The incident wavefront impinging upon each cell introduces a phase variation that depends on the cell's position relative to the array center, calculated as

$$\psi_n^{\text{inc}} = k\left[\Delta_x(x - x_c)\cos\varphi_i + \Delta_y(y - y_c)\sin\varphi_i\right]\sin\theta_i, \tag{3.4}$$

where $k = 2\pi/\lambda$ is the wavenumber associated with the operational frequency. By leveraging both the RIS-imposed phase shift and the spatially dependent incident phase, the total contribution from each unit cell is formulated as

$$\Psi_n = e^{j\left(\Phi_n + \psi_n^{\text{inc}}\right)}. \tag{3.5}$$

Subsequently, the far-field scattered response of the RIS is obtained by summing the contributions of all unit cells, yielding

$$E(\theta, \varphi) = E_0 \cos^{2\rho}(\theta) \sum_{n=1}^{N_{el}} \Psi_n, \tag{3.6}$$

where E_0 is a normalization constant related to the incident field amplitude (measured in V/m), and ρ characterizes the angular selectivity of the scattering pattern. Specifically, $\rho = 0$ corresponds to an isotropic response, while larger values of ρ sharpen the angular focus. To quantify the effectiveness of a given RIS configuration in steering energy toward a specific observation direction (θ_d, φ_d), the field magnitude at the corresponding point of interest (POI) is evaluated as

$$E_{\text{POI}} = \left|E(\theta_d, \varphi_d)\right|, \tag{3.7}$$

where $E(\theta_d, \varphi_d)$ represents the scattered field computed using (3.6) at the desired angles. Finally, under the far-field approximation, POIs correspond to angular sectors (directional cones) measured in steradians, not physical locations within a spatial volume. Depending on the operator's preference at a given time, POIs are used either to maximize signal strength for beam steering/focusing or to minimize it for suppression in physical layer security applications.

3.1.2 Near-Field

When RIS behavior must be simulated in confined environments, the Fresnel–Kirchhoff diffraction (FKD) framework accurately captures intricate wave interactions [110]. Unlike far-field models that rely on angular propagation, near-field modeling resolves the spatial distribution of EM fields, for both signal strength maximization and minimization in a specific area within the communication network. For the near-field model, we assume that the network is a cubic environment of side length L. For the accurate EM computation the environment is splitted in a grid-based domain. For the analysis, we adapt a dual-grid scheme that balances efficiency and local accuracy: a coarse grid with parameters N_x, N_y, N_z uniformly discretizes $(x, y, z) \in [0, L]^3$, while a finer grid—constructed by refining the coarse spacing—covers a spherical zone of interest (ZoI) of radius r_{ZoI} centered at (x_c, y_c, z_c). This refinement captures field variations within the ZoI while keeping computation manageable.

The four vertical walls of the space can be covered by RIS units. In a full-covered case, each RIS is composed of $N_{\text{el}} = N_r \times N_{\text{c}}$ uniformly distributed reflecting elements. These elements are arranged with a margin m from the space boundaries, resulting in a total of $N_{\text{RIS}} = 4N_{\text{el}}$ scatterers, where the position of the nth RIS element is given by (x_n, y_n, z_n). Due to the significant physical dimensions of the RIS surfaces relative to the space size and operating wavelength, the far-field approximation becomes invalid, and a spherical-wave near-field modeling based on FKD is adopted to ensure accurate field estimations within the confined space. Thus, assuming that the source is located inside the room at (x_s, y_s, z_s) and emits an omnidirectional spherical EM wave at frequency f, the scattered field within the space is calculated on both grids given by:

$$E_{\text{inc},n} = \frac{E_0}{d_n} e^{jkd_n}, \tag{3.8}$$

where E_0 represents the amplitude of the source field, and d_n denotes the distance between the source and the nth RIS element, which is equal to

$$d_n = \sqrt{(x_n - x_s)^2 + (y_n - y_s)^2 + (z_n - z_s)^2}. \tag{3.9}$$

Each RIS element reflects the incident wave by applying a controllable phase shift Φ_n, while maintaining a unitary reflection magnitude. Consequently, the complex reflection coefficient for the nth element is defined as

$$\Gamma_n = e^{j\Phi_n}. \tag{3.10}$$

The contribution of the nth RIS element to the scattered field at each observation point is approximated by spherical-wave propagation as

$$E_{\text{scat},n}(x_p, y_p, z_p) = E_{\text{inc},n}\Gamma_n \frac{e^{jkR_n(x_p,y_p,z_p)}}{R_n(x_p, y_p, z_p)}, \tag{3.11}$$

where $R_n(x_p, y_p, z_p)$ denotes the propagation distance from the nth RIS element to the observation point, computed as

$$R_n(x_p, y_p, z_p) = \sqrt{(x_p - x_n)^2 + (y_p - y_n)^2 + (z_p - z_n)^2}. \tag{3.12}$$

The total scattered field E_{scat} at each grid point is then obtained by summing the contributions from all RIS elements, according to

$$E_{\text{scat}}(x_p, y_p, z_p) = \sum_{n=1}^{N_{\text{RIS}}} E_{\text{scat},n}(x_p, y_p, z_p). \tag{3.13}$$

For the purposes of this model, each RIS element is treated as an ideal isotropic point scatterer with individually adjustable reflection phase, while mutual coupling and polarization effects are neglected. Additionally, reflections from uncovered surfaces and secondary reflections among RIS-covered walls are disregarded. The source is assumed to emit a continuous spherical wave at a fixed frequency, maintaining constant amplitude throughout the simulation domain.

Using this modeling framework, we further define a zone of interest in the room as a spatial region where the scattered field is minimized. Specifically, the *ZoI* is formally defined as the spherical region satisfying

$$(x_p - x_{\text{c}})^2 + (y_p - y_{\text{c}})^2 + (z_p - z_{\text{c}})^2 \leq r_{\text{ZoI}}^2, \tag{3.14}$$

where (x_p, y_p, z_p) corresponds to the coordinates of each observation point. For the observation points that satisfy (3.14), the calculation of (3.13) is conducted using the finer and denser computational grid for more accurate results. To quantify the EM field strength within this region, we evaluate the average magnitude of the scattered field, denoted as $E_{\text{field}}^{\text{avg}}$. This metric directly reflects the residual field energy inside the *ZoI* and is computed as

$$E_{\text{field}}^{\text{avg}} = \frac{1}{N_{\text{ZoI}}} \sum_{(x_p,y_p,z_p)\in\text{ZoI}} |E_{\text{scat}}(x_p, y_p, z_p)|^2, \tag{3.15}$$

where N_{ZoI} denotes the total number of discretized grid points, and $|E_{\text{scat}}(x_p, y_p, z_p)|$ represents the magnitude of the scattered field at each corresponding point. This formulation enables a detailed volumetric analysis of the field distribution within enclosed environments and serves as the foundation for developing RIS control strategies.

3.1.3 Realistic RIS Modelling

Since the model of the previous section neglects additional propagation effects in RIS-enabled environments—most notably mutual coupling between adjacent RIS elements and secondary reflections—we introduce an enhanced model for the thesis parts that require higher-fidelity RIS characterization. Also in this case, we consider a three-dimensional rectangular cavity of edge length L_{th} aligned with the Cartesian axes. Each wall can be optionally coated by a RIS panel consisting of $N = N_\text{d} \times N_\text{d}$ square unit cells gird with center-to-center spacing d. This is typically chosen on the order of $\leq \lambda/4$ to satisfy the Nyquist sampling criterion [112]. We consider that $\mathbf{p}_n = (x_n, y_n, z_n) \in [0, L_{th}]^3$ is the center of the n^{th} element and $\mathbf{n}_w$ denotes the outward normal vector of the wall to which the element belongs. The transmitter sits at position $\mathbf{s} \in [0, L_{th}]^3$, radiating at wavelength $\lambda = c_0/f$ with wavenumber $k = 2\pi/\lambda$. Each RIS element acts as a passive scatterer characterised by a complex reflection coefficient $\Gamma_n = \rho_n e^{\mathrm{i}\phi_n}$, where the amplitude ρ_n is typically set to unity and the phase ϕ_n is tunable.

The incident field at element n arising from the transmitter is modelled as a spherical wave with amplitude decay and free-space phase advance [110]. Denoting the vector from the transmitter to the element by $\mathbf{d}_n = \mathbf{p}_n - \mathbf{s}$ and its norm by $r_n = \|\mathbf{d}_n\|$, the direct term reads

$$E_{\text{dir},n} = \frac{e^{\mathrm{i}kr_n}}{r_n} \underbrace{\left|\hat{\mathbf{d}}_n \cdot \mathbf{u}_\text{TX}\right|^m}_{\text{transmitter pattern}} \underbrace{\left[\max\{0, (\hat{\mathbf{d}}_n \cdot \mathbf{n}_w)\}\right]^p}_{\text{element cosine law}}, \tag{3.16}$$

where $\hat{\mathbf{d}}_n = \mathbf{d}_n/r_n$ is the unit vector pointing from the transmitter to the element, $\mathbf{u}_\text{TX}$ is the unit vector of the main lobe of the transmitter, m controls the beamwidth of the transmitter radiation pattern and is typically chosen such that $|\hat{\mathbf{d}}_n \cdot \mathbf{u}_\text{TX}|^m$ approximates a raised cosine lobe, and p is the order of the cosine law used to model the effective aperture of the patch with respect to the angle of incidence. When $p = 1$ the element obeys the classical Lambert cosine law; larger values of p produce a more specular response. The transmitter phase centre offset, if any, is absorbed into $\mathbf{s}$.

One major deviation in respect of the previous model is that this one takes into account the mutual coupling among RIS elements. This phenomenon arises because the field scattered by one element can re-illuminate its neighbors. Taking into account that a full EM solution of the Maxwell integral equations is computationally prohibitive for large arrays, we adopt a discrete model that captures the dominant short-range interactions. Each element interacts primarily with its immediate neighbors on the same panel. Denote by $\mathcal{N}_n$ the set of up to four direct neighbors of element n in the two-dimensional lattice. The coupled illumination received by element n is expressed as

$$E_{\text{cpl},n} = \alpha \sum_{m \in \mathcal{N}_n} \Gamma_m E_{\text{inc},m} \frac{e^{\mathrm{i}k\|\mathbf{p}_n - \mathbf{p}_m\|}}{\|\mathbf{p}_n - \mathbf{p}_m\|} \tag{3.17}$$

where α is a dimensionless coupling coefficient that quantifies the strength of these near-field interactions. The sum captures the fact that each neighbor acts as a secondary source whose contribution decays with distance and accumulates a phase delay. Besides neighbor coupling, the room boundaries create additional illumination paths via specular reflections. When a wall is not coated by a RIS it is modeled as a smooth reflector. The signal from the transmitter reflects off such a wall before reaching the RIS, and this secondary path contributes to the incident field. Let $\mathbf{s}_w$ denote the mirror image of the transmitter with respect to wall w and let $\mathbf{n}_w$ be that wall's normal. The secondary contribution arriving at element n from wall w is

$$E^{(w)}_{\text{sec},n} = \beta_w \frac{e^{\mathrm{i}k\|\mathbf{p}_n - \mathbf{s}_w\|}}{\|\mathbf{p}_n - \mathbf{s}_w\|} \left|\hat{\mathbf{d}}_{n,w} \cdot \mathbf{n}_w\right|^p \tag{3.18}$$

where $\hat{\mathbf{d}}_{n,w}$ is the unit vector from $\mathbf{s}_w$ to $\mathbf{p}_n$ and β_w is the reflectivity of wall w (for perfect electrical conductors $\beta_w = 1$). The exponent p is the same as in the direct term, ensuring that grazing incidence is de-emphasized. All secondary reflections over the available bare walls are included; the total incident field is the sum of the direct illumination, the coupled contributions and the secondary reflections. Putting these pieces together, the incident field on element n is

$$E_{\text{inc},n} = E_{\text{dir},n} + E_{\text{cpl},n} + \sum_{w \in \mathcal{W}} E^{(w)}_{\text{sec},n} \tag{3.19}$$

where $\mathcal{W}$ indexes the reflecting walls. Because $E_{\text{cpl},n}$ depends on the unknown coefficients Γ_m, the problem is nonlinear and the incident fields must be recomputed whenever the configuration changes. Once the E_{inc} arrives in RIS surface, each element re-radiates it with phase shift Γ_n. The complex electric field at an arbitrary observation point $\mathbf{r} \in [0, L_{th}]^3$ is obtained by superposing the contributions of all RIS elements:

$$E(\mathbf{r}) = \sum_{n=1}^{N} \Gamma_n E_{\text{inc},n} \frac{e^{\mathrm{i}k\|\mathbf{r}-\mathbf{p}_n\|}}{\|\mathbf{r}-\mathbf{p}_n\|} \tag{3.20}$$

This model implicitly enforces the boundary conditions on the covered walls via the cosine factors in Eq. (3.16).

As a result, the received signal quality at a given location is quantified through the signal-to-noise ratio (SNR), which is obtained from the squared magnitude of the total electric field evaluated at the receiver position can be computed as:

$$\text{SNR} = 10 \log_{10}\left(|E(\mathbf{r})|^2\right) \tag{3.21}$$

where, we consider that the noise is normalized.

3.2 FDTD Simulations

The FDTD method is a numerical technique solving Maxwell's equations in the time domain to model EM wave propagation through arbitrary structures. FDTD offers time-resolved, spatially detailed EM field evolution, making it suitable for simulating complex sub-wavelength phenomena such as mutual coupling, scattering, and resonance effects significant in RIS design.

FDTD's ability to model wave interaction with intricate RIS unit cells, including tunable elements (e.g., PIN diodes, varactors), enables accurate characterization of EM parameters essential for RIS optimization. Parameter sweeps within FDTD facilitate exploring how variations in unit cell geometry, material properties, or control states affect the macroscopic RIS behavior. This capability is crucial for designing RIS unit cells with tailored impedance responses and reflection/transmission properties.

Notably, FDTD simulations guide RIS design by quantifying near- and far-field responses and confirming theoretical models of RIS behavior under realistic scenarios incorporating mutual coupling and scattering by environmental objects. Such physics-based simulations augment system-level optimizations by providing EM consistent inputs rather than relying solely on idealized or statistical channel models.

The advantages of FDTD usage in RIS modeling compared with other methods (full-wave solvers, ray-tracing, statistical channel modeling) include [113]:

- Mutual coupling effects among RIS elements, which impact the element's effective scattering parameters. This is paramount for large RIS arrays, where ignoring coupling leads to suboptimal designs.

- Frequency-selective and angle-dependent responses arising from unit cell geometry and tuning elements, which critically affect overall RIS beamforming capabilities.

- Transient behavior and broadband response, offering insights into the RIS performance over frequency bands relevant to 5G/6G systems and beyond.

However, FDTD requires significant computational resources, especially for electrically large RIS structures or high-frequency scenarios. Careful meshing and time-stepping are needed to ensure numerical stability and accuracy, which can lead to long simulation times. Moreover, it requires strong expertise in EM theory and numerical methods to set up and interpret results correctly.

The main metric used during FDTD simulations is transmission coefficient S_{21} between the receiver and the transmitter towards the RIS path. Generally, S_{21} is a standard scattering measure that characterizes transmission between ports in multiport networks. In RIS unit cells and arrays, S_{21} captures how an incident wave couples through or around the element toward another port or direction, encoding transmittance magnitude and phase. Accurate S_{21} knowledge quantifies power transfer and phase shift imparted by an element, which is critical for configuring the RIS to shape the EM environment. S_{21} data integrate naturally into multiport network models that treat RIS as coupled scattering elements, enabling realistic computation of end-to-end channel characteristics in RIS-assisted links.

In terms of this thesis, the open source FDTD solver, openEMS is used [114]. OpenEMS is a widely used EM field solver implementing the FDTD method. It excels in simulating complex electromagnetic problems including RIS unit cells, metasurfaces and antenna arrays. Briefly, the main advantages of openEMS usage are:

- **Flexible Geometry and Material Modeling**: openEMS supports complex 3D geometries and non-linear materials, enabling detailed modeling of RIS element shapes, substrates, and tuning elements.

- **S-Parameter Extraction**: The tool provides native calculation of S-parameters, including S21 (forward transmission coefficient), forming a fundamental basis for characterizing RIS element scattering behavior.

- **Parametric Sweeps and Optimization Integration**: Users can automate parameter sweeps facilitating optimization loops for RIS design variables such as element dimensions or bias voltages controlling phase shifts.
- **Compatibility with MATLAB**: openEMS integrates with MATLAB, enabling seamless coupling with optimization algorithms and system-level analysis platforms.
- **Community-Driven Development**: Being open-source, openEMS benefits from continuous improvements, making it cost-effective and customizable for RIS research and prototyping.

In this thesis, the RIS unit is inspired by [24]. This design, depicted in Fig. 3.1, embodies a controlled departure from simple RIS-reflectarray approach, by interconnecting the horizontally adjacent cells with tunable loads. The RIS unit cells are gaps between square metal patches that are bridged by lumped complex impedance loads. The patches lie on a metal-backed dielectric substrate and a thin metal sheet that acts as a groundplane. Specifically, the substrate is Rogers RT/Duroid 5880, with electric permittivity, ε_r, 2.2 tangent loss, $\tan\delta$, 0.0009 and thickness of 1.016 mm. The groundplane is a perfect conductor, akin to the square patches, and possesses a thickness of 17.5 μm. This cell design negates transmission, provides wide angular and bandwidth, while the value of its embedded load imparts control over the amplitude and phase of reflected wave in the x polarization.

Table 3.1: RIS design dimensions.

Parameter	**Value**
Width of unit-cell substrate (Ly)	11.25 mm ($= 0.15\lambda$)
Length of unit-cell substrate (Lx)	11.25 mm ($= Ly$)
Gap between patches (d)	1 mm
Patch width (w)	10.25 mm
Thickness of patch	17.5 μm
Tunable element width	1 mm

The RIS unit has been designed to operate at 4 GHz, and after the fine-tuning procedure, the dimensions of the patches are presented in Table 3.1. The tunable elements

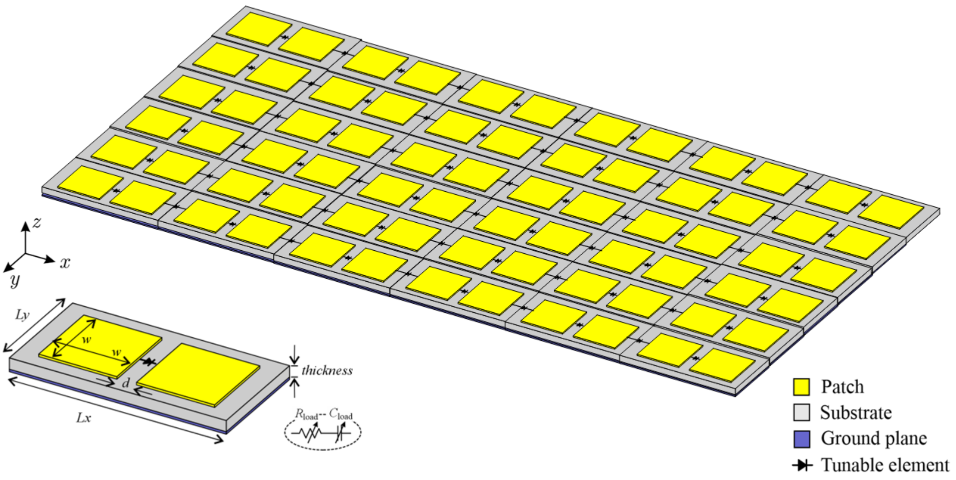

Figure 3.1: Overview of the RIS unit's element that is used for simulation in this thesis.

connecting the horizontal adjacent patches can simulate any type of integrated circuit chips that could be used as configuration mechanisms in the RIS. In our approach, the tunable loads assume continuous values in the complex plane, i.e., for both reactance and resistance. Note that the tunable load is strategically placed between each pair of horizontally adjacent patches, thus modifying the surface impedance produced by the combination of gap capacitance and the grounded dielectric inductance.

3.3 Real Testbed

Apart from the simulation tools and methods, results from a real, experimental results are presented for the validation of the developed algorithms and components. The experimental testbed comprises two Software-Defined Radios (SDRs)—one acting as the base station and the other as a network user—both implemented with USRP B210 units.[1] The USRP B210 is a full-duplex SDR platform that provides continuous RF coverage from 70 MHz to 6 GHz, supports 2×2 MIMO operation through the Analog Devices AD9361 RFIC, and offers up to 56 MHz of instantaneous bandwidth. Its FPGA (Xilinx Spartan-6 XC6SLX150) enables real-time baseband processing, while connectivity to the host computer is achieved via SuperSpeed USB 3.0. Each SDR in

---

[1] `https://www.ettus.com/all-products/ub210-kit/`

Figure 3.2: Overview of the real RIS testbed.

the testbed is equipped with Wi-Fi antennas configured to operate at 4.2 GHz. The main specifications of the USRP B210 are summarized in Table 3.2.

GNU Radio[2] serves as the signal processing framework of the testbed. It is an open-source software toolkit widely used for rapid prototyping of communication systems. GNU Radio provides a modular flow-graph architecture where signal processing operations (e.g., filtering, modulation, synchronization, channel estimation) are implemented as blocks that can be linked to form complex transceiver chains.

---

[2] `https://www.gnuradio.org/`

Table 3.2: USRP B210 Specifications

Feature	Specification
Frequency Range	70 MHz – 6 GHz
RF Chains	2 Tx, 2 Rx (2×2 MIMO)
Instantaneous Bandwidth	Up to 56 MHz
RF Front-End	AD9361 transceiver
FPGA	Spartan-6 XC6SLX150
Connectivity	USB 3.0 (SuperSpeed)
Host Software	UHD driver, GNU Radio integration

This flexibility allows seamless integration with USRP hardware via the UHD driver, enabling real-time baseband processing and over-the-air experiments. Furthermore, GNU Radio supports multi-threaded execution and GPU acceleration when coupled with high-performance host platforms.

The host computers employed are NVIDIA Jetson AGX Orin units,[3] chosen for their powerful parallel processing capabilities and compatibility with GNU Radio's Linux-based ecosystem. The Jetson AGX Orin features an NVIDIA Ampere GPU with up to 2048 CUDA cores and 64 Tensor Cores, delivering several TFLOPs of FP32 compute performance. This hardware acceleration is critical for managing the computationally intensive workloads of SDR-based 5G/6G experiments. Table 3.3 summarizes the specifications of the Jetson AGX Orin.

Table 3.3: NVIDIA Jetson AGX Orin Specifications

Feature	Specification
GPU	Ampere architecture, 2048 CUDA cores, 64 Tensor Cores
CPU	12-core ARM Cortex-A78AE v8.2 64-bit
Memory	Up to 64 GB LPDDR5
Storage	NVMe SSD support
Connectivity	PCIe Gen4, USB 3.1, Ethernet
OS Support	Linux (JetPack SDK), GNU Radio compatible

---

[3] `https://www.nvidia.com/content/dam/en-zz/Solutions/gtcf21/jetson-orin/nvidia-jetson-agx-orin-technical-brief.pdf`

The RIS hardware integrated in the testbed is the TMYTEK XRifle Dynamic RIS,[4] designed to operate within the first 5G frequency band. It consists of a 16×16 array of PIN diodes that allow binary phase configurations—the most restrictive and challenging case for multiplexing schemes. The device covers 4.2–5.2 GHz, supports linear polarization, and provides incidence/reflection steering capabilities from -60° to $+60^\circ$ along both vertical and horizontal planes. Table 3.4 lists the main specifications of the RIS.

Table 3.4: TMYTEK XRifle Dynamic RIS Specifications

Feature	**Specification**
Array Size	16×16 PIN diodes
Control Scheme	Binary phase configuration
Frequency Range	4.2 – 5.2 GHz
Polarization	Linear
Incidence/Reflection Range	-60° to $+60^\circ$ (vertical and horizontal)
Target Application	5G sub-6 GHz experimental platforms

An overview of the complete experimental testbed is shown in Fig. 3.2. The measurement campaign is conducted in two phases: RIS-off baseline collection and RIS-on evaluation. In the baseline phase, the transmitter and each receiver are positioned at predefined angular offsets relative to the RIS center, while the RIS remains deactivated. GNU Radio is employed to capture the in-phase (I) and quadrature (Q) samples of the received waveform. From these samples, the instantaneous amplitude is computed as

$$A[n] = \sqrt{I[n]^2 + Q[n]^2}, \quad n = 1, 2, \ldots, N, \tag{3.22}$$

where N denotes the total number of acquired samples. This baseline reflects both the direct LoS components and any uncontrolled multipath reflections from the inactive RIS surface.

In the evaluation phase, the RIS is activated using manufacturer-provided optimal configurations tailored for each transmitter–receiver angular pair. To quantify the performance enhancement attributable to the RIS, we compare the largest 10,000 amplitude values obtained in the RIS-off case with those from the RIS-on case. This

---

[4] `https://tmytek.com/products/components/xrifle-dynamic-ris`

approach mitigates the influence of small-scale fading and transient fluctuations, focusing instead on the most representative received power levels. The gain for each user k is defined as

$$G_k = \frac{\max_{n \in \mathcal{S}_{\text{on}}} A[n]}{\max_{n \in \mathcal{S}_{\text{off}}} A[n]}, \tag{3.23}$$

where $\mathcal{S}_{\text{on}}$ and $\mathcal{S}_{\text{off}}$ denote the sets of $10{,}000$ highest-amplitude samples under RIS-on and RIS-off conditions, respectively. This metric isolates the gain introduced by the RIS and enables fair comparison across different link geometries.

3.4 Terminology

For consistency, we use the terms *emulations*, *simulations*, and *measurements* to refer to the physics-based EM frameworks, full-wave FDTD (openEMS), and real testbed experiments, respectively (Table 3.5). In addition, for rapid debugging and early-stage evaluation, we employ *tests*, where the RIS phases are drawn i.i.d. from $[0, 2\pi)$. This quick-test baseline is used for fast indicative results, while the final evaluation relies on the more accurate EM models and/or measurements.

Table 3.5: Reminder terminology used in this thesis.

Term	**Meaning in this thesis**
Emulations	Physics-based EM frameworks (far-/near-field).
Simulations	Full-wave FDTD (openEMS).
Measurements	Real testbed experiments.
Tests	Fast baseline using i.i.d. phases $\Phi_n \sim \mathcal{U}(0, 2\pi)$ for quick testing.

3.5 Optimization Tools in RIS-enabled Networks

The usage of optimization methods is essential in the RIS lifecycle, especially in the codebook compilation and the determination of the optimal RIS configuration for specific functionality within the network. The main two categories of the used optimization algorithms for RIS configuration and codebook compilation are the classical gradient-based methods and population-based evolutionary algorithms.

Gradient-based optimization operates in the continuous domain, typically leveraging differentiable objective functions and constraints to iteratively improve RIS configurations using gradient information. These methods can exploit problem smoothness to converge rapidly but are often vulnerable to local optima due to the highly non-convex nature of RIS optimization landscapes. Additionally, handling the unit-modulus constraint of RIS configurations requires specialized techniques such as projected gradient methods or optimization on manifolds [115].

On the other hand, evolutionary algorithms encompass metaheuristic search processes inspired by natural evolution, such as Genetic Algorithms (GA) and multi-objective evolutionary algorithms including Strength Pareto Evolutionary Algorithm 2 (SPEA-2) and Non-dominated Sorting Genetic Algorithm II (NSGA-II). These population-based methods explore the solution space robustly without relying on gradient information, making them particularly suitable for discrete, quantized phase settings and multi-objective formulations. While they offer a superior ability to avoid local minima traps, they generally incur higher computational costs and require careful parameter tuning for mutation rates, population sizes, and selection pressures [116].

Each optimization category offers distinct advantages and limitations. Gradient-based techniques provide fast convergence and modest computational complexity but may fail in global exploration. Evolutionary algorithms ensure broader search coverage and multi-objective optimization but face scalability and processing time issues in large-scale RIS scenarios. Understanding these trade-offs is crucial for designing efficient hybrid or adaptive optimization frameworks tailored for the codebook compilation. Specifically, the following four optimization algorithms are used in this thesis component for RIS configuration optimization and codebook compilation [71]:

1. **The Gradient Descent Algorithm (GDA)**: GDA represents one of the foundational optimization techniques applied in RIS configuration due to its iterative, conceptually straightforward nature. It operates by computing the gradient of an objective function—often related to the defined communication metric with respect to RIS configuration and updating parameters in the direction of steepest descent. For RIS, the optimization problem commonly involves continuous configurations subject to the unit-modulus constraint (i.e., the magnitude of each reflection coefficient is fixed at unity), which renders the problem non-convex

and necessitates adaptations to the standard gradient descent method.

Despite its attractiveness, gradient descent approaches must carefully select step sizes and initialization points due to the possibility of becoming trapped in local minima inherent to RIS optimization tasks. They are particularly effective when integrated with joint waveform and beamforming optimization, where the continuous parameters can be tuned cohesively to maximize overall system objective functions. Overall, gradient-based methods remain integral to RIS configuration due to their computational efficiency and theoretical grounding [116].

2. **Genetic Algorithm (GA):** GA is a prominent evolutionary algorithm inspired by the principles of natural selection and genetics, widely utilized for complex optimization problems including RIS configuration. In the context of RIS, GA represents configurations (e.g., value of each RIS unit cell) as chromosomes—strings typically encoding the discrete or quantized phase states. A population of candidate solutions evolves over generations via selection, crossover (recombination), and mutation, with fitness evaluation guiding survival and reproduction. This iterative process explores the search space by balancing exploitation of high-quality solutions and exploration via genetic diversity mechanisms [117].

   Key GA operators for RIS include carefully designed selection strategies to prioritize promising configurations, crossover to exchange genetic material between solutions promoting diversity, and mutation introducing random variation to prevent premature convergence. Fitness functions quantify how well a configuration enhances system goals such as communication throughput, transmission coefficient, or combined multi-objective trade-offs. The population-based search does not rely on gradient information, enabling GA to handle highly nonlinear, discrete, or even non-differentiable problems naturally arising in RIS phase quantization [118].

3. **Strength Pareto Evolutionary Algorithm 2 (SPEA2):** The SPEA2 is a state-of-the-art multi-objective evolutionary algorithm widely adopted for problems involving competing objectives, such as RIS optimization for beam steering in a direction and suppression in another one for physical layer security of the transmitted information. SPEA-2 improves upon its predecessor by incorporating enhanced fitness assignment based on dominance strength, density estimation,

and elitism through an external archive that stores non-dominated solutions, preserving diversity and promoting convergence toward the Pareto front [118].

SPEA-2 calculates individual fitness by considering both the number of solutions it dominates (”strength”) and the density of solutions in its vicinity, allowing the maintenance of high-quality yet diverse solution sets. The external archive facilitates elitism by retaining historically best solutions across generations, avoiding regression in optimization performance. Together, these features enable SPEA-2 to approximate Pareto fronts effectively in complex and high-dimensional solution spaces typical of RIS configuration under multiple conflicting goals.

4. **Non-dominated Sorting Genetic Algorithm II (NSGA-II)**: The NSGA-II is a widely acclaimed multi-objective evolutionary algorithm renowned for its efficient non-dominated sorting approach and crowding distance mechanism that jointly preserve convergence pressure toward the Pareto front and maintain population diversity. NSGA-II operates by categorizing solutions into fronts based on dominance levels and further ranking them using crowding distance to balance exploration-exploitation trade-offs effectively [117]. The algorithm’s fast sorting method significantly reduces computational requirements relative to earlier multi-objective evolutionary methods. NSGA-II has found broad application in wireless communications and related fields due to its robust performance, scalable design, and ease of parameter tuning.

Whenever run-times are reported in this thesis, all experiments are implemented in MATLAB and executed on a 13th Gen Intel® Core™ i9-13900F CPU with an NVIDIA GeForce RTX 4090 GPU.

Chapter 4

Codebook Compilation

In this section, we present the developed methods for physics-based codebook compilation. First, we introduce the Physics-informed Codebook Compilation Software (PICCS), whose evaluation is based on simulations, as described in Section 3.2. Building on the main outcomes of the PICCS analysis, we propose a novel physics-based algorithm, MATCH. MATCH is evaluated via physics-based EM emulations, as it has been described in Section 3.1.3. Finally, we introduce OSIRIS, a unified strategy for codebook generation, that includes the extracted by MATCH entries and minimize the required computation time for online usage.

4.1 Physics-informed Codebook Compilation Software

As mentioned previously, the determination of the RIS configuration that yields a prescribed macroscopic response in a given topology is a complex optimization problem. Three factors govern codebook-compilation performance: the initial state from which the procedure starts; the number of optimization variables—when large, the time required becomes insufficient for practical completion; and the optimization method itself. In this context, integrating codebook compilation with state-of-the-art metaheuristic algorithms can enhance its performance.

4.1.1 PICCS Description

In our case, the main purpose of codebook compilation is the definition of RIS configuration for a beam steering case with different network topologies, i.e., locations of transmitter and receiver. Beam steering constitutes the building block for more complex functionalities, such as beam splitting and multi-directional steering [2]. In other words, complex functionalities can be decomposed into a set of simple beam steerings. The process that can merge/multiplex beam steerings into a single configuration that yields a required more complex functionality is discussed in more detail in Chapter 5.

A crucial tool for a first overview of the RIS-assisted signal trajectory in beam steering is geometric optics (GO). GO, known also as ray tracing, is a well-established framework for simulating EM propagation in domains larger than the wavelength. Although less accurate than physics-based EM frameworks and FDTD, it is ideal for producing initial unit-cell values that seed the optimization. In beam steering, GO implements the generalized Snell's law: for a homogeneously configured RIS, the incidence and reflection angles coincide; with a phase gradient along the surface—via proper tuning of the subwavelength cells—the reflection angle changes, enabling controllable anomalous steering of the incident wavefront. For point-to-point links, GO traces ray pathways between each tunable RIS element and the transmitter (Tx) and receiver (Rx) by accounting for their relative distances and angles. It progressively computes the incident and required reflected angles per cell under a decoupled-element assumption, providing the initial estimates for the optimizer. GO further reveals that symmetry in the RIS pattern arises solely from the geometric symmetry of Tx and Rx with respect to the RIS.

The RIS unit design used in the simulations is the one presented in Section 3.2. This design departs from the simple RIS-reflectarray model by adopting an overlapping phase-shifter model (see Fig. 3.1), whereby interconnections between adjacent unit cells yield a dual contribution of each RIS element to the imposed phases. Nevertheless, GO's applicability is independent of the specific phase-shifter model, as it relies only on the relative distances and angles from Tx to Rx via each RIS element.

The same concept applies to any metasurface type, e.g., by placing a virtual phase shifter at the position of each active element (e.g., PIN diode).

PICCS also addresses the second parameter governing the optimization procedure, informed by observations at the physical layer. PICCS exploits correlations between specific groups of tunable elements and the macroscopic RIS response, quantified by the path loss in the $Tx \rightarrow \text{RIS} \rightarrow Rx$ link. Two element groups are expected to correlate strongly with path loss. The first consists of elements located near the receiver and transmitter, especially those closely aligned vertically with them. The second comprises peripheral elements at the RIS edges. These peripheral elements operate in a non-homogeneous environment due to their placement and are influenced by surrounding structures and edge effects; the resulting abrupt variations in surface current distribution drive diffraction and strongly affect scattering. Simulations confirm these correlations, with the peripheral elements exerting the dominant impact. Accordingly, PICCS uses GO to compute the central-element values (where GO is acceptably accurate) and to initialize the peripheral values. During optimization, only the peripheral values are updated, within a range also determined by the GO analysis. This focused strategy streamlines the procedure, enabling targeted adjustments and significant compilation speed-up.

Regarding the optimization method, PICCS can accommodate any communication metric as the objective. In the present analysis, we use the transmission coefficient, as discussed in Section 3.2. Given the lack of benchmarking on which metaheuristic is best suited to codebook compilation, we present comparative results for the most promising candidates—methods already used in RIS configuration optimization and beamforming, described in Section 3.5—to identify the most suitable approach for integration with PICCS.

4.1.2 Evaluation of PICCS

In this section, we present the results of the PICCS evaluation, conducted in two different setups. The first setup is used to construct PICCS and determine which optimization tool is best suited to the physics-based codebook compilation; the second verifies the proposed methodology. The RIS unit is shown in Fig.4.1, with the tunable elements participating in the optimization highlighted in red. The RIS has 6 unit cells per column and 12 cells per row, resulting in a total of 66 tunable loads; these

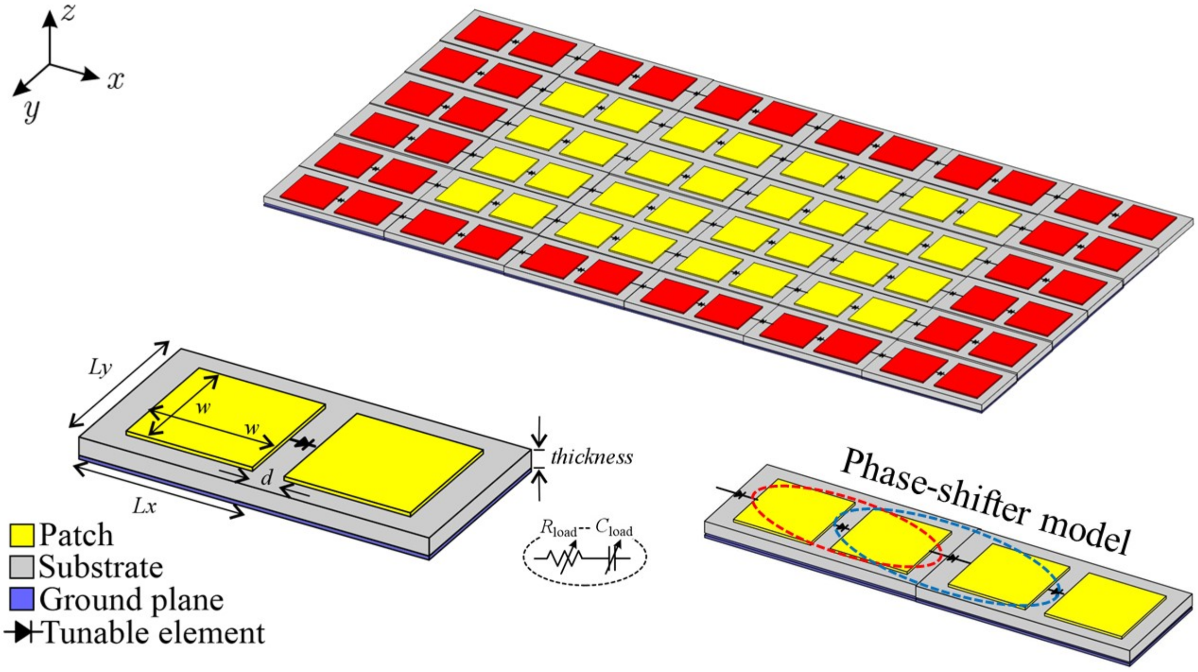

Figure 4.1: Schematic of the RIS used in this work, where the horizontally adjacent metal patches are interconnected with tunable lumped loads. The peripheral phase-shifters (red) and the central ones (yellow). Despite the lack of strict periodicity per element, virtual phase shifters can be defined to expedite GO (bottom right).

dimensions trade off computation time and resources against a sufficiently large RIS implementation. As evident from Fig.4.1, of the total 66 elements, only 20 participate in the optimization, reducing the optimization parameters during PICCS usage by about 70%.

The evaluation setups differ primarily by the Tx–Rx placement. In the first setup (denoted as Setup A), Tx and Rx are symmetrically located with respect to the RIS center; Setup B deliberately breaks this symmetry to assess PICCS under non-trivial conditions where no further reduction of optimization variables is feasible. In both setups, the antennas lie in the near field of the RIS, intentionally stressing any simplifying far-field assumptions. In practice, accurate Tx/Rx positions can be obtained via standard wireless localization techniques [119].

The results for both setups are shown in Table 4.1. For Setup A, we compare five scenarios: (i) $|S_{21}|$ measured between Tx and Rx without the RIS, i.e., along the natural path, serving as the benchmark; (ii) the RIS operated as a static reflectarray; (iii) the RIS configured via GO without any optimization; (iv) an optimization that ignores physics-derived knowledge, labeled NPAO (no-physical-aware optimization);

Table 4.1: Transmission between Tx/Rx for different scenarios conducted by full-wave EM simulations.

	Scenario	**Transmission Coefficient (dB)**	**Number of simulation runs**
Setup A	Antennas	-25.80	1
	Reflectarray	-21.69	1
	GO	-17.40	1
	NPAO	-15.64	1,193
	PICCS	-15.47	25
Setup B	Antennas	-23.07	1
	Reflectarray	-22.88	1
	GO	-19.03	1
	PICCS	-15.32	25

and (v) PICCS. The comparison between NPAO and PICCS is dual, considering both the $|S_{21}|$ behavior and the number of simulation runs. For the second setup, we compare four scenarios, excluding NPAO one.

For Setup A, in the absence of the RIS, the transmission coefficient $|S_{21}|$ is -25.8dB, reflecting the LoS efficiency between the two antennas. Introducing the RIS as a reflectarray, without tunable elements, improves the link by 4.11 dB to -21.69 dB due to passive reflection from the metallic patches. Configuring the RIS via GO yields a further 4.29 dB gain, reaching -17.4 dB, owing to the GO-imposed phase profile that steers the reflected wave toward the receiver. With NPAO, $|S_{21}|$ improves by another 1.76 dB to -15.64 dB, at the cost of $1{,}193$ simulation runs. Finally, PICCS attains -15.47 dB (i.e., $+1.93$ dB over GO) with only 25 simulation runs—about 2% of the NPAO total runs.

The large dataset results obtained from the NPAO using simulations can be processed to verify the correlation between the groups of tunable RIS elements and the overall response. The correlation is quantified between the tunable element values and the resulting $|S_{21}|$ value, as acquired in each iteration along the optimization process. As reminder, the correlation is normalized in the $[0,1]$ range, with the two extremities denoting no impact and direct impact, respectively. As expected, all tunable elements show some correlation with $|S_{21}|$. The average correlation value is 0.24, with a minimum of 0.02 and a maximum of 0.58. For the peripheral elements, the mean correlation is 0.25, while elements closer to normal incidence/reflection with respect to the antennas have a mean correlation of 0.20. The peripheral elements that are also in close proximity to the antennas exhibit the strongest correlation, with a mean value of 0.49.

In order to benchmark which of the optimization tools are more suitable with the codebook compilation process, we compare the performance of PICCS when integrated with each of the optimization tools described in Section 3.5. The results are shown in Fig. 4.2 for Setup A. The boxplots summarize the $|S_{21}|$ values obtained after 25 simulation runs for each optimization tool, both with and without PICCS. The red line indicates the $|S_{21}|$ value obtained from GO, serving as a reference point. Mainly, the integration of PICCS to all optimization methods yields improved performance as compared to optimization with the same methods and same iteration count but without using physical insights. Specifically, PICCS with all methods exhibit an enhancement in $|S_{21}|$ of 2.5 dB. Another important observation from Fig. 4.2 is that employing PICCS leads to a reduced variance around the optimal solution. This ensures that the optimal solution lies within a narrower variable range, allowing the optimization to converge after only a few adjustments. Finally, among the various algorithms tested, NSGA-II emerged as the most effective in consistently finding the optimal solution.

Having demonstrated that the proposed PICCS leads to improved performance with minimized computational resources in a setup with four-fold symmetry, we move to the more complicated Setup B: This setup has only two-fold symmetry, so the optimization solution space is doubled in relation to Setup A. The PICCS is integrated only with the NSGA-II that was previously identified as the most promising optimization method. Examining the results for Setup B in Table 4.1, we observe that with only 25 measurements using PICCS we achieve an enhancement of over 3.6 dB, as compared to the computation solely from GO. The overall enhancement from using PICCS amounts to 7.65 dB.

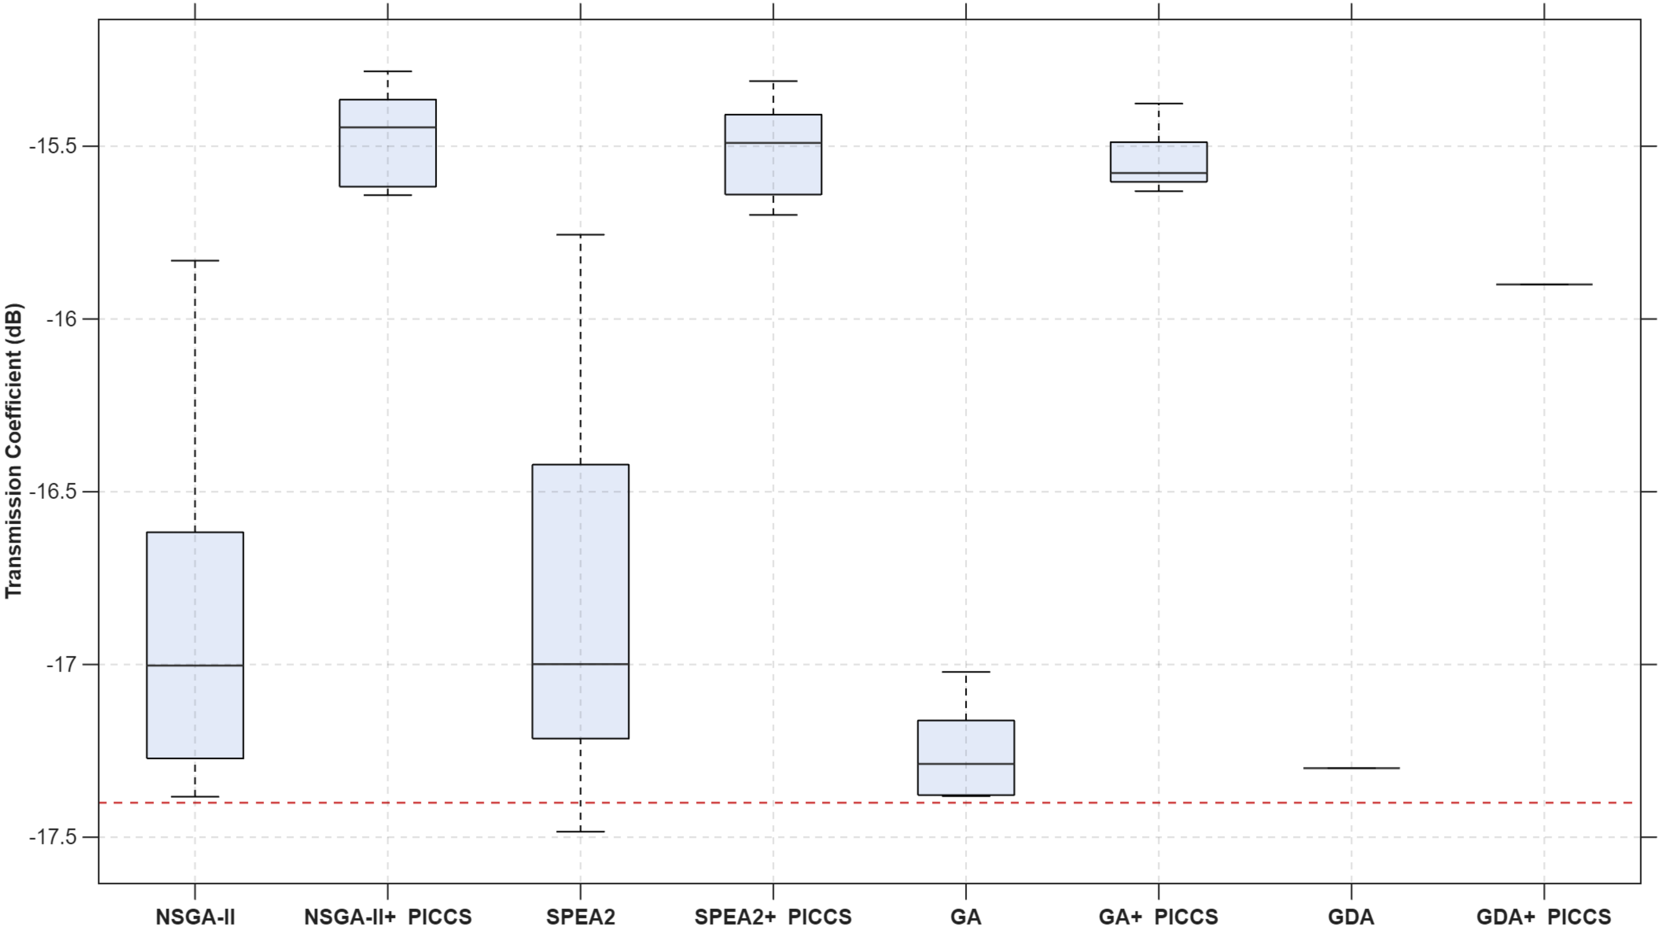

Figure 4.2: Results of benchmarking among optimization methods with and without PICCS usage.

4.1.3 Discussion

The evaluation results demonstrate that incorporating physical-layer insights can substantially accelerate RIS codebook compilation. PICCS achieves this by combining (i) geometric-optics (GO) initialization and (ii) physics-driven dimensionality reduction with (iii) established metaheuristic algorithms. In particular, PICCS exploits the empirically validated correlation between tunable element states and the end-to-end response: peripheral elements—especially those also close to the Tx/Rx alignment—exert a disproportionate impact on the transmission coefficient. Accordingly, PICCS uses GO to produce a strong initial configuration and then restricts the online optimization to a small subset of peripheral variables within physics-informed bounds, which is especially beneficial for finite-sized RISs exhibiting non-periodic phase profiles and edge-driven diffraction effects.

Across the considered setups, this strategy delivers near-NPAO performance with orders-of-magnitude fewer simulation runs, yielding an overall compilation-time reduction of about 98% relative to physics-agnostic baselines. Benchmarking further shows that the choice of metaheuristic remains consequential even under physics-informed pruning: among the tested methods, NSGA-II most consistently converges to the best configuration, while other candidates (e.g., SPEA2) also benefit from PICCS but exhibit less reliable optimality. These findings suggest two actionable directions: (1) physics-aware initialization and variable selection should be treated as first-class components of codebook compilation pipelines, and (2) RIS-tailored metaheuristics—designed to exploit the problem structure revealed by EM behavior—are a credible route to further reducing compilation time without sacrificing link performance.

4.2 MATCH: A Novel Codebook Compilation Algorithm

Continuing the exploration of physics-based codebook compilation, we present MATCH, a novel algorithm for RIS codebook compilation. Unlike PICCS, MATCH couples physics-informed insights with a staged synergy of distinct optimization approaches across its phases. It does not assume a priori which RIS element groups exhibit high correlation with the macroscopic response; instead, it estimates these correlations in an initial phase and then leverages them to structure the main optimization. Finally, MATCH adopts a holistic objective: beyond strengthening the field within the user

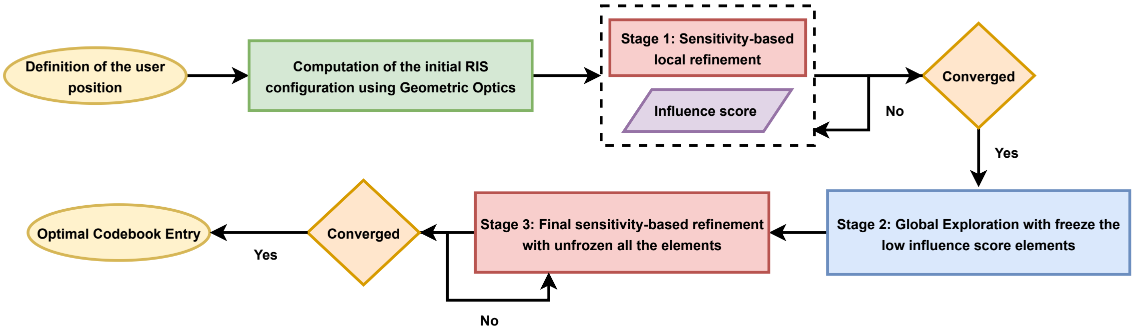

Figure 4.3: Workflow of MATCH algorithm.

region, it concurrently suppresses energy elsewhere so that power is redirected from undesired directions toward the target. MATCH workflow is illustrated in Fig. 4.3.

4.2.1 MATCH Description

In this section, we present MATCH algorithm. The step-by-step description of MATCH is included in Alg. A.1 in Appendix A. MATCH is a physics-driven algorithm for near-field beam focusing that explicitly incorporates mutual coupling among RIS elements and specular reflections within the environment. As outlined in Fig. 4.3, MATCH follows a multi-stage refinement process that not only derives the optimal reflection coefficients $\Gamma = (\Gamma_1, \dots, \Gamma_N) \in \mathbb{C}^N$for a given topology but also leverages the physics knowledge gathered during execution to improve convergence and performance. Furthermore, it formulates beam focusing as a holistic EM energy concentration problem—rather than merely maximizing energy within the focus area, it actively redirects power from undesired directions toward the region of interest.

To formulate the design objective of MATCH, the desired EM behavior is defined over specific regions of the environment where field enhancement or suppression is required. The receiver area corresponds to a focus region $\mathcal{F}$ centered at one or more points $\{\mathbf{c}^{(\ell)}\}\ell = 1^L$, each modeled as a small sphere of radius r that represents the spatial extent within which a strong and uniform field distribution is desired. To complement this focusing objective, a second region $\mathcal{O}$ is defined to represent the remaining part of the environment, excluding all focus spheres, where the field intensity should be minimized. However, since the EM field distribution is inherently coupled across space, the concentration of energy within $\mathcal{F}$ cannot be achieved independently of how the field behaves elsewhere. Consequently, an exclusive maximization of the field inside the target region may not effectively redirect the scattered energy toward it, as undesired constructive interference outside $\mathcal{F}$ could still dominate the overall response. A formulation that simultaneously enhances the field within $\mathcal{F}$ and attenuates it across $\mathcal{O}$ therefore ensures that the optimization acts coherently on the global energy balance of the environment, steering the configuration toward physically consistent focusing behavior.

To evaluate these objectives efficiently, the entire space is discretized into observation points at which $E(\mathbf{r})$ is computed, with a total of $N_{\mathcal{F}}$ samples distributed within the focus regions and $N_{\mathcal{O}}$ samples placed throughout the outer region. The sampling density is intentionally increased inside $\mathcal{F}$ to capture fine spatial variations of the field, whereas a coarser grid is used in $\mathcal{O}$ to maintain computational tractability. Based on this discretization, the first design goal of MATCH is to maximize the mean field magnitude $\bar{E}_{\mathcal{F}}$ within the focus regions, expressed as

$$\max_{\Gamma} \bar{E}_{\mathcal{F}} = \max_{\Gamma} \sum_{j=1}^{N_{\mathcal{F}}} w_j \left|E(\mathbf{r}_j)\right|, \tag{4.1}$$

where $w_j = 1/N_{\mathcal{F}}$ is a weighting factor introduced for the spatial averaging of the field over the focus region under uniform sampling, while the second goal is to minimize the corresponding mean field $\bar{E}_{\mathcal{O}}$ magnitude across the outer region, given by

$$\min_{\Gamma} \bar{E}_{\mathcal{O}} = \min_{\Gamma} \sum_{j=1}^{N_{\mathcal{O}}} w\text{'}_j \left|E(\mathbf{r}\text{'}_j)\right|, \tag{4.2}$$

where $w\text{'}_j = 1/N_{\mathcal{O}}$ denotes the analogous weighting coefficient associated with each sampling point in the outer region. Therefore, through the joint optimization of $\bar{E}_{\mathcal{F}}$ and $\bar{E}_{\mathcal{O}}$, MATCH establishes a physically consistent trade-off that achieves strong field focusing at the receiver while suppressing unwanted radiation throughout the rest of the environment.

To ensure that the optimization starts from a physically meaningful state, as in the PICCS, MATCH employs an initialization inspired by GO, where each RIS element is assigned a reflection phase that promotes constructive interference at the desired focus regions (Initialization in Alg. A.1). In this preliminary stage, secondary reflections and mutual coupling effects are disregarded, allowing the phase of each element to be determined solely by its propagation distance to the center of the focus region.

Following the GO initialization, MATCH transitions to a sensitivity-based local refinement stage (Stage 1) that builds directly upon the physics of the full propagation model in (3.20). At this stage, mutual coupling among unit cells and secondary wall reflections are incorporated to account for the complete EM interactions that characterize the RIS-assisted channel. To consistently integrate these effects within the optimization process, each RIS element undergoes small iterative phase perturbations of magnitude $\delta\phi$, and after each adjustment, the total field $E(\mathbf{r})$ is recomputed across the sampling sets to evaluate the corresponding variation in the focusing objective $\bar{E}_{\mathcal{F}}$. By continuously comparing successive field realizations, the procedure selectively retains phase updates that enhance the objective while reverting those that degrade it, thereby steering the configuration toward a physically consistent local optimum. Through this mechanism, the refinement process effectively performs a gradient-driven adaptation of the RIS phases, following the established principles of beamforming optimization, suitably extended to the near-field and strongly coupled regime of indoor RIS environments [120, 121]. The iterations proceed until the relative improvement in $\bar{E}_{\mathcal{F}}$ falls below a local tolerance $\varepsilon_{\mathrm{loc}}$, indicating convergence to a locally stable configuration. In parallel, the same perturbation–response mechanism quantifies the contribution of each RIS element to the overall focusing performance through its sensitivity

$$C_n = \frac{\partial \bar{E}_{\mathcal{F}}}{\partial \phi_n}, \tag{4.3}$$

which can be obtained analytically from the field derivative

$$\frac{\partial E(\mathbf{r})}{\partial \phi_n} = j\,\Gamma_n E_{\mathrm{inc},n}\, G(\mathbf{r}, \mathbf{p}_n), \tag{4.4}$$

where $G(\mathbf{r}, \mathbf{p}_n)$ denotes the Green's function linking the n-th RIS element with the observation point $\mathbf{r}$. Aggregating this response over all focus samples yields

$$\frac{\partial |E|}{\partial \phi_n} = \mathrm{Re}\left\{ \frac{E^*(\mathbf{r})}{|E(\mathbf{r})|}\, \frac{\partial E(\mathbf{r})}{\partial \phi_n} \right\}, \tag{4.5}$$

providing a closed-form relation between local phase adjustments and the global field strength inside $\mathcal{F}$. For each sampling point $\mathbf{r}_j$, the vector of element-wise sensitivities $\mathbf{c}_j$ is defined with entries

$$[\mathbf{c}_j]_n = \mathrm{Re}\left\{ \frac{E^*(\mathbf{r}_j)}{|E(\mathbf{r}_j)|}\, j\,\Gamma_n E_{\mathrm{inc},n} G(\mathbf{r}_j, \mathbf{p}_n) \right\}, \tag{4.6}$$

and the ensemble of these local responses is then aggregated to form the overall correlation matrix

$$\mathbf{H} = \sum_{j=1}^{N_{\mathcal{F}}} \mathbf{c}_j \mathbf{c}_j^{\top}, \tag{4.7}$$

which captures the joint influence of all RIS elements on the mean field within the focus region. The diagonal entries of $\mathbf{H}$ quantify the individual contribution of each element to the focusing performance, whereas the off-diagonal terms characterize inter-element coupling due to mutual interactions and shared propagation paths. Consequently, $\mathbf{H}$ provides a physically grounded foundation for identifying low-influence unit cells that can be frozen to improve efficiency and for defining high-impact unit cells that can guide a global optimization phase.

After reaching local convergence, MATCH naturally transitions to a global exploration phase (Stage 2) that extends the optimization beyond the neighborhood shaped by the sensitivity-based refinement. Building upon $C_n^{(t)}$, this stage seeks to reconcile the two physically intertwined objectives of enhancing the field intensity within the focus region $\bar{E}_{\mathcal{F}}$ while simultaneously suppressing the residual radiation across the outer region $\bar{E}_{\mathcal{O}}$. However, since constructive interference and leakage suppression are inherently coupled EM phenomena, their joint control requires a Pareto-based optimization framework that captures the fundamental trade-off between them. Within this setting, the sensitivity data guide the search process by identifying elements with limited impact that can be temporarily frozen, thereby constraining the exploration to the most influential degrees of freedom without compromising physical accuracy. This targeted reduction of the search space enables the algorithm to concentrate computational resources on the most influential unit cells and efficiently traverse the Pareto landscape, where each candidate configuration is evaluated through the pair of metrics $(\bar{E}_{\mathcal{F}}, \bar{E}_{\mathcal{O}})$.Finally, to balance convergence efficiency and solution diversity, MATCH employs multi-objective metaheuristics, with NSGA-II—identified during PICCS as the optimal choice—proving particularly effective for preserving Pareto-front diversity and achieving rapid convergence in physics-driven settings. Once a trade-off between focusing gain and leakage suppression is reached, a final sensitivity-based refinement (Stage 3) applies gradient ascent to fine-tune all RIS phases under the full propagation model, and the resulting configuration $\Gamma^{\star}$ is stored in the codebook, completing one MATCH cycle.

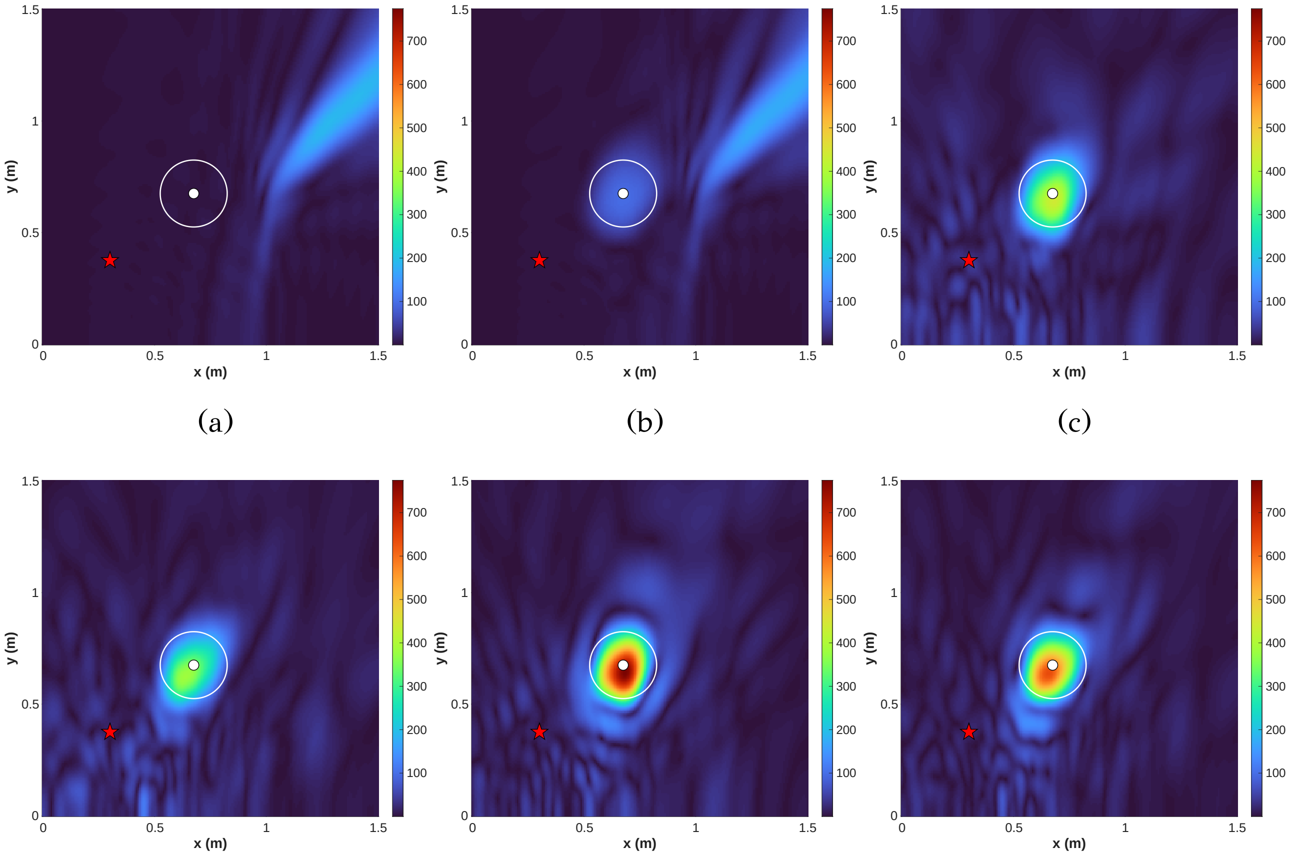

(a) (b) (c)

(d) (e) (f)

Figure 4.4: Radiation patterns for near-field beam focusing: (a) GO initialization, (b) Stage 1, (c) Stage 2 with and (d) without the minimization objective, and (e) Stage 3 with and (f) without the minimization objective.

4.2.2 Evaluation of MATCH

In this section, we assess the performance of MATCH using emulations, as described in Section 3.1.3, under two cases: a single-RIS and a full-coverage where all room walls are covered with RIS units. The corresponding physical and algorithmic parameters used throughout all experiments are summarized in Table 4.2. The initialization of the phase perturbation step is set to $\delta\phi = 0.2$ with a local convergence tolerance of $\varepsilon_{\text{loc}} = 10^{-2}$, ensuring a balanced trade-off between exploration and stability during the early optimization stages, while in the subsequent refinement iterations these values are reduced to $\delta\phi = 0.1$ and $\varepsilon_{\text{fin}} = 10^{-5}$ to enable smooth convergence toward the final optimum without oscillatory behavior. Moreover, each optimization run is restricted to a maximum of 10^4 iterations, providing a consistent balance between computational efficiency and accuracy across all configurations. Furthermore, the perturbation–response information obtained during the sensitivity analysis is leveraged to gradually focus the optimization on the most relevant degrees of freedom by freezing the least-contributing 5% of RIS elements every 25% of the total generations, thereby allocating computational effort to the most influential unit cells, enabling faster convergence and improved field control. Finally, after convergence, the resulting field distributions are evaluated on a horizontal plane intersecting the receiver level, and all radiation patterns are normalized with respect to an incident field magnitude of $E_{\text{n}} = 1$ V/m to enable consistent visual and quantitative comparisons across different setups.

Fig. 4.4 illustrates the evolution of the near-field beam focusing process throughout the stages of the MATCH algorithm for the single-RIS configuration. The transmitter radiates signal targeting the RIS unit. The sequence begins in Fig. 4.4a, where the GO initialization establishes a phase distribution that directs the reflected energy toward the intended region, forming a coherent baseline for the subsequent refinement. However, as this stage does not account for mutual coupling and specular reflections, the resulting field still exhibits significant leakage along the main beam, and the energy concentration within the focal region remains limited. When these EM interactions are incorporated through the sensitivity-based local refinement shown in Fig. 4.4b, the reflected wavefront becomes progressively reorganized within the space, with the RIS adapting its response in a manner consistent with mutual coupling and specular reflections. As a result, the energy density within the focus

Table 4.2: EMULATION PARAMETERS

Parameter	**Value**
Room dimensions	$L_{\text{th}} = 1.5$ m
RIS dimensions	120×120 unit cells
Inter-element spacing	$d = \lambda/4$
Operating frequency	$f = 6$ GHz
Focus region radius	$r = 0.15$ m
Focus sampling points	$N_{\mathcal{F}} = 12{,}500$
Outer sampling points	$N_{\mathcal{O}} = 15{,}000$
Mutual coupling coefficient	$\alpha = 0.15$
Field normalization reference	$E_{\text{n}} = 1$ V/m
Population size	300
Generations	75
Crossover / mutation indices	8 / 8
Mutation probability	25%

increases by 22.9 dB, while the off-axis energy decreases, though residual sidelobes remain due to the optimization acting only on the inner region. This imbalance naturally leads to the global exploration stage illustrated in Fig. 4.4c, where MATCH simultaneously maximizes the field within the focus area and minimizes it outside of the focus area, achieving an additional 11.3 dB improvement and a noticeably sharper confinement of energy around the receiver. However, Fig. 4.4d highlights that when this outer-field minimization is omitted, the attainable improvement drops to 7.99 dB, confirming that suppressing leakage indirectly reinforces the focal intensity. Finally, as the optimization converges, the configuration shown in Fig. 4.4e fine-tunes the global solution through the final sensitivity-based refinement, adding another 3.5 dB gain and stabilizing the beam without reintroducing sidelobes. In contrast, when the same refinement follows the single-objective formulation, as in Fig. 4.4f, the gain reaches only 2.5 dB, showing that the coupled optimization is essential for accurate near-field beam focusing.

Table 4.3 complements these observations by quantifying the redistribution of the total scattered energy across the focused, directed, and unexploited components at each stage. Starting from the geometric-optics initialization, where only 0.1% of the total energy lies within the focal region, the sensitivity-based refinement already redirects 12.6% toward the target and reduces the leakage fraction from 93.1% to 77.0%,

Table 4.3: Energy fractions by region for each stage of MATCH w/ & w/o the minimization objective in the single-RIS case

Stage	η_{focus}	η_{dirOut}	η_{unexp}
GO (w & w/o)	0.1%	6.8%	93.1%
Stage 1 (w & w/o)	12.62%	10.36%	77.02%
Stage 2 (w)	75.5%	16.36%	8.14%
Stage 2 (w/o)	71.75%	18.19%	10.06%
Stage 3 (w)	85.66%	11.62%	2.72%
Stage 3 (w/o)	80.5%	14.56%	4.94%

indicating that coupling and reflections begin to cooperate constructively. Building on this, the global exploration further aligns the scattered energy, raising the focused fraction to 75.5% while lowering the unexploited energy below 10%. The final refinement consolidates this effect, achieving 85.7% of the total energy within the focus and reducing the remaining directed and unexploited components to 11.6% and 2.7%, respectively. Compared to the single-objective case, which concentrates only 80.5% and exhibits higher leakage, these results show that effective near-field beam focusing arises not merely from constructive interference but from the coordinated adjustment of all interactions within the environment.

Subsequently, Fig. 4.5 presents the evolution of the near-field beam focusing process for the full-RIS case and in Table 4.4 the redistribution of the EM energy at each stage is reported. In this scenario, the transmitter radiates omnidirectionally, making the propagation control within the room considerably more challenging due to the increased complexity of multipath interactions and reflections. The GO initialization (Fig.4.5a) forms an initial phase distribution that roughly directs the field toward the focus but leaves strong leakage and limited energy concentration (2.73% of total energy within the focus). Incorporating EM interactions via sensitivity-based refinement (Fig.4.5b) reorganizes the wavefront, increasing the focal intensity by 4.97 dB and raising the focused energy to 7.31%, while leakage drops from 64.63% to 57.23%. The global exploration of MATCH (Fig.4.5c) jointly maximizes the in-focus field and suppresses out-of-focus energy, yielding an additional 4.1 dB gain and 18.87% energy confinement within the focus, with unexploited energy falling below 43.3%.

When outer-field minimization is omitted (Fig.4.5d), there is no gain within the

Table 4.4: Energy fractions by region for each stage of MATCH w/ & w/o the minimization objective in full-RISs case

Stage	η_{focus}	η_{dirOut}	η_{unexp}
GO (w & w/o)	2.73%	32.65%	64.63%
Stage 1 (w & w/o)	7.31%	35.46%	57.23%
Stage 2 (w)	18.87%	38.8%	42.33%
Stage 2 (w/o)	7.31%	35.53%	57.16%
Stage 3 (w)	26.6%	36.75%	33.65%
Stage 3 (w/o)	16.39%	41.51%	42.1%

focus are, confirming once more the necessity of leakage suppression, especially in environments with rich multipath effects. The final refinement (Fig.4.5e) adds 1.9 dB, concentrating 26.6% of total energy within the focus while limiting directed and unexploited fractions to 36.75% and 33.65%, respectively. In contrast, the single-objective variant (Fig. 4.5f) attains 3.59 dB not succeeding though to achieve focus efficiency greater than 16.39%. Overall, MATCH achieves superior near-field focusing by coupling EM-aware optimization with global leakage suppression, revealing how mutual coupling and specular reflections jointly govern spatial energy confinement in RIS-assisted environments.

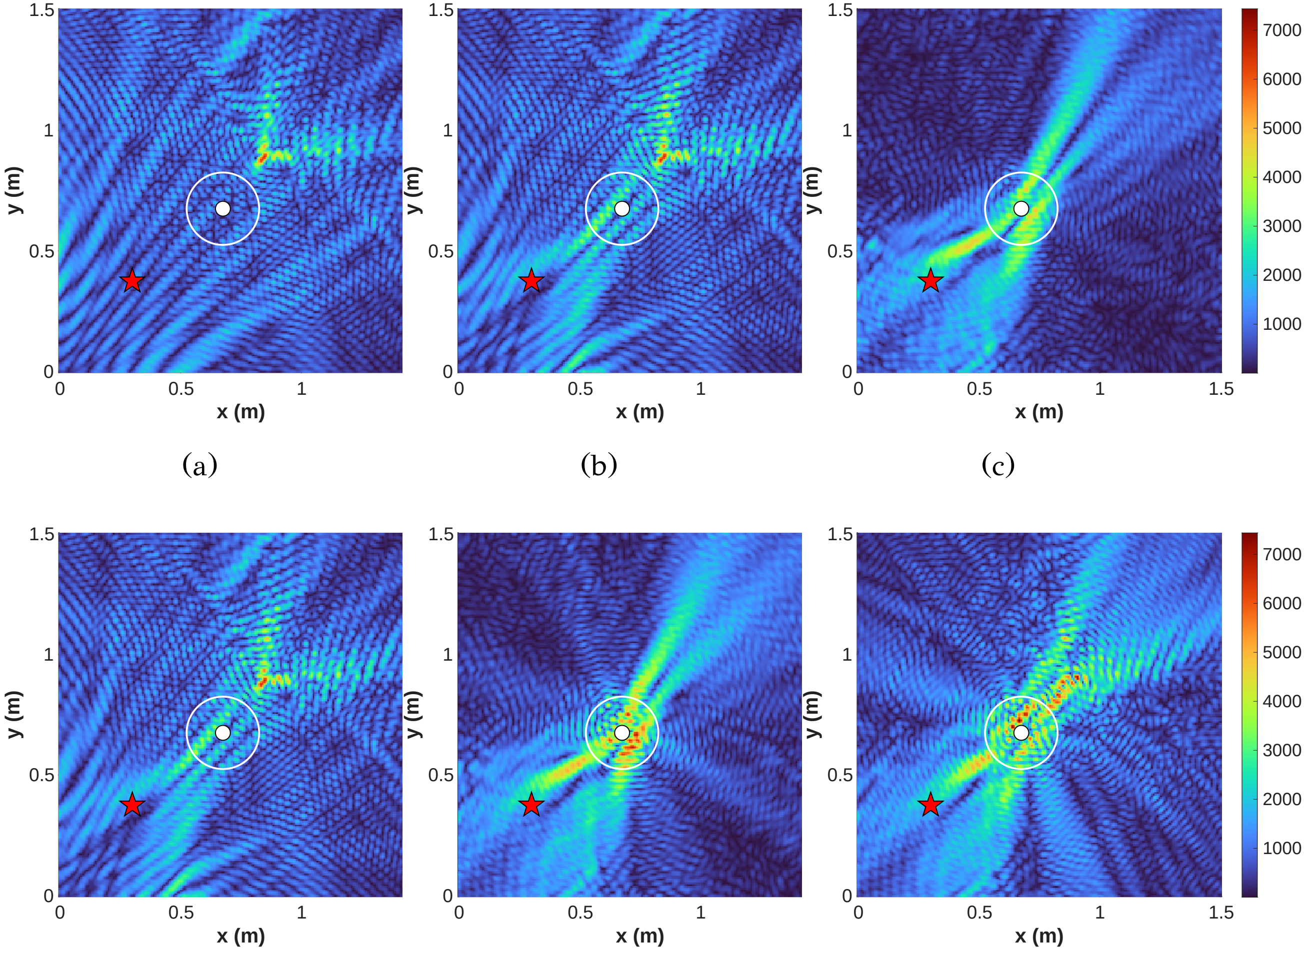

(a) (b) (c)

(d) (e) (f)

Figure 4.5: Radiation patterns for near-field beam focusing: (a) GO initialization, (b) Stage 1, (c) Stage 2 with and (d) without the minimization objective, and (e) Stage 3 with and (f) without the minimization objective.

4.2.3 Discussion

The proposed MATCH algorithm extends the physics-based codebook-compilation direction beyond PICCS by coupling physics extraction with a staged synergy of optimization tools. MATCH begins with a GO-based initialization that provides a coherent phase baseline, then performs a sensitivity-driven local refinement under the full propagation model (including mutual coupling and specular reflections), and subsequently transitions to a guided global exploration that concentrates the search on the most influential degrees of freedom. A final sensitivity-based refinement completes the cycle by consolidating convergence and stabilizing the near-field focus.

Emulation results verify that explicitly accounting for coupling and secondary reflections is not a secondary detail but a primary enabler of controllable near-field focusing. Across both the single-RIS and full-coverage scenarios, MATCH progressively redistributes the scattered energy toward the focal region, with the joint objective—simultaneous enhancement inside $\mathcal{F}$ and suppression across $\mathcal{O}$—consistently reducing leakage and strengthening the achieved focus compared to single-objective variants. This confirms the core design premise of MATCH: effective focusing in strongly coupled indoor environments requires global field control rather than maximization within the target region alone.

Overall, MATCH demonstrates that codebook compilation can be made both efficient and physically interpretable when the optimization is driven by on-the-fly physics knowledge (element sensitivities/correlations) and when leakage suppression is treated as a first-class objective. The resulting configurations achieve higher energy confinement with fewer wasted degrees of freedom, supporting the broader thesis claim that physics-consistent, codebook-based pipelines are a practical route to low-latency RIS operation in PWEs while providing clear insight into near-field energy redistribution mechanisms.

4.3 OSIRIS: A Two Stage Strategy for Online Codebook Generation

MATCH is a useful tool for codebook compilation. A key question, however, is whether MATCH must run at every point in the RIS-assisted network to attain the per-user optimum; if so, the load is prohibitive. We therefore propose OSIRIS to enable low-time-complexity selection from the codebook. OSIRIS follows a two-phase design: (i) offline, physics-consistent configurations are generated via MATCH on a sparse 3D sampling of the user space; (ii) these configurations are compressed into a compact, searchable codebook via phase-domain clustering and spatial anchoring. By explicitly structuring the codebook and restricting online evaluation to a small, position-relevant subset, OSIRIS achieves efficient configuration selection with bounded, predictable complexity. The main logic of OSIRIS is illustrated in Fig. 4.6.

4.3.1 OSIRIS Description

This section presents in details the flow of proposed strategy, OSIRIS. The step-by-step procedure is presented in Alg. B.1 that is included in Appendix B. As mentioned before, OSIRIS is built around a two-phase structure that reflects the fundamental separation between physics-intensive processing and real-time decision making. In particular, the generation of physics-consistent RIS configurations is confined to an offline phase using MATCH, while their structured organization and selection are handled during operation, enabling low-complexity configuration without sacrificing accuracy.

In the offline phase, OSIRIS operates over the admissible user region defined by deployment geometry and service requirements, which naturally constrain user locations within bounded spatial limits. In practice, these constraints restrict the vertical extent of the user space, yielding a finite range of admissible user heights and delimiting the region over which RIS focusing must be supported. To capture the dominant spatial variation of optimal RIS responses within this bounded region, spatial sampling is required, which is achieved by discretizing each admissible height level using a regular grid whose resolution balances representational fidelity against offline computational cost. Under this discretization, MATCH is executed to compute a unit modulus RIS phase configuration $\boldsymbol{\Gamma}_n^\star \in \mathbb{C}^N$, with $|\Gamma_{n,i}^\star| = 1$ for all elements i, at each sampled location. Collectively, these configurations form a training database

$$\mathcal{D} = \{(\mathbf{x}_n, \boldsymbol{\Gamma}_n^\star)\}_{n=1}^{M}, \tag{4.8}$$

which constitutes a maximal physics consistent codebook determined by the chosen spatial sampling. However, while $\mathcal{D}$ enables accurate RIS operation at the sampled locations, it also exposes a fundamental limitation of finite codebooks, since an optimal configuration cannot be stored for every possible user position in a continuous

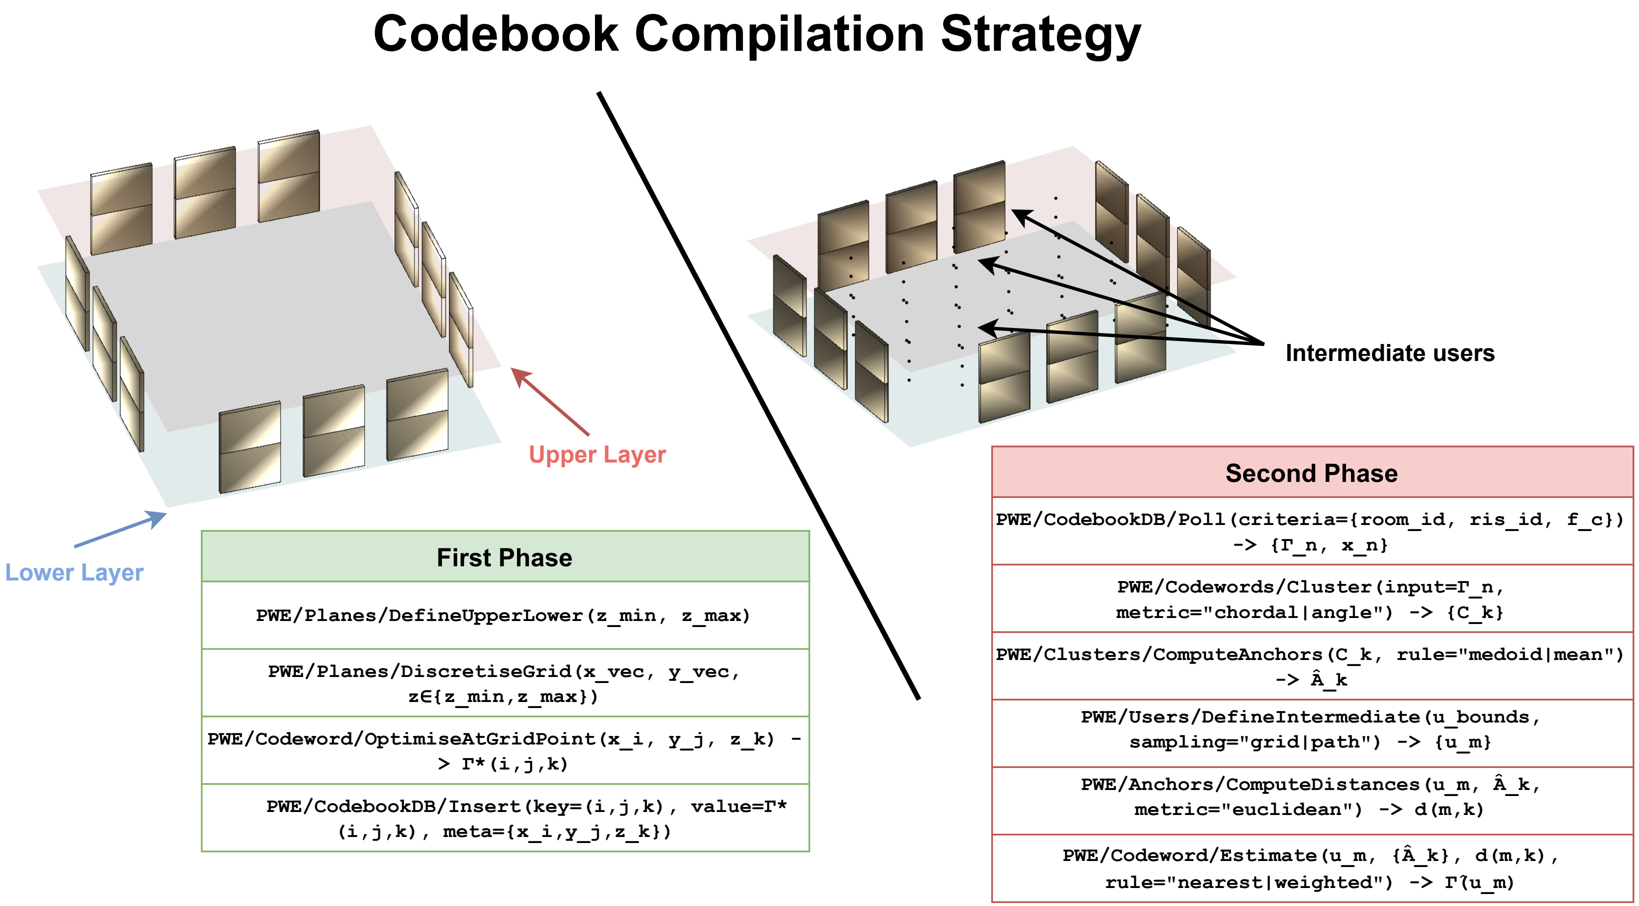

First Phase
PWE/Planes/DefineUpperLower(z_min, z_max)
PWE/Planes/DiscretiseGrid(x_vec, y_vec, z∈{z_min,z_max})
PWE/Codeword/OptimiseAtGridPoint(x_i, y_j, z_k) -> Γ*(i,j,k)
PWE/CodebookDB/Insert(key=(i,j,k), value=Γ*(i,j,k), meta={x_i,y_j,z_k})

Second Phase
PWE/CodebookDB/Poll(criteria={room_id, ris_id, f_c}) -> {Γ_n, x_n}
PWE/Codewords/Cluster(input=Γ_n, metric="chordal|angle") -> {C_k}
PWE/Clusters/ComputeAnchors(C_k, rule="medoid|mean") -> Â_k
PWE/Users/DefineIntermediate(u_bounds, sampling="grid|path") -> {u_m}
PWE/Anchors/ComputeDistances(u_m, Â_k, metric="euclidean") -> d(m,k)
PWE/Codeword/Estimate(u_m, {Â_k}, d(m,k), rule="nearest|weighted") -> Γ̃(u_m)

Figure 4.6: Workflow of the proposed codebook compilation strategy within PWE. In the first phase (a), the upper and lower user planes are defined and grid points are computed via MATCH. In the second phase (b), the grid is queried online and the retrieved configurations are post-processed to estimate near-optimal codebook entries on the fly.

space, and increasing the sampling density to improve spatial coverage directly enlarges the codebook and renders high dimensional processing impractical. As a result, scalability requires exploiting structure within the configurations already contained in the database. In this direction, the stored configurations associated with different spatial locations often induce similar electromagnetic responses. To capture this redundancy in a physically meaningful manner, configurations are represented in the phase domain and mapped to the unit phasor space, where distance reflects similarity in the induced response, revealing an intrinsic low dimensional structure that enables compression of $\mathcal{D}$ into a finite set of representative phase only codebook entries, each associated with a spatial anchor.

After obtaining the compressed $\mathcal{D}$, OSIRIS performs the selection of a near optimal RIS configuration for a given user position by exploiting the structured representation obtained offline. In more detail, given the user position, OSIRIS retains only those codebook entries whose anchors lie close to that position, while entries associated with distant anchors are excluded from consideration. Thus, by restricting evaluation to this position relevant subset, the online computational effort remains bounded and independent of the total codebook size. Finally, among the retained candidates, OSIRIS selects the configuration that maximizes the received SNR at the user location, yielding the discrete decision

$$\mathbf{\Gamma}_{\mathrm{MAP}}(\mathbf{x}) \in \{\mathbf{C}_1, \dots, \mathbf{C}_K\} . \tag{4.9}$$

This constitutes the final online action for a single location, and the same inference step is then repeated independently for each of the L user positions

4.3.2 Evaluation of OSIRIS

In this section, we evaluate OSIRIS against the optimal computed by MATCH. Since a straightforward approach for intermediate user locations is 3D interpolation, we include it as a baseline. For interpolation, we first find the eight surrounding grid points (four on the lower plane and four on the upper plane) that form the smallest axis-aligned box containing the intermediate user. We then compute the user's relative position inside this box along each axis and use it to assign standard trilinear weights to the eight corners. Moreover, we report the compilation time of OSIRIS versus exhaustive optimization. The experimental setup follows Section 3.1.3.

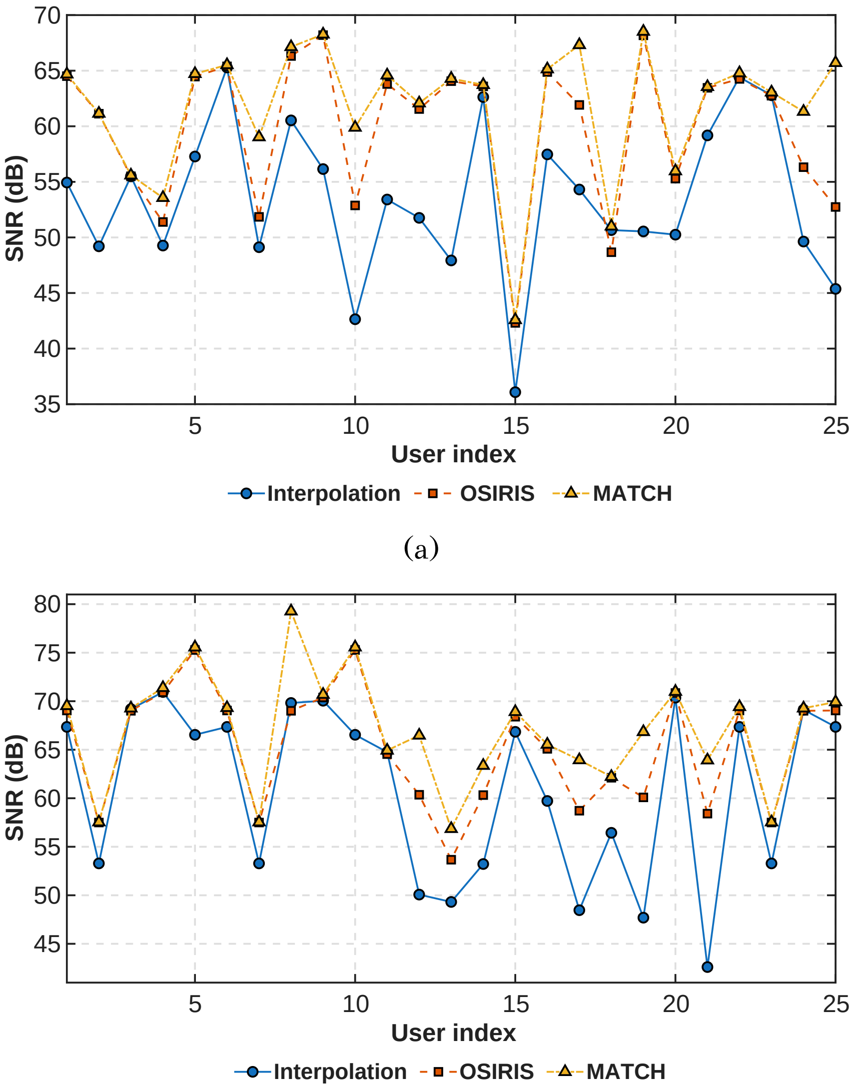

(a)

(b)

Figure 4.7: Comparison of SNR among 3D interpolation, OSIRIS and optimal with grid points $L/6$ (up) and $L/10$ (down).

Table 4.5: SNR loss in dB between 3D interpolation/OSIRIS and optimal case.

Grid	**Method**	**ΔMean**	**ΔMedian**	**ΔMin**	**ΔMax**
$L/6$	3D Interpolation	8.27	7.67	0.35	17.64
$L/6$	OSIRIS	1.91	0.34	0.02	12.97
$L/10$	3D Interpolation	4.37	4.21	0.8	15.81
$L/10$	OSIRIS	1.83	0.28	0.04	10.26

First, we assess the quality of the OSIRIS-generated codebook entries in respect of optimal values, as computed by MATCH, and interpolation. The evaluation is performed under two offline grids: one with spacing $L/6$ along both the $x-$ and $y-$ axes (where L denotes the side length of the cubic RIS-assisted space) and one with spacing $L/10$. For each grid, we set the number of clusters K approximately one-tenth of the number of grid points, i.e., $K = 5$ for the $L/6$ grid and $K = 16$ for the $L/10$ grid, and we fix $K_{\text{top}} = 5$. The lower and upper planes are located at $z_{min} = 0.15\,\text{m}$ and $z_{max} = 1.35\,\text{m}$, respectively. For both cases, we perform a Monte Carlo evaluation over 25 intermediate user locations randomly placed between the two planes. The corresponding results are shown in Fig. 4.7.

For the $L/6$ grid, Fig. 4.7a shows that OSIRIS achieves near-optimal SNR for the majority of users. The mean SNR gap between OSIRIS and the optimal values, as computed by MATCH, is 1.91 dB, with a median of 0.34 dB; the minimum gap is 0.02 dB. The maximum deviation occurs at user index 10, where the gap reaches 12.97 dB; however, the absolute SNR for that user remains above 40 dB, indicating that the achieved performance is still sufficient. In contrast, 3D interpolation exhibits a substantially larger loss: the mean gap from optimal is 8.27 dB with a median of 7.67 dB, and its worst-case deviation is approximately 4.5 dB higher than that of OSIRIS. This stems from interpolation's implicit smooth-phase assumption—although it provides a local, closed-form, low-cost estimate of RIS configurations, it can fail in regions where the EM response changes abruptly such as near walls, corners, or room transitions. These differences are reported in detail in Table 4.5.

For the denser $L/10$ grid, interpolation improves as expected, confirming the benefit of increased offline sampling density: the mean gap from optimal decreases to 4.37 dB and the median to 4.21 dB, while the maximum deviation is reduced to 15.81 dB. OSIRIS also benefits from the denser grid, since it can select among a richer set of candidate patterns; its mean gap becomes 1.83 dB with a median of 0.28 dB. Importantly, the improvement for OSIRIS relative to the $L/6$ case is modest, indicating diminishing returns with grid densification and supporting the key advantage of OSIRIS: it can deliver near-optimal performance, thereby reducing offline codebook-generation time for a fixed network topology.

Finally, we analyze the compilation time of the exhaustive optimization that is performed via MATCH against the OSIRIS inference procedure. The measured times, shown on a logarithmic scale, are reported in Fig. 4.8. For a broader perspective,

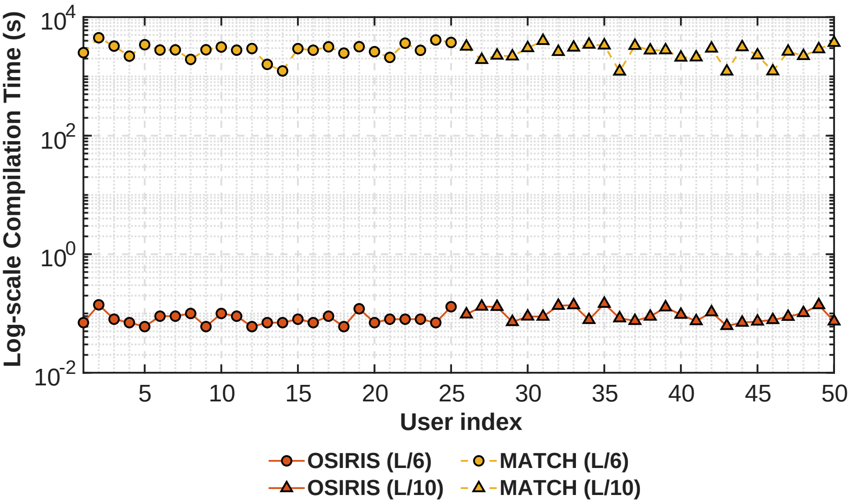

Figure 4.8: Comparison of compilation time in logarithmic scale between OSIRIS and MATCH.

we aggregate results from both the $L/6$ and $L/10$ cases, noting that the offline grid density does not affect the online runtime. The results show a clear trade-off: the controllable performance loss incurred by OSIRIS is exchanged for a drastic reduction in compilation time. In particular, MATCH requires $2615.34\,\mathrm{s}$ on average (median $2751.9\,\mathrm{s}$), whereas OSIRIS requires $0.097\,\mathrm{s}$ on average (median $0.089\,\mathrm{s}$), corresponding to an average runtime reduction of approximately 2.8×10^4. This confirms that OSIRIS can serve as a practical tool for fast codebook entry generation while retaining near-optimal performance.

4.3.3 Discussion

OSIRIS addresses the scalability bottleneck exposed by MATCH by explicitly separating physics-intensive configuration generation from real-time decision making. It first computes physics-consistent configurations offline via sparse MATCH runs over a bounded 3D user region, and then compresses the resulting database through phase-domain clustering and spatial anchoring. This yields a compact, structured codebook in which each representative entry is tied to a spatial anchor, enabling online selection to be performed over a small, position-relevant subset rather than over the full database or via exhaustive re-optimization.

The evaluation confirms that this structure delivers near-optimal performance across a continuous user space without requiring MATCH to run per user location. For intermediate user positions, OSIRIS consistently achieves SNR close to the MATCH optimum, with mean losses below 2 dB and median losses below 1 dB even under coarse spatial sampling, while substantially outperforming 3D interpolation in both average and worst-case behavior. These results support the underlying premise that the space of physics-consistent RIS configurations exhibits exploitable redundancy that can be captured reliably in the phase domain.

Crucially, OSIRIS exchanges a controlled and quantifiable performance loss for a drastic reduction in runtime: the reported inference time is on the order of 10^{-1} s versus 10^{3} s for exhaustive MATCH-based optimization, corresponding to a reduction exceeding four orders of magnitude with bounded, predictable complexity. Overall, OSIRIS demonstrates that structured compression of physics-consistent configurations is a practical route to scalable RIS codebook design, preserving EM fidelity while enabling real-time operation in PWEs.

CHAPTER 5

RESOURCE ALLOCATION IN RIS-ASSISTED PWEs

In this section, we first formalize RIS as a network resource—an aspect missing from prior work. Building on this, we introduce the codebook-entry multiplexing approach for multi-user operation, validated by analysis and the setups of Sec. 3. On this basis, we propose three dedicated algorithms: Codebook Multiplexing Algorithm (COMMON), PRIcing-based Multiplexing of codebook Entries (PRIME), and MuLtiplexing codebook Entries with controllable performance leVELs (LEVEL) for priority-based resource allocation in a multi-user, RIS-assisted PWE. We evaluate and compare the proposed algorithms in terms of user performance, fairness, and computational time for realistic deployment.

5.1 RIS as a Network Resource

In order to define the notion of the RIS resource, we employ the notion of EM functions [121]. As noted, these are macroscopic descriptors of the manipulation exerted by the RIS to the impinging wave. They are expressed in software terms, i.e., as functions with well-defined input and output data.

For instance, the `PLANAR_ABSORB`$(\vec{r}) \rightarrow \emptyset$ function describes the full absorption of a planar wave impinging on the RIS from a direction $\vec{r}$. The `PLANAR_STEER`$(\vec{r}) \rightarrow \vec{d}$ describes the reflection or refraction of a wave impinging from direction $\vec{r}$ to direction $\vec{d}$. The `PLANAR_SPLIT`$(\vec{r}) \rightarrow \left[\{\vec{d_1}, p_1\}, \ldots\right]$ describes the reflection or refraction of the impinging wave into a set of different directions, $\vec{d_i}$, each carrying a portion, p_i, of the power of the original, impinging wave. Notice that function f must have a corresponding set of RIS element states, σ_{ij}, such that $\sigma_{ij} \rightarrow f$. In other words, setting the RIS elements to the states σ_{ij} yields the macroscopic function f. The definition of σ_{ij} depends on the RIS design and can be, e.g., cell phases, cell impedances, or even voltages applied to embedded actuators.

The above EM function examples are ideal in terms of performance. In practice, each has a degree of efficiency, ϵ, when implemented, e.g., due to material losses, RIS design specifications or RIS multitasking (i.e., serving two or more EM functions at the same time). Extending the notation, we can write `PLANAR_ABSORB`$(\vec{r}) \rightarrow \epsilon$, and ϵ : `PLANAR_STEER`$(\vec{r}) \rightarrow \vec{d}$. The definition of ϵ can be adapted depending on the set system specifications. E.g., for `PLANAR_ABSORB`, it can be defined as either the maximum reflected or refracted power towards any direction, or as the attenuation degree of the impinging wave. For `PLANAR_STEER` it can be defined as the gain towards the intended direction of departure, the ratio of power towards $\vec{d}$ divided by the total reradiated power, etc.

Complimenting the requirements for defining the RIS resource notion, note that a RIS has been shown to be able to serve many EM functions simultaneously [2]. In particular, for a set of functions $\epsilon_1 : f_1 \rightarrow \sigma_{ij}^{(1)}, \ldots, \epsilon_k : f_k \rightarrow \sigma_{ij}^{(k)}, \ldots, \epsilon_n : f_n \rightarrow \sigma_{ij}^{(n)}$, there exists a single $\sigma_{ij} : \sigma_{ij}\big(\sigma_{ij}^{(1)}, \ldots, \sigma_{ij}^{(k)}, \ldots, \sigma_{ij}^{(n)}\big)$ such that:

$$\sigma_{ij} \rightarrow \{\epsilon_1' : f_1, \ldots, \epsilon_k' : f_k, \ldots, \epsilon_n' : f_n\}.$$

Notice that the merging of $\sigma_{ij}^{(1)}, \ldots, \sigma_{ij}^{(k)}, \ldots, \sigma_{ij}^{(n)}$ into σ_{ij} yields the same functions, but alters the efficiency of each.

With the above definitions, we proceed to define the notion of the RIS resource, by first drawing a parallelism to a more common and well-understood resource, the CPU. A CPU is a resource shared by applications. When a single application is running on a CPU, it is allotted the full number of clock ticks per second that the specific CPU can offer. When a number of additional applications is running on the same CPU, the original application gets a reduced number of clock ticks. The resource slice is the ratio of the new allotted clock ticks, divided by the full number of clock ticks allotted when ran alone on the CPU.

Definition 5.1 (RIS as a resource). A RIS is a resource that is shared by macroscopic EM functions.

Definition 5.2 (Resource slice). The resource slice allotted to an EM function f is defined as the ratio r such that

$$r_f = \frac{\epsilon'}{\epsilon}, \tag{5.1}$$

where ϵ is the efficiency of f when hosted alone on the RIS, and ϵ' is its efficiency when hosted with other EM functions on the same RIS, $\sigma_{ij} \to \{\epsilon' : f, \ldots, \epsilon'_k : f_k, \ldots, \epsilon'_n : f_n\}$.

5.2 Codebook Entries Multiplexing Approach

The most straightforward strategy for a RIS unit to support multiple functionalities is the multiplexing of the codebook entries computed during the manufacturing phase. This method, as supported by the current thesis, is more efficient than the current RIS sharing methods and can be established in realistic deployments thanks to the millisecond-level required computational time. The multiplexing codebook entries are sourced by the observation that the deviation with respect to the RIS configuration for a specific EM functionality is strongly linked with the deviation from the optimal performance.

Specifically, when multiple users are served simultaneously, the multiplexing of codebook entries seeks a common RIS configuration CC that balances fairness and efficiency according to pricing policies. To evaluate the impact of sharing, we define the *configuration deviation* of the user that demands the function f as

$$\Delta_f = d(\mathrm{CC}, \mathrm{CE}_f) \in [0, 1], \tag{5.2}$$

where CE_f is the stand-alone optimal configuration and $d(\cdot, \cdot)$ is a normalized distance function determined by the deviation of RIS configurations in the two cases. By definition, $\Delta_f = 0$ when the shared and optimal configurations coincide, while $\Delta_f = 1$ corresponds to the maximum possible mismatch. We state that the resource slice r_f is strongly connected with the configuration deviation Δ_f. Formally,

$$\begin{aligned} r_f &= g_f(\Delta_f), \\ g_f &: [0,1] \to [0,1], \\ g_f(0) &= 1, \\ g_f(1) &\geq 0, \end{aligned} \tag{5.3}$$

where $g_f(\cdot)$ denotes the mapping that captures the dependence of retained performance on configuration deviation. The exact form of g_f can be obtained by emulations or simulations, but its existence expresses the fundamental link between Δ_f and r_f.

Firstly, we support this statement via analysis in both far and near field. For the former, an $(M \times N)$ RIS manipulates an incident plane wave originating from angle (θ_i, ϕ_i) to redirect it toward a desired direction (θ_d, ϕ_d). The ideal continuous RIS phase profile ϕ_{opt} ensuring constructive interference at (θ_d, ϕ_d) is defined by the geometric relationship:

$$\phi_{opt}(m,n) = \mod\left((\boldsymbol{k}_d - \boldsymbol{k}_i) \cdot \boldsymbol{r}(m,n), 2\pi\right) \tag{5.4}$$

where $\boldsymbol{k}_i$ and $\boldsymbol{k}_d$ are the wave-vectors for the incident and desired directions, respectively, and $\boldsymbol{r}(m,n)$ represents the spatial position of RIS element (m,n). The electric field amplitude at the targeted far-field direction, resulting from this ideal configuration, $|E_{ideal}(\theta_d, \phi_d)|$, can be computed by summing contributions from all RIS elements coherently and assuming the optimal phase precisely compensates for propagation differences as [110]:

$$E_{ideal}(\theta_d, \phi_d) = E_{inc} \sum_{m,n} e^{j\boldsymbol{k}_d \cdot \boldsymbol{r}(m,n)} \tag{5.5}$$

In the case of the multiplexing of codebook entries and for a specific RIS functionality, we result in a deviation from the optimal configuration $\Delta\phi(m,n)$. Then, the actual EM field is computed as:

$$E_{act}(\theta_d, \phi_d) = E_{inc} \sum_{m,n} e^{j[\boldsymbol{k}_d \cdot \boldsymbol{r}(m,n) + \Delta\phi(m,n)]} \tag{5.6}$$

Going back in the definition of RIS as network resource, the Eq. (5.1) can be expressed as:

$$r_f = \frac{\epsilon'}{\epsilon} = \frac{E_{act}(\theta_d, \phi_d)}{E_{ideal}(\theta_d, \phi_d)} = \frac{\left|\sum_{m,n} e^{j\Delta\phi(m,n)}\right|}{MN} \tag{5.7}$$

leading to the conclusion that the resource slice allocated function can be computed by the normalized in respect of RIS dimensions deviation from the optimal phase shifts.

Expanding our analysis in the near-field case, the RIS configuration is designed to focus EM energy at a nearby focal point $\boldsymbol{f}$. The ideal phase configuration ϕ_{opt} to achieve maximum constructive interference at $\boldsymbol{f}$ is determined by:

$$\phi_{opt}(m,n) = \mod(-k_0 R_{m,n}(\boldsymbol{f}), 2\pi) \tag{5.8}$$

where $R_{m,n}(\boldsymbol{f})$ is a constant entity and expresses the distance between RIS element (m,n) and focal point $\boldsymbol{f}$, and $k_0 = \frac{2\pi}{\lambda}$. The optimal EM field amplitude at the focal point can be expressed as:

$$E_{ideal}(\boldsymbol{f}) = \sum_{m,n} \frac{E_{elem}(m,n)}{R_{m,n}(\boldsymbol{f})} \tag{5.9}$$

where $E_{elem}(m,n)$ the electric field contribution for each specific RIS element. For configurations with deviations $\Delta\phi(m,n)$, the EM field at the focal point $|E_{act}(\boldsymbol{f})|$ becomes [110]:

$$E_{act}(\boldsymbol{f}) = \sum_{m,n} \frac{E_{elem}(m,n) e^{-j\Delta\phi(m,n)}}{R_{m,n}(\boldsymbol{f})} \tag{5.10}$$

Again, assuming uniform element excitation magnitudes $E_{elem}(m,n)$, the simplified near-field efficiency is also a function of the $\Delta\phi(m,n)$ as before. Mainly:

$$r_f = \frac{\epsilon'}{\epsilon} = \frac{\left|\sum_{m,n} \frac{e^{-j\Delta\phi(m,n)}}{R_{m,n}(\boldsymbol{f})}\right|}{\sum_{m,n} \frac{1}{R_{m,n}(\boldsymbol{f})}} \tag{5.11}$$

The findings of this analysis show that multiplexing among pre-computed codebook entries faithfully captures the underlying EM behavior. In both far-field and near-field cases, the relationship between phase-shift deviation and efficiency confirms the algorithms' theoretical validity and practical applicability.

In order to validate the above analysis, we perform a set of emulations. Specifically, we compute 18 codebook entries for different users' position via MATCH. The optimal SNR values for each user position is illustrated in Fig. 5.1. To validate Eq.(5.3), we

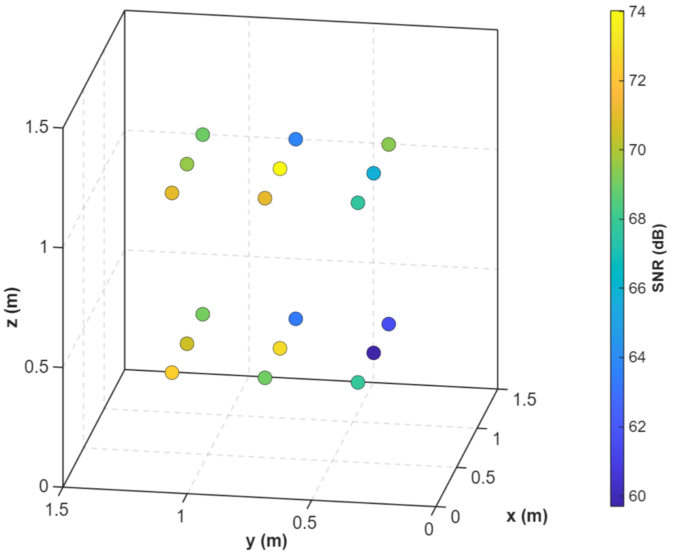

Figure 5.1: Users' position within the emulation setup and optimal SNR values as computed during codebook compilation.

deliberately perturb the RIS configuration away from its optimal value. At each stage, we impose a maximum acceptable deviation and repeat the procedure 50 Monte Carlo trials over all user positions to ensure statistical robustness. For each emulation trial, we compute the resource slice defined in Eq.(5.1), using ϵ as the optimal SNR (linear scale) and ϵ' as the SNR observed under the perturbed configuration. The resulting trends are reported in Fig. 5.2.

The results confirm that larger deviation bounds consistently reduce user efficiency, as quantified by r_f. For maximum deviations below 40°, the mean r_f decreases from 0.99 to 0.83, with closely matched medians from 0.98 to 0.82. Increasing the bound to 80° further degrades performance, with the mean r_f dropping to approximately 0.76–0.48 and the median to 0.75–0.45. Beyond this regime, efficiency collapses: as the maximum deviation approaches 180°, the mean r_f falls from 0.29 to 0.0133 and the median from 0.28 to 0.0062. These results reveal a clear correlation between configuration deviation and degradation of efficiency r_f. It should be underlined that we do not claim that configuration deviation exactly measures performance deviation, only that it provides a fast and robust estimation. This validation underpins codebook-entry multiplexing as a practical approach for real-time resource allocation and complex RIS functions.

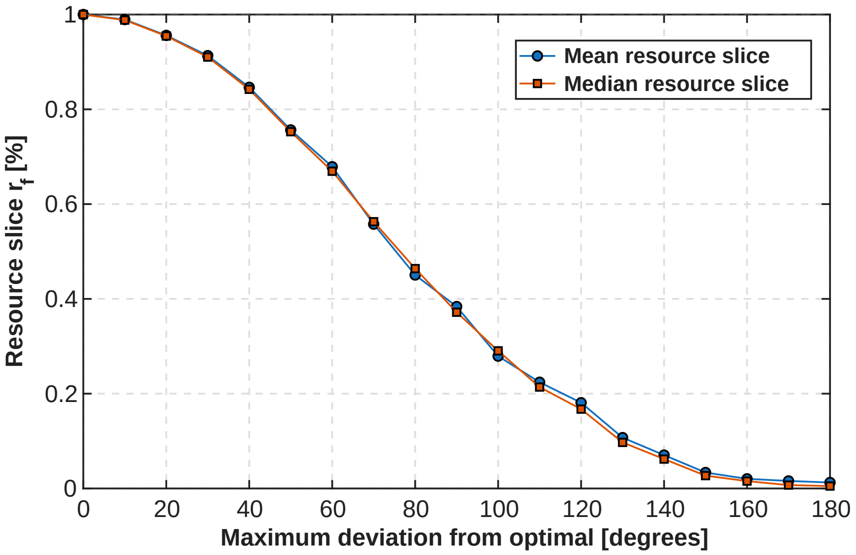

Figure 5.2: Emulation results about the resource slice behavior per maximum acceptable deviation from optimal RIS configuration.

5.3 COMMON & PRIME

In this section, we introduce the first algorithms for multiplexing codebook entries: COMMON and PRIME. The workflows of COMMON and PRIME are illustrated in Fig. 5.3

5.3.1 COMMON & PRIME Descriptions

The flow of COMMON is quite simple. Mainly, it retrieves the codebook entries for the users from the respective database, pre-processes them, and then sets each RIS element to the most frequent value among the codebook entries resulting to the final, common RIS configuration. Its workflow is presented in Alg. C.1 in Appendix C.

The RIS resource allocation procedure, in addition to the definition of the common RIS configuration, also requires a pricing mechanism to ensure that the sharing of RIS elements leads to the proper distribution of the resources to each user. Since COMMON lacks such a mechanism, we introduce the PRIME algorithm- its workflow is presented in Alg. D.1, that is included in Appendix D. It is designed for allocation of the RIS resources via the pricing-based multiplexing of codebook entries. PRIME ensures that each user's performance reflects their pricing level.

The required inputs for the algorithm include the dimensions of the RIS (M, N), the total number of users K utilizing the RIS unit for enhanced performance within

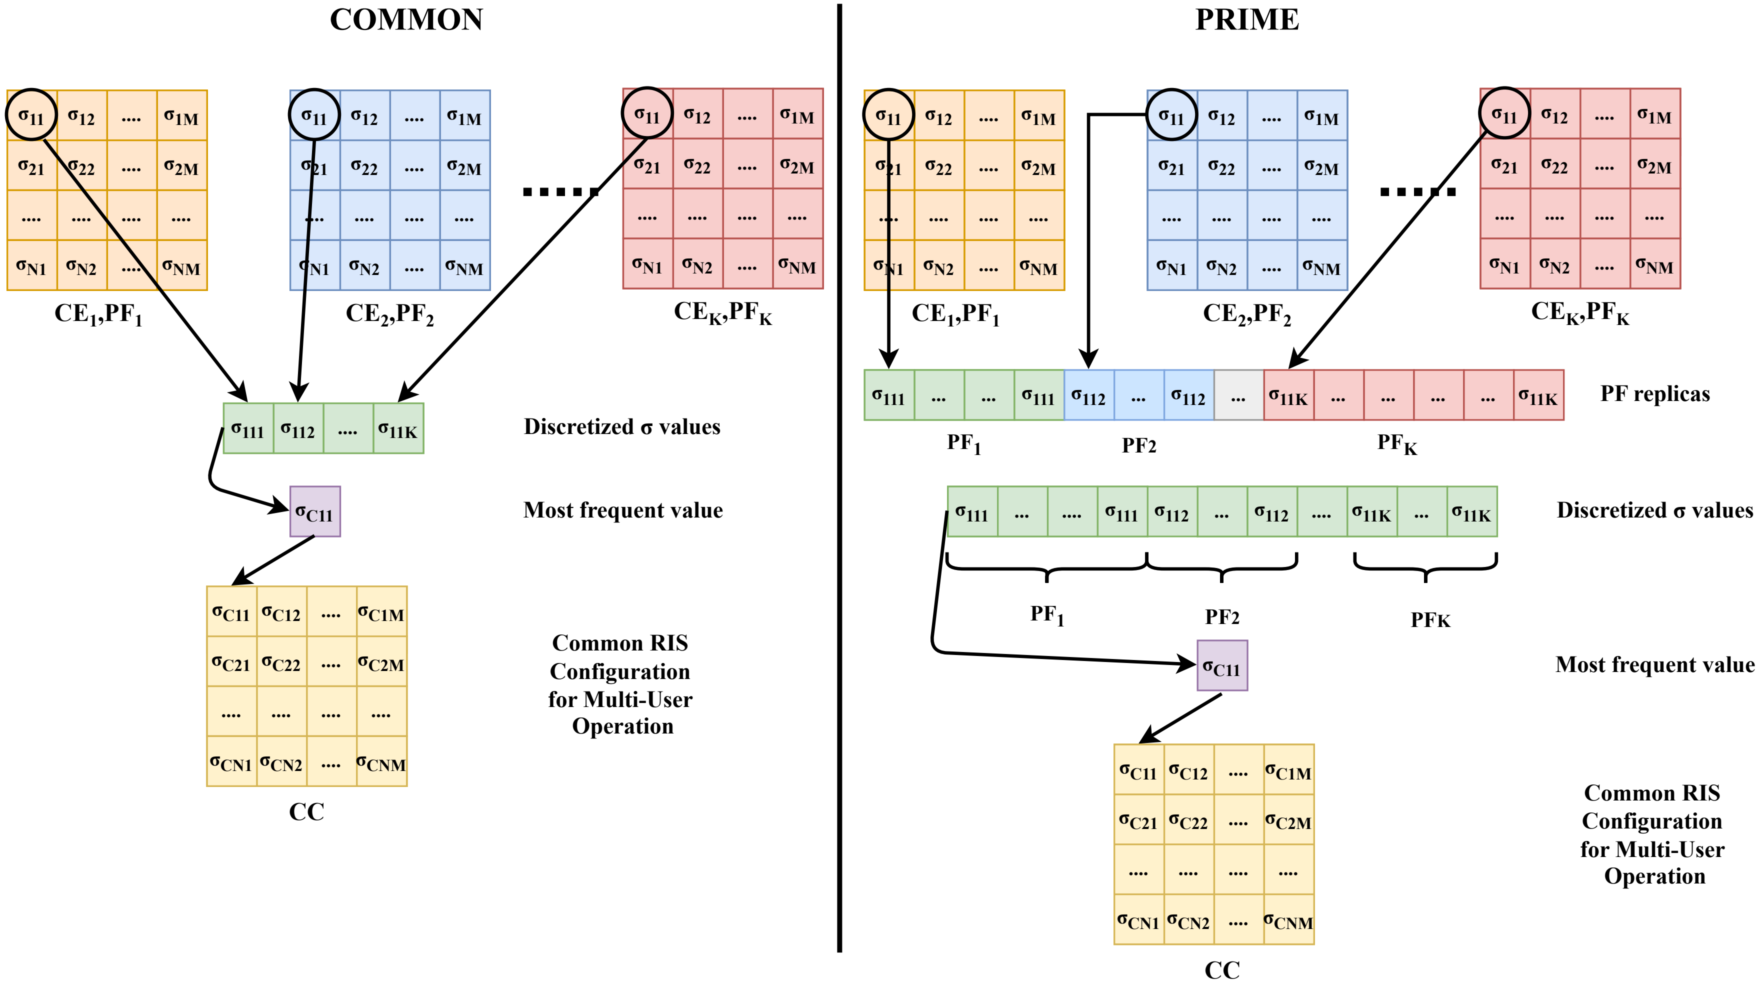

Figure 5.3: Workflow of COMMON (left) & PRIME (right) algorithms.

the communication network, and a discretization parameter N_d. Additionally, the algorithm must have access to the Codebook Database, which contains the entries for the optimal RIS configuration for each user. It is important to note that the computation of these codebook entries is not part of the algorithm and is assumed to be an input.

The adaptation of the PRIME algorithm to existing stakeholder pricing policies is achieved through the definition of the Payment Factor (PF). The PF is a positive integer assigned to each user upon entering the resource allocation process. Stakeholders can define multiple discrete PF values, assigning them in a way that quantitatively reflects the varying levels of performance granted to users based on their pricing tiers. This approach accommodates heterogeneous user demands and requirements by ensuring that users with higher PFs receive proportionally greater influence in the allocation process. All users' PFs must be provided as input to the algorithm to accurately incorporate these pricing levels into the resource allocation mechanism.

PRIME begins by retrieving the codebook entries, CE, matching the user requests for a RIS functionality. Next, it discretizes all the elements of CE based on the discretization parameter N_d leading to the DE matrices. Using the PF values, the algorithm then creates an equal number of replicas for each DE_k, which is why PFs must be natural numbers. After generating all the replicas of the discretized values, the algorithm determines the common RIS configuration, CC, by selecting the most frequent values among the replicas for each RIS element. As is evident, users with higher PF values have a greater influence on the CC due to the increased number of replicas of their entries that participate in the final selection, ensuring the pricing part of the PRIME.

The evaluation of the algorithm's feasibility in real-world applications is closely tied to the computational time required for pricing-based codebook entries multiplexing. Analyzing the algorithm's workflow, the computational complexity depends on PF values, the dimensions of the RIS unit, and the total number of users.

First, the algorithm discretizes the codebook entries for each user across the RIS elements, resulting in a computational complexity of $O(KMN)$. Next, the algorithm generates configuration replicas based on each user's PF values, adding another layer of complexity. These replicas are then combined, resulting in a computational complexity proportional to the number of users and their respective PFs. Consequently, this part of the algorithm incurs a complexity equivalent to the sum of all PF values, denoted as $O(\sum PF)$. Therefore, the overall computational complexity of the PRIME algorithm is expressed as $O(KMN + \sum PF)$.

5.3.2 Evaluation of PRIME & COMMON

In this section, firstly, the PRIME algorithm is evaluated using simulations, as this method has been described in Section 3.2. We investigate the critical relationship between each user's received performance and their respective PF. We compare PRIME with COMMON based on tests.

For the simulations and evaluation of PRIME in pricing-based resource allocation, we consider a communication network that consists of 4 different base stations. The total number of users are 5 and there is one RIS unit. To encompass a more general concept, two users are served by the same base station, while the remaining three are served by different ones. The LoS connection between the users and their respective base stations is blocked. The users are positioned opposite a RIS unit with the goal of restoring the LoS connection. The RIS unit elements are shared among the five users. The resource allocation is established via a PF assigned to each user. Each user's PF is unique.

As mentioned previously, only the pricing-based multiplexing, not the computation of the codebook entries, is part of the PRIME algorithm. Therefore, the optimal RIS configuration for each user when served individually by the RIS must be measured and saved in the Codebook Dataset. For this procedure, we use the PICCS. This procedure generates a codebook entry for each user, consisting of 66 values for impedance capacitance. The impedance values range from 0.4 to 0.9 μF. The PRIME algorithm simply retrieves the codebook entries from the Codebook Database and multiplexes them according to the PFs assigned to each user. For the sake of generality of the evaluation, 205 Monte Carlo combinations of PFs are assigned to the users. For the remainder, the term combination will refer to the set of PF values that are assigned to the users each time. The assignment of the PFs is a random procedure. The values vary from 1 to 20. The PRIME uses a discretization parameter of 0.05 for the determination of the common RIS configuration, CC.

In Fig. 5.4, the values of the PF combinations are depicted. Users with higher PF values have an average PF of 16.8 with a standard deviation of 3.07, whereas users with lower PF values have an average PF of 4.8 with a deviation of 2.82. Within

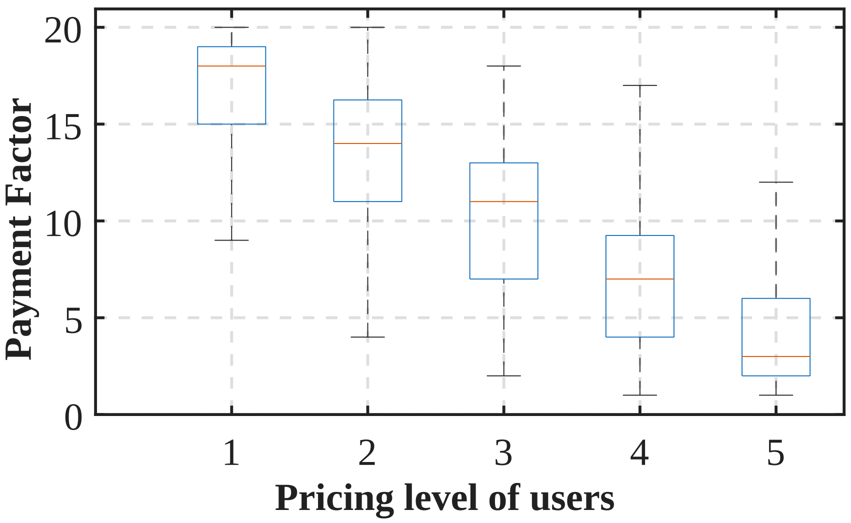

Figure 5.4: Boxplot of the PF per pricing level of the combinations.

each combination, since no user shares the same PF value with another, five distinct pricing levels are formed.

The measurements of S21 for each user, once the RIS is set with CC, are computed for each combination of PFs. Since the use of PICCS ensures that the codebook entries include the optimal RIS configuration for serving each user individually, the performance for the allocation of RIS resources with respect to user k during combination c is calculated as follows:

$$P_{kc} = \frac{S21_{opt_k} - S21_{kc}}{S21_{opt_k}} \tag{5.12}$$

The performance values for each user during the different combinations are depicted in Fig. 5.5. The average values for users, from highest to lowest, are 98%, 94%, 93%, 92%,and 87%, with variations of 2%, 5%, 4%, 3%, and 2%, respectively. As is evident, the measurements confirm that the PRIME can effectively allocate the RIS resources via the proper sharing of its elements for multiple users with different pricing levels.

An aspect that requires further assessment is the relationship between each user's PF and their performance. PRIME offers a linear relationship between PF assigned to each user and the received performance. This relationship is verified in Fig.5.6. Specifically, three out of a total of 205 scenarios, each involving five users with varying PF values, are randomly selected and illustrated. It is evident that in all cases there is a linear proportionality between the assigned PF and the received performance.

The intercept of this line primarily represents the minimum efficiency level, in-

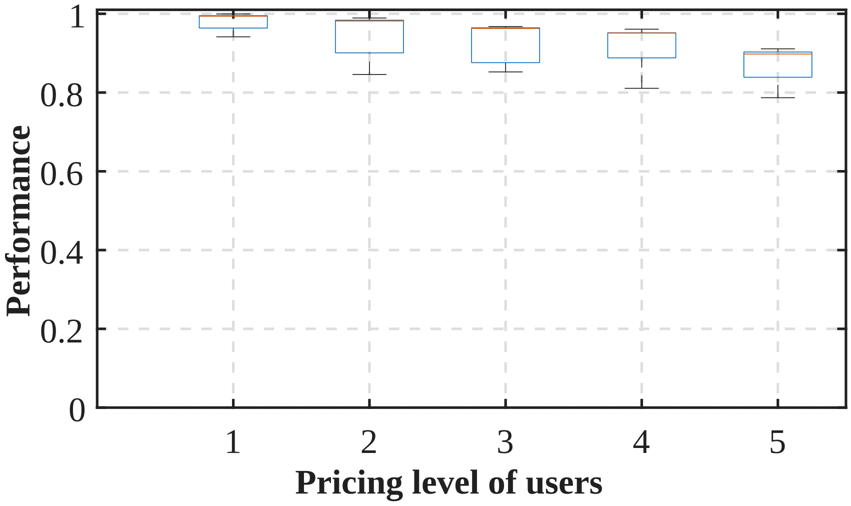

Figure 5.5: Boxplot of the users' performance per pricing level with PRIME.

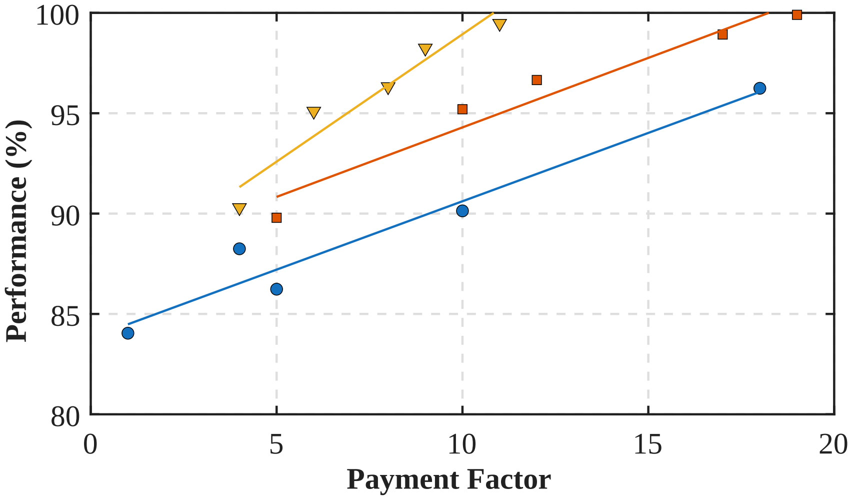

Figure 5.6: The linear relationship between users' PF and performance with the usage of PRIME.

dicating the performance that a basic user receives. The slope of the line denotes the performance enhancement achieved with the addition of each PF unit, as defined by telecommunication stakeholders. Accurate estimation of the intercept and slope in each network, combined with the PF assigned to each user, allows for easy approximation of their performance.

The intercept and the slope of PRIME's performance line are strongly influenced by the total number of users. As a next step, we also investigate how the intercept and the slope are affected by the users' PF. This analysis aims to determine how the performance of a specific user with a given pricing level can be affected by the other users participating in the same network. To this end, we perform a correlation

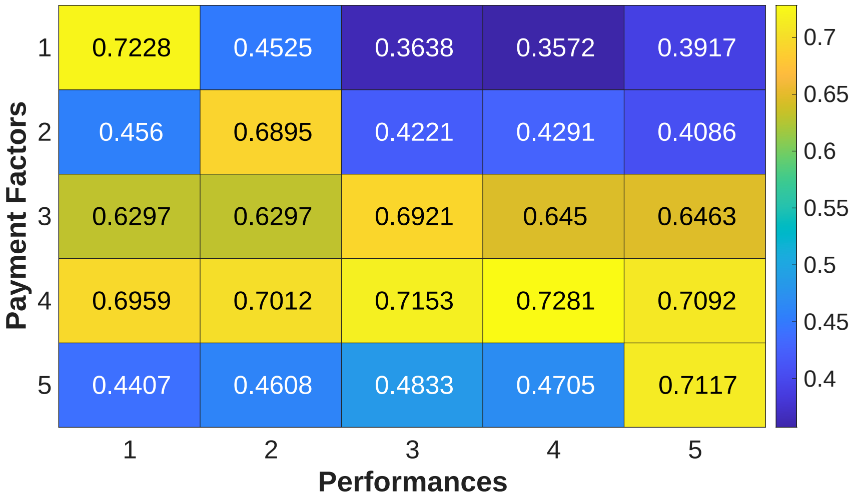

Payment Factors \ Performances	1	2	3	4	5
1	0.7228	0.4525	0.3638	0.3572	0.3917
2	0.456	0.6895	0.4221	0.4291	0.4086
3	0.6297	0.6297	0.6921	0.645	0.6463
4	0.6959	0.7012	0.7153	0.7281	0.7092
5	0.4407	0.4608	0.4833	0.4705	0.7117

Figure 5.7: The correlation matrix among all users' performance and PF values.

analysis between the PF assigned to all users and the performance received by an individual user.

Subsequently, we investigate the correlation between each PF value and the corresponding performance. As mentioned previously, no users have the same PF value in any combination. Therefore, we can shape five different pricing levels. The results are illustrated in Fig. 5.7. The highest correlation is observed between each user group's PF values and their respective performance (major diagonal elements in Fig. 5.7). However, there is also a significant correlation with other PF user groups, ranging from 0.3572 to 0.7153. The highest correlations are seen in the groups with mid-level PF values, with a notable correlation greater than 0.6297, indicating a relatively strong relationship.

Having detected the correlation between user groups and performances, the objective analysis for intercept and slope of the performance line follows. For this task, we scale the PF values during each combination. Therefore, each scaled PF value is mapped in relation to the PF values of the rest users. After this action, each scaled PF is aligned with the respective performance, and the intercept and the slope factors are computed. The respective boxplot is illustrated in Fig. 5.8.

The mean value of the slope is 0.0928, with a standard deviation of 0.0185. For the intercept, the mean value is 0.8820, with a standard deviation of approximately 0.04. This indicates that PRIME ensures that all users can achieve performance levels between 84% and 92%, while users with higher PF values can approach 100% performance. A higher PF ensures that, irrespective of the PF values of the other users

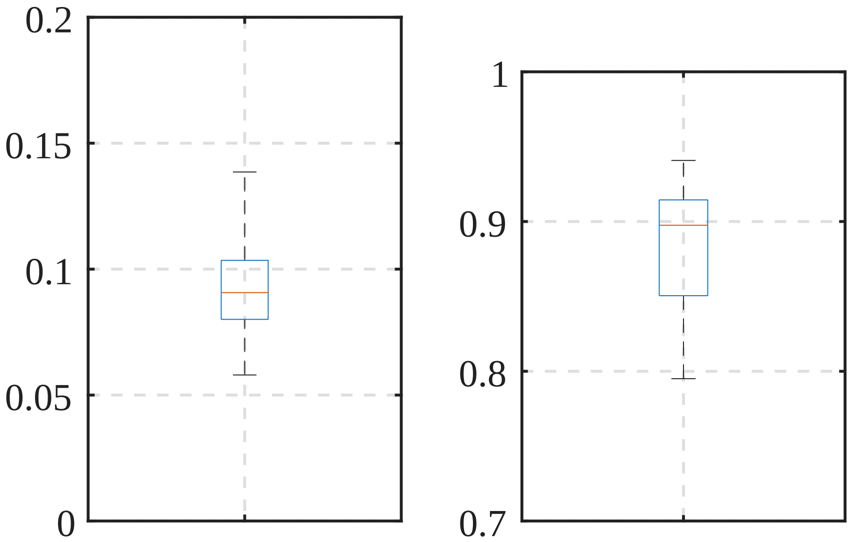

Figure 5.8: Boxplot of the slope (left) and the intercept (right) of the performance line based on the PFs of the pricing levels.

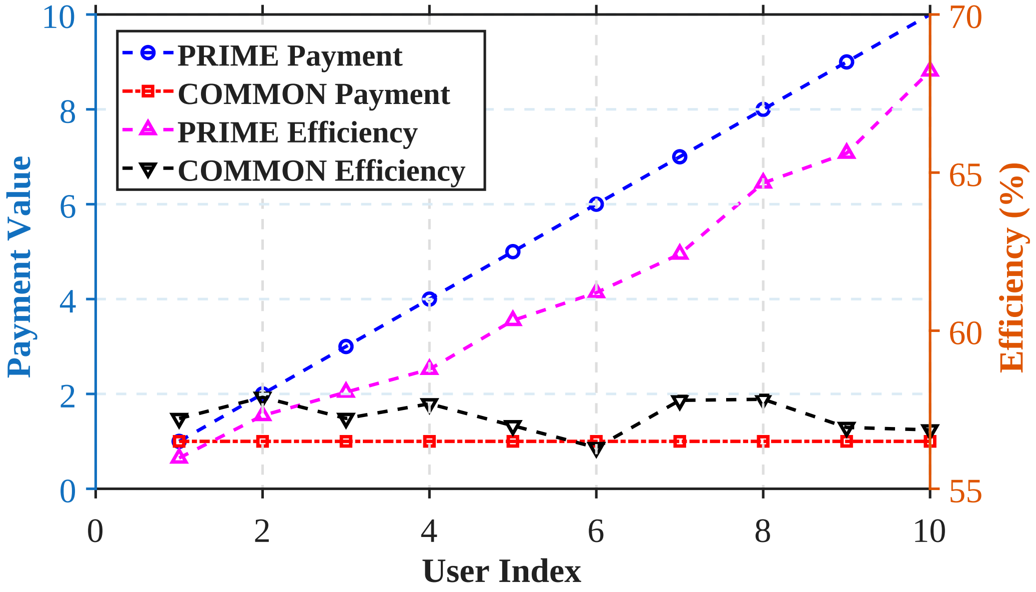

Figure 5.9: User efficiency for PRIME and COMMON algorithms

and the total number of them, the performance will be close to optimal.

As concerns the comparison of the two algorithms, we use tests. As already discussed, COMMON can multiplex codebook entries so that the same RIS unit can simultaneously serve multiple users. PRIME can accomplish the same task, but with the added capability of considering unique pricing levels for each user and assigning priorities accordingly. Essentially, COMMON can be viewed as a special case of PRIME where all users share the same pricing level. Specifically, we assume that a 100×100 RIS unit is serving 10 users simultaneously, with a discretization parameter of 0.05. As shown in Fig. 5.9, PRIME assigns PFs to each user incrementally, starting from 1 and increasing to 10 (with User ID=1 assigned PF=1, User ID=2 assigned PF=2, and so on). The codebook entries are simulated with varying values from 0 to 1.

Fig. 5.9 demonstrates that COMMON can effectively share the RIS by multiplexing codebook entries, resulting in an efficiency of approximately 57% for all 10 users. Regarding the PRIME algorithm and the associated sharing and pricing results, it is clear that the improved efficiency levels of high-PF users do not come at the expense of others, ensuring network fairness. The higher PF values can achieve efficiencies about 27% higher.

We also compare the computational time required by COMMON and PRIME. In this set of tests, the number of users varies from 5 to 100, and the RIS dimensions range from $M = N = 10, 25, 50, .., 200$. In the PRIME algorithm, the PFs are randomly assigned values from 1 to 10. All computations are accelerated using the GPU unit. The results are presented in Fig.5.10a for COMMON and Fig.5.10b for PRIME.

Because of its workflow, the PRIME algorithm benefits significantly from GPU acceleration. As a result, the computational time required for PRIME is only about 40% greater than that required for COMMON. Even for very large RIS units serving many users, the computation time remains limited to approximately 0.1 ms. The ms-level latency of the PRIME renders it applicable also in dynamically changing network conditions, e.g., vehicular networks. We remind that all runtime measurements are implemented in MATLAB, which is typically slower than C/C++ or Fortran. This highlights the practicality of the proposed algorithm for computing RIS configurations to efficiently serve multiple users with large RIS dimensions in realistic scenarios, especially regarding computational demands.

5.3.3 Discussion

Defining the RIS as a network resource is a necessary step for multi-user PWEs, because it provides an explicit abstraction of what is shared (macroscopic EM functions), how it is shared (via a common configuration), and how much each user retains under sharing (via the resource slice). Within this abstraction, codebook-entry multiplexing becomes a practical operating-phase mechanism: instead of per-element online re-optimization, the RIS combines precomputed single-user codebook entries into one

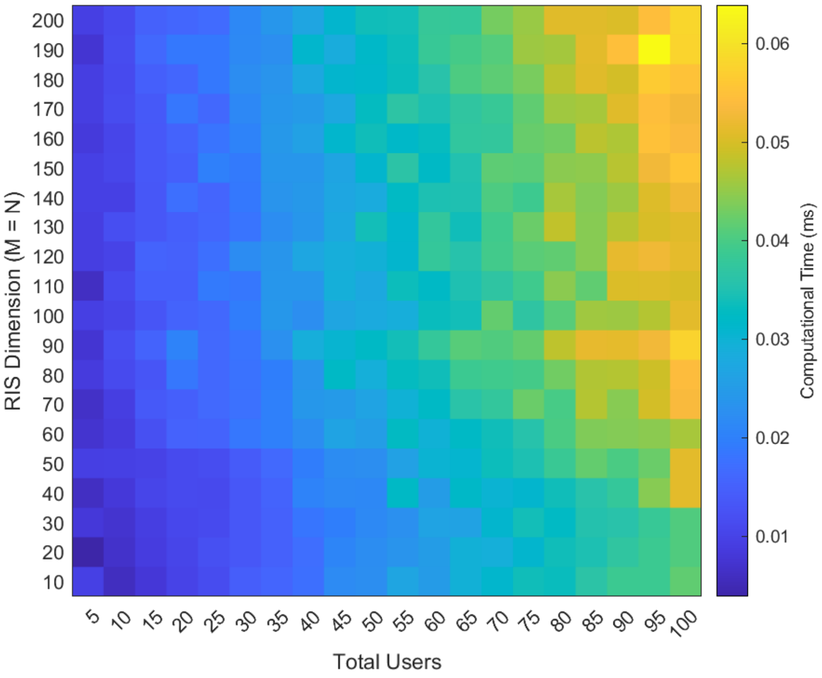

(a) COMMON algorithm.

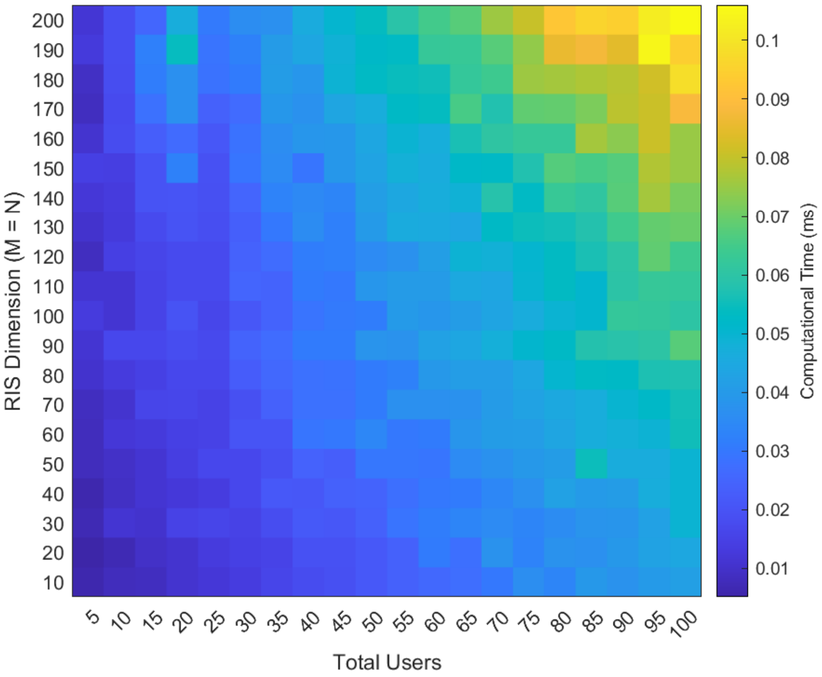

(b) PRIME algorithm.

Figure 5.10: Heatmaps of the required computational time of COMMON (top) and PRIME (bottom) with respect to the total number of users and the RIS dimensions.

common configuration, and the resulting performance loss is interpretable through the configuration-deviation principle developed earlier.

Building on this framework, COMMON and PRIME constitute the first dedicated algorithms for codebook-entry multiplexing in multi-user RIS operation. COMMON provides the baseline mechanism for concurrent service by constructing a common configuration through element-wise consensus across user entries. PRIME extends this baseline by incorporating operator-defined pricing tiers through the Payment Factor (PF), which weights each user's influence on the shared configuration in a transparent and easily deployable manner. Importantly, PRIME does not require tight receiver–RIS control loops or frequent CSI-driven re-optimization; it relies only on stored codebook entries and lightweight post-processing, aligning with the latency constraints of PWEs.

Simulation results validate both the pricing fidelity and the practicality of PRIME. Across Monte Carlo PF combinations, user performance exhibits a clear and approximately linear dependence on PF, indicating that operator policy can be translated into predictable performance differentiation. At the same time, the measured runtimes remain at the sub-millisecond level even for large RIS sizes and many users, and the additional overhead relative to COMMON is modest under GPU acceleration. Collectively, these outcomes support the main claim of this section: codebook-entry multiplexing is a viable operating-phase primitive for multi-user RIS control, and PRIME offers a low-latency, policy-consistent instantiation that enables concurrent service with controlled performance trade-offs and deployable computational cost.

5.4 Benchmarking of RIS sharing methods

Multiplexing of RIS codebook entries, as an efficient RIS-sharing and resource-allocation method, is one of the main contributions of this thesis. Accordingly, we must benchmark its beam-splitting performance—specifically under pricing-based control—against established methods. As noted in Chapter 2, the baselines are TDMA and RIS-surface segmentation. The comparison is conducted on the real testbed of Section 3.3.

5.4.1 Benchmarking procedure Description

First, for each of the k users, we compute the standalone gain G_k^{std} when served by the RIS alone, using the manufacturer-optimal configuration defined in Eq.3.23. The common RIS configuration CC is then derived under TDMA, segmentation, and RIS multiplexing. Using this CC for each case, we compute the multi-user gain G_k^{multi}. Aligned with the definition of the RIS as a network resource (Eq.5.1), the per-user efficiency in the multi-user scenario is

$$r_f = \frac{\epsilon'_k}{\epsilon_k} = \frac{G_k^{multi}}{G_k^{std}}. \tag{5.13}$$

5.4.2 Benchmarking results

Firstly, we compare the performance of three RIS sharing methods for serving multiple users equally. We assume the RIS simultaneously serves four users belonging to the same pricing tier, aiming to equally steer the beam towards each user. For multiplexing, we employ the COMMON algorithm. The segmentation approach evenly allocates 25% of RIS elements per user, with each sub-panel configured identically to their standalone optimal settings. In the TDMA approach, RIS configurations switch cyclically, adopting each user's ideal standalone configuration for 100 ms intervals. Measurement results are shown in Fig. 5.11. Clearly, multiplexing with COMMON presents a reliable alternative to the established segmentation and TDMA methods. The mean r_f values achieved are comparable, recorded as 54.18%, 48.97%, and 48.22% for COMMON, equal segmentation, and equal TDMA, respectively. The standard deviation of r_f for COMMON is 32.67% with a minimum and maximum of 21.35% and 95.92%, respectively. The segmentation approach shows a higher deviation of approximately 42.26%, achieving peaks up to 100% but also causing users with zero gains, signifying considerable unfairness. TDMA exhibits a standard deviation of 48.22%, with a minimum of 29.57% and maximum of 75.81%, thus closely matching the multiplexing performance.

Next, we explore RIS resource allocation across users of different pricing tiers. Each of the four users belongs to distinct pricing levels, and we present results for three scenarios. The multiplexing method utilizes PRIME algorithm with payment factors from highest to lowest tier as $5, 4, 2, 1$. The segmentation method assigns RIS elements proportionally at 50%, 25%, 15%, and 10% for the respective pricing tiers.

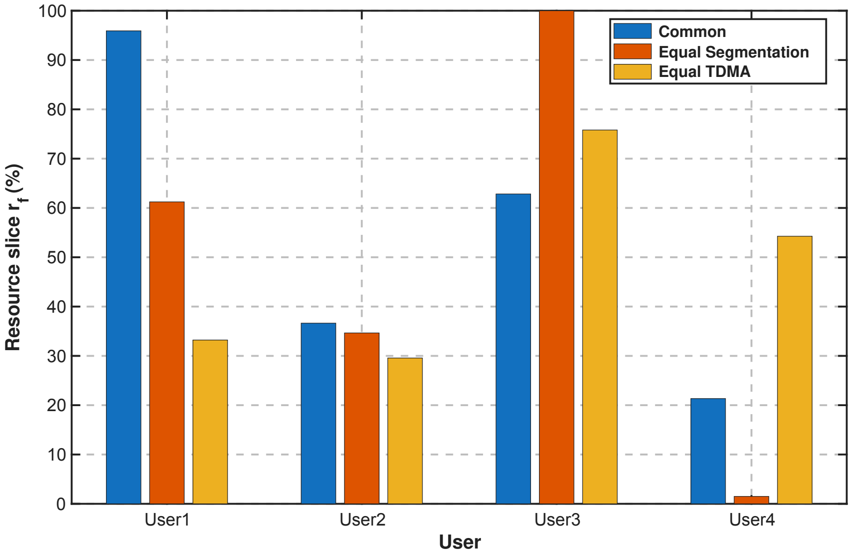

Figure 5.11: Measurement results for equal beam-splitting using multiplexing algorithm, segmentation and TDMA

Finally, the TDMA approach applies each user's optimal configuration cyclically for durations of 300 ms, 200 ms, 100 ms, and 50 ms according to their pricing tier. Results for these scenarios are illustrated in Figs. 5.12, 5.13, and 5.14, respectively, with user ordering from highest to lowest pricing tier.

Analyzing the results, PRIME uniquely achieves fair resource allocation aligned with user pricing tiers, although slight deviations occur, especially between the second and third tiers. Specifically, for PRIME Case 1, the r_f values are 95.65%, 45.18%, 54.80%, and 41.17%, while Case 2 results are 95.63%, 38.14%, 49.89%, and 47.92%, highlighting the challenges of multiplexing binary RIS configurations yet preserving general fairness. Case 3 achieves better fairness with r_f values of 51.47%, 50.11%, 45.44%, and 40.75% with lower performance behavior for all the users.

For segmentation, only Case 1 follows an acceptable fairness trend, but with substantially limited r_f for lower-tier users (16.81% for third-tier, 36.36% for lowest-tier). In segmentation Case 2, the allocation appears random with r_f values at 31.76%, 73.99%, 38.72%, and 52.72% respectively. The third segmentation case further indicates failure in fairness, where the top-tier user achieves just 30.75%, whereas others approach 100%.

Lastly, TDMA also fails to consistently deliver fair allocation. In Case 1, all users experience extremely low performance, with the second user receiving no improvement compared to no RIS scenario. Case 2 partially succeeds, benefiting only the

Table 5.1: Performance per Ranking Level for Different RIS Sharing Approaches

Approach	**Rank**	**Mean**	**Std**	**Min**	**Max**
Multiplexing (PRIME)	Highest	80.91	25.50	51.47	95.65
	2nd	43.81	6.01	38.14	50.11
	3rd	50.05	4.68	45.44	54.80
	Lowest	43.28	4.02	40.75	47.92
Segmentation	Highest	45.66	24.95	30.75	74.46
	2nd	78.35	19.84	61.04	100.00
	3rd	51.42	42.41	16.81	98.73
	Lowest	63.03	33.05	36.37	100.00
TDMA	Highest	43.15	36.27	16.53	84.46
	2nd	37.11	31.63	0.98	59.79
	3rd	14.15	15.08	0.00	30.01
	Lowest	31.26	48.45	0.00	87.07

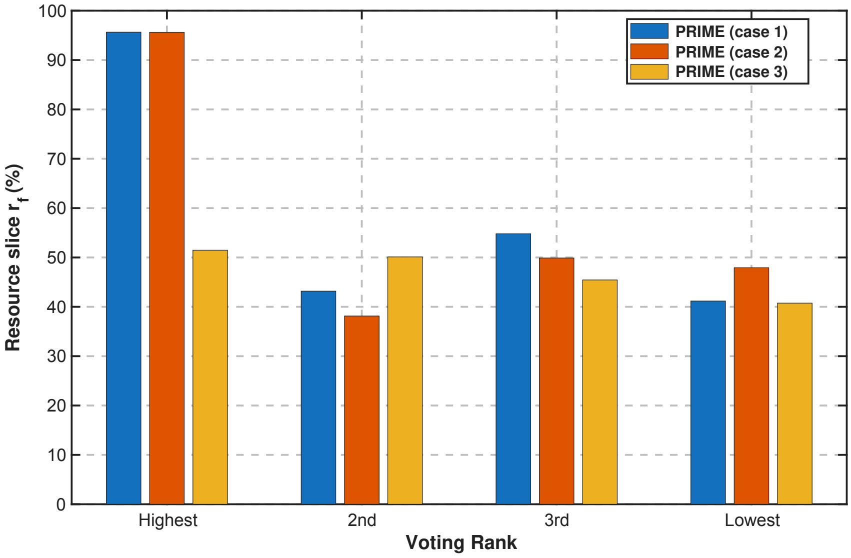

Figure 5.12: Measurement results for RIS resource allocation among different user-tiers using multiplexing algorithm.

first two users significantly (84.45% and 50.76%), whereas lower-tier users gain nothing. In Case 3, allocation remains random, yielding performance of 16.53%, 59.75%, 30%, and 87.07%, respectively. A statistical summary of these results is presented in Table 5.1, reinforcing the superiority of the multiplexing PRIME algorithm.

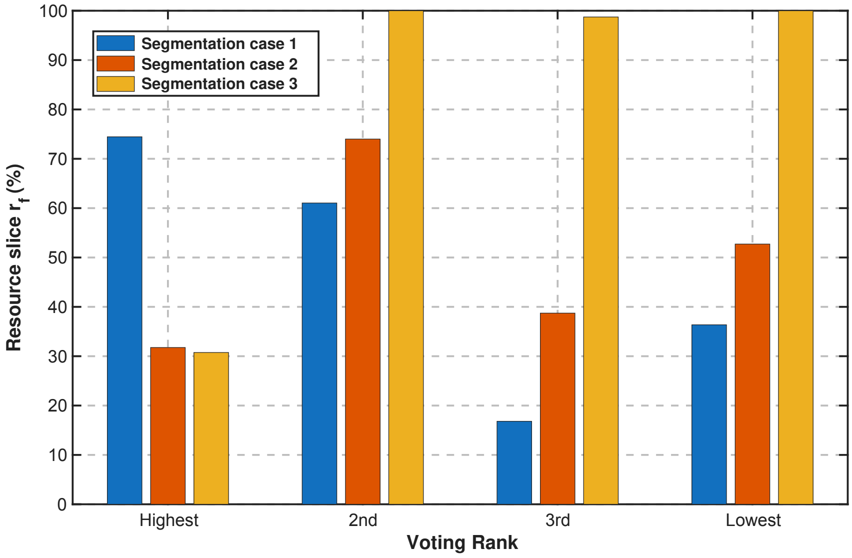

Figure 5.13: Measurement results for RIS resource allocation among different user-tiers using segmentation approach.

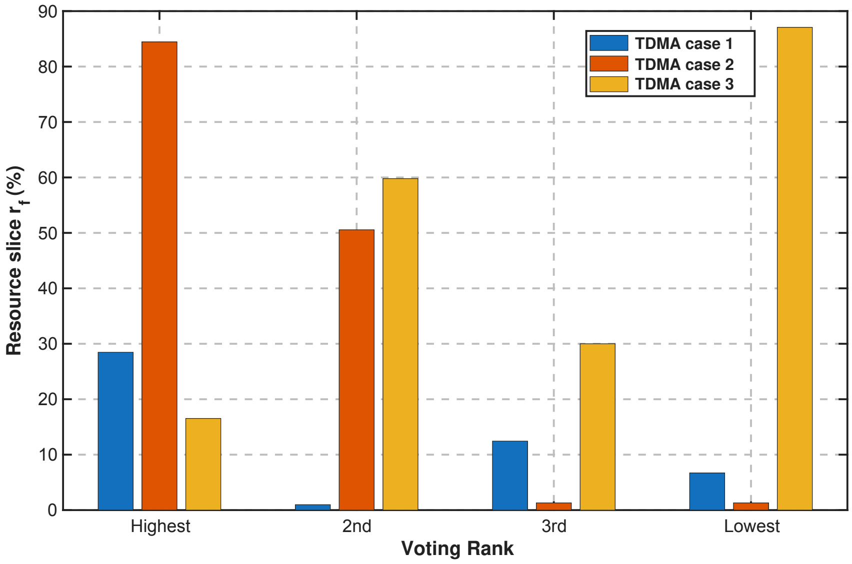

Figure 5.14: Measurement results for RIS resource allocation among different user-tiers using TDMA approach.

5.4.3 Discussion

The benchmarking on the real RIS testbed validates codebook-entry multiplexing as a practical operating-phase mechanism for multi-user PWEs, and positions it favorably against established sharing baselines (segmentation and TDMA). Using the resource-slice definition, the equal-tier experiment shows that COMMON achieves mean efficiency comparable to (and slightly higher than) equal segmentation and equal TDMA, while avoiding the pronounced instability observed in the baselines. In particular, segmentation can yield extreme unfairness (including zero-gain users) due to rigid aperture partitioning, whereas TDMA exposes users to temporal service fluctuations despite its simplicity.

More importantly, when users belong to different pricing tiers, PRIME is the only evaluated method that consistently translates tiering into a corresponding ordering of realized efficiencies. Across the examined scenarios, PRIME maintains the intended priority structure with modest deviations between adjacent middle tiers—an expected consequence of multiplexing under discrete (binary) configuration constraints—while preserving non-trivial performance for all users. In contrast, proportional segmentation and tiered TDMA frequently fail to enforce pricing-consistent allocations: segmentation can behave unpredictably because sub-panel isolation breaks the joint aperture gains, and TDMA can collapse in practice because time-slicing does not guarantee per-user gain once switching, synchronization, and topology-dependent focusing constraints are accounted for.

Overall, these measurements support three conclusions. First, multiplexing is a viable alternative to segmentation/TDMA for concurrent multi-user RIS service with millisecond-level control. Second, PRIME operationalizes stakeholder pricing policies through a lightweight weighting mechanism that yields pricing-consistent resource allocation in practice. Third, the observed failure modes of segmentation and TDMA highlight that "intuitive" sharing rules do not reliably map to EM performance; codebook-entry multiplexing, grounded in the RIS-as-a-resource abstraction, provides a more robust path to fairness-aware RIS integration in 6G PWEs.

5.5 LEVEL

The chapter's initial findings, together with the encouraging results from COMMON and PRIME, underscore that controlling the deviation between individual codebook entries and the common RIS configuration is central to multiplexing performance. However, COMMON and PRIME, mainly, multiplex codebook entries in order to realize Δ_f in Eq. (5.3). In contrast, LEVEL computes Δ_f directly, yielding controllable, predictable deviations from the per-user optimum. Tailored for resource allocation in RIS-assisted networks, LEVEL aligns with RIS as a shared resource and operationalizes resource slicing by explicitly managing performance trade-offs among concurrent users with heterogeneous service tiers. The full workflow is given in Alg.E.1.

5.5.1 LEVEL Description

LEVEL strategically minimizes the deviation described in Eq. (5.3) through a simple, real-time optimization process. This approach systematically accounts for varying user priorities, preserving the performance hierarchy intended by their respective pricing tiers. Specifically, consider K users simultaneously served by an $M \times N$ RIS, with their stand-alone optimal configurations denoted as $\mathbf{CE}(\cdot, k) \in \mathbb{R}^{MN}$. When LEVEL is applied, the RIS is set to a single common configuration $\mathbf{CC}_{\text{opt}} \in \mathbb{R}^{MN}$, which inevitably deviates from the individual optima. For each user k, we define

$$\Delta_{fk} = \frac{1}{MN} \sum_{i=1}^{MN} \Big| \mathbf{CC}_{\text{opt}}(i) - \mathbf{CE}(i, k) \Big|, \tag{5.14}$$

as the user-specific deviation, consistent with Eq. 5.2. This deviation yields an effective efficiency ϵ'_k. Assuming the stand-alone efficiency equals one, the efficiency and the corresponding resource slice allocated to user k are

$$\epsilon'_k = 1 - \Delta_{fk}, \tag{5.15}$$

$$r_k = \epsilon'_k. \tag{5.16}$$

For the determination of $\mathbf{CC}_{\text{opt}}$, LEVEL groups the users into three service tiers $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ with $|\mathcal{L}_1| < |\mathcal{L}_2| < |\mathcal{L}_3|$. Stakeholders assign positive weights $w_1 > w_2 > w_3$ that can be adjusted according to their pricing strategies, ensuring the algorithm directly integrates with economic models prevalent in network resource allocation

$$|\mathcal{L}_3|\, w_3 \;<\; |\mathcal{L}_2|\, w_2 \;<\; |\mathcal{L}_1|\, w_1 \;\leq\; |\mathcal{L}_3|\, w_3 + |\mathcal{L}_2|\, w_2.$$

This inequality guarantees balanced yet distinct performance differentiation among user levels, ensuring users with higher priorities attain superior service quality without overly penalizing lower-priority users and ensuring fairness within the network. Subsequently, LEVEL performs an independent optimization of each RIS unit cell, significantly reducing computational complexity compared to global optimization approaches and enabling real-time application of the algorithm. This scalar, single-variable optimization within each cell determines an optimal configuration value that minimizes the weighted deviation from individual user preferences. This process is formalized by the optimization objective at each RIS unit cell i:

$$\sigma_i^* = \arg\min_{\sigma} \left(\frac{\sum_{k=1}^{K} w_k \cdot |\sigma - \mathbf{CE}_i(k)|}{\sum_{k=1}^{K} w_k} \right)$$

where σ_i^* denotes the optimized RIS configuration for cell i. It then computes each user's efficiency with Eqs. (5.14), (5.15) and reports the average by tier. Finally, to respect hardware constraints, the resulting common configuration is quantized according to the RIS specifications.

Algorithmically, LEVEL exhibits computational complexity primarily determined by the RIS dimensions, the number of users, and the optimization tolerance. Since LEVEL performs a per-element optimization across an $M \times N$ RIS, and each per-element objective evaluation aggregates over K users, the overall complexity scales as $\mathcal{O}(KMN)$. In practice, runtime is dominated by the MN per-element solves, so RIS dimensions largely govern computational time, while K induces only a weaker, near-linear overhead within each solve. Notably, the algorithm remains practically efficient even in extensive deployment scenarios, demonstrating suitability for real-world RIS-assisted network implementations.

In summary, LEVEL systematically operationalizes RIS resource slicing, directly aligning RIS network resource allocation with stakeholder-driven economic considerations and multi-user performance management. Its design ensures transparent, predictable, and fair resource allocation, positioning it as a robust solution for next-generation RIS-assisted network deployments.

5.5.2 Evaluation of LEVEL

In this section, we evaluate LEVEL. Mainly, we demonstrate LEVEL and quantify its runtime for practical deployment via tests. Moreover, we return to emulations, as they have been described in Section 3.1.3, to compare LEVEL against COMMON/PRIME, focusing on multi-user efficiency and its correlation with the corresponding pricing tier.

For LEVEL demonstration, we consider a scenario with twenty users in total: three in Level 1, five in Level 2, and the rest in Level 3. The RIS is composed of 50×50 elements, with respective priority weights of 5, 3, and 1 for each level. The resulting user efficiencies are illustrated in Fig. 5.15. Specifically, the Level-1 users achieve

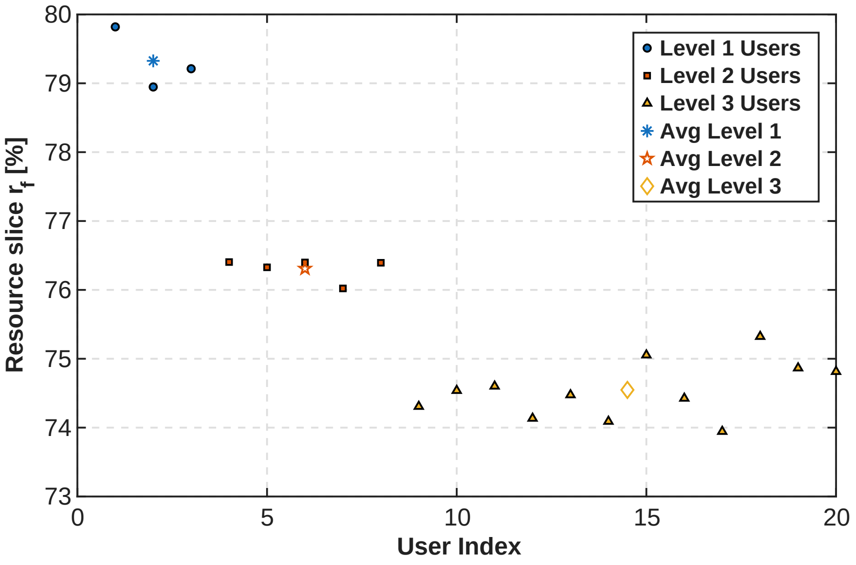

Figure 5.15: Resource slicing r_f with LEVEL usage in a scenario with 20 users belonging to different pricing tiers.

approximately 79% efficiency, the Level-2 users achieve about 76.5% on average, and the Level-3 users achieve approximately 74.5% on average.

Following, we evaluate LEVEL's capability for efficient resource allocation across different user combinations with the total number of users ranging from 5 to 30. For each case, approximately 10% of the users are assigned to the first level, 30% to the second level, and the remaining users to the third level. The RIS unit is consistently set to a 50×50 grid. The average efficiency results for each user level and user combination are shown in Fig. 5.16. In all cases, there is a discrete efficiency difference between the user levels. Specifically, when the number of users is fewer than 15, first-level users demonstrate about 10% higher efficiency compared to second-level users and up to 16% higher efficiency compared to third-level users. In scenarios with denser communication networks, the efficiency difference between each level decreases to approximately 3%. Second-level users maintain similar performance trends across different user combinations, while third-level users show slight improvements in these denser network cases.

As mentioned earlier, a key advantage of codebook-entry multiplexing is the negligible runtime overhead introduced during operation. To assess this overhead for LEVEL under practical RIS dimensions, we capture the computational time for RIS sizes 16×16, 20×20, 25×25, and— as an upper bound—50×50, while varying the number of users from $K = 4$ to $K = 10$ in steps of one. As shown in Fig. 5.17a, the

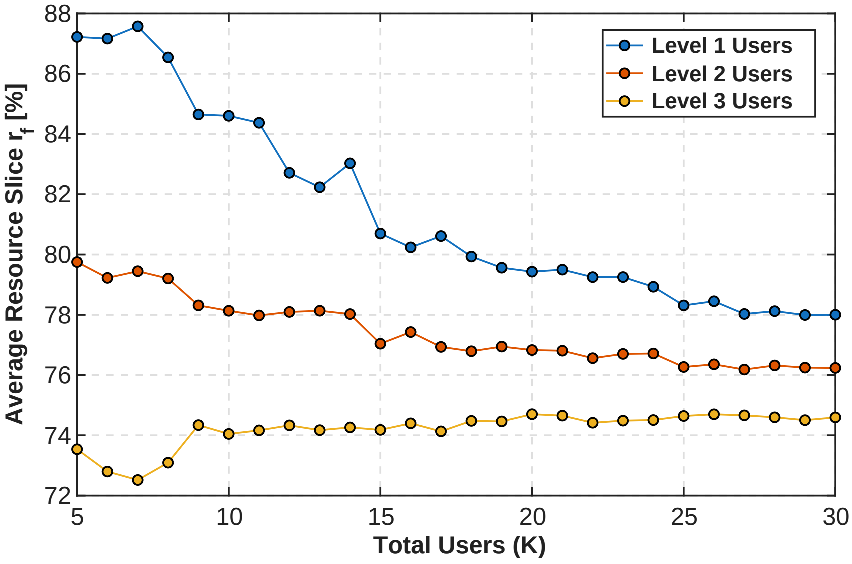

Figure 5.16: Resource slicing behavior per user pricing tier with respect to the total users with LEVEL usage.

runtime remains below 3 ms for RIS sizes up to 25×25, and below 5 ms even for the 50×50 case, confirming that LEVEL operates in the millisecond regime for realistic deployments.

To further assess whether the RIS dimension dominates the computational time of LEVEL, we repeat each configuration, with the same RIS size and the same number of users, over 50 Monte Carlo trials, and record the runtime distribution. The results are reported in Fig. 5.17b. For an RIS with $M \times N$ cells, LEVEL executes essentially the same per-element optimization procedure across all $M \times N$ elements; therefore, the overall runtime scales primarily with the RIS dimensions. By contrast, increasing the number of users K mainly increases the cost of the per-element objective evaluation (e.g., a weighted aggregation over K codebook entries), which remains a lightweight inner-loop operation. Over the considered K range, this introduces only a mild multiplicative factor that is largely masked by fixed solver overhead. Consequently, the RIS dimension determines the dominant workload, whereas the total number of users has a substantially weaker effect on the total computational time.

As a next step, we compare the multi-tier resource-allocation capability of LEVEL against COMMON and PRIME. Using emulations, we assess the codebook-entry multiplexing paradigm and quantify the performance of LEVEL relative to these established baselines. This experiment also evaluates the ability of LEVEL to orchestrate multiple RIS units under multi-user contention while preserving fair tier-based dif-

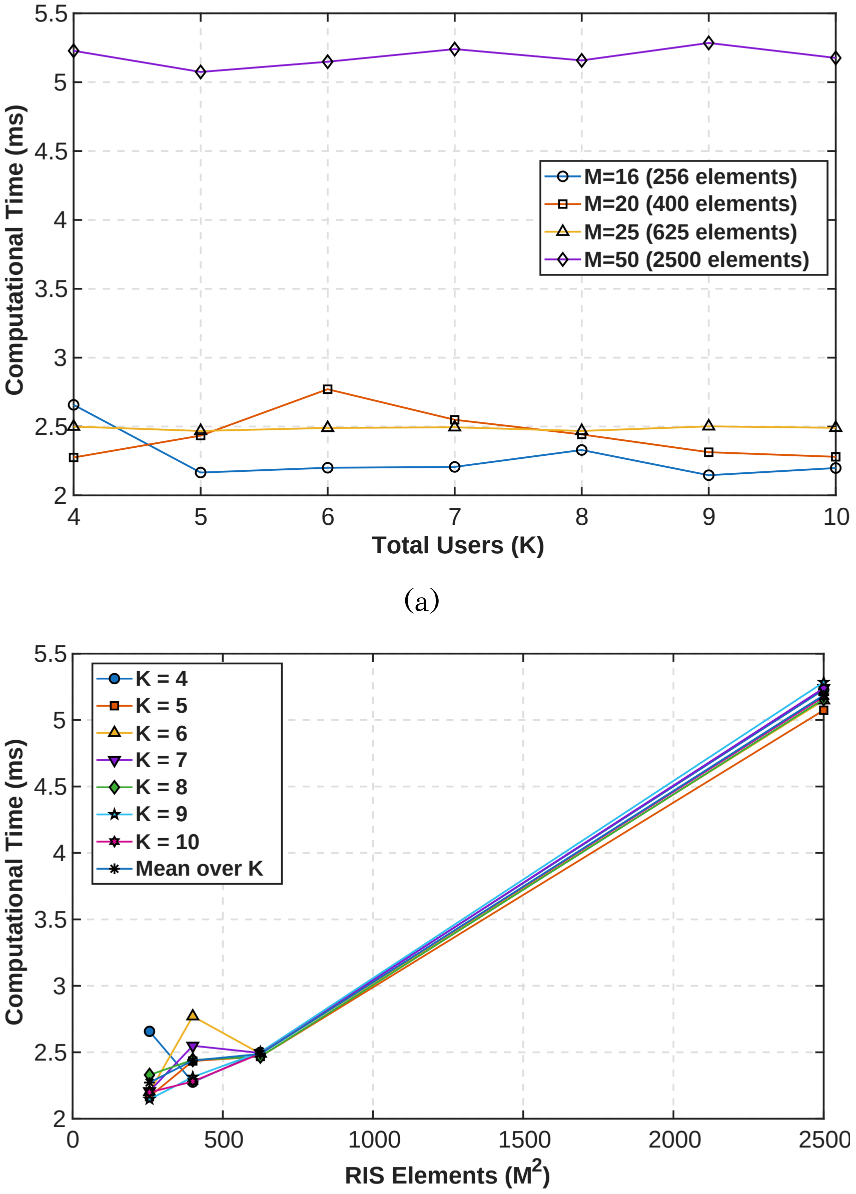

(a)

(b)

Figure 5.17: Computational time of LEVEL algorithm (a). Trend of the computational time with respect to RIS dimensions for different number of total users (b).

ferentiation. To account for practical hardware constraints, we consider both 1-bit and 2-bit RIS quantization. Moreover, during multi-user operation, users are assumed to be scheduled on orthogonal resource blocks (e.g., OFDMA) within the considered interval. Therefore, the reported SNR and SNR-loss are interference-free and capture only the degradation induced by imposing a single common RIS configuration across users—i.e., the deviation from each user's stand-alone optimum—rather than co-channel multiuser interference.

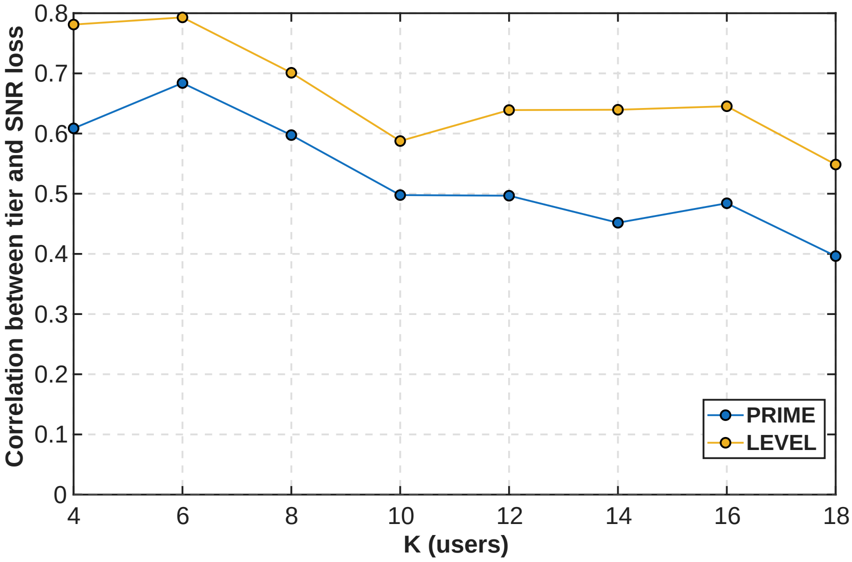

Figure 5.18: Correlation behavior between user tier and SNR loss using PRIME (blue) and LEVEL (yellow).

We first study the correlation between user tier and realized SNR loss as a fairness indicator, i.e., whether higher-tier users consistently experience smaller losses. We consider user populations ranging from $K = 4$ to $K = 18$. Users are randomly assigned either to LEVEL tiers or, for PRIME, to payment factors PF $\in \{1, 3, 5\}$. For each K, we run 50 Monte Carlo trials with random user selection and tier (or PF) assignment to ensure statistical validity. The cumulative results for both quantization options are summarized in Fig. 5.18. Overall, LEVEL yields a stronger correlation between tier and achieved efficiency, indicating more consistent tier-to-performance mapping. For up to 8 users, PRIME attains a mean correlation of 0.63, whereas LEVEL reaches 0.78. As the network becomes denser, LEVEL maintains correlations of 0.65 (for users from 10 to 14) and 0.60 (for users 16 to 18), while PRIME drops to 0.482 and 0.40, respectively. This behavior indicates that LEVEL transfers tier weights to the final RIS configuration more effectively, preserving a largely monotonic tier–performance relationship as contention increases, whereas PRIME becomes more prone to tier inversion when many users are jointly served.

We next examine the per-tier SNR loss itself, which directly quantifies multi-user degradation relative to the standalone (single-user) optimum. We again vary the total number of users from 4 to 18, with random tier assignments and 50 Monte Carlo trials for each total number of users K. The main trends for the 1-bit and 2-bit cases are shown in Fig.5.19a, with detailed tier-wise results in Fig.5.19b. The

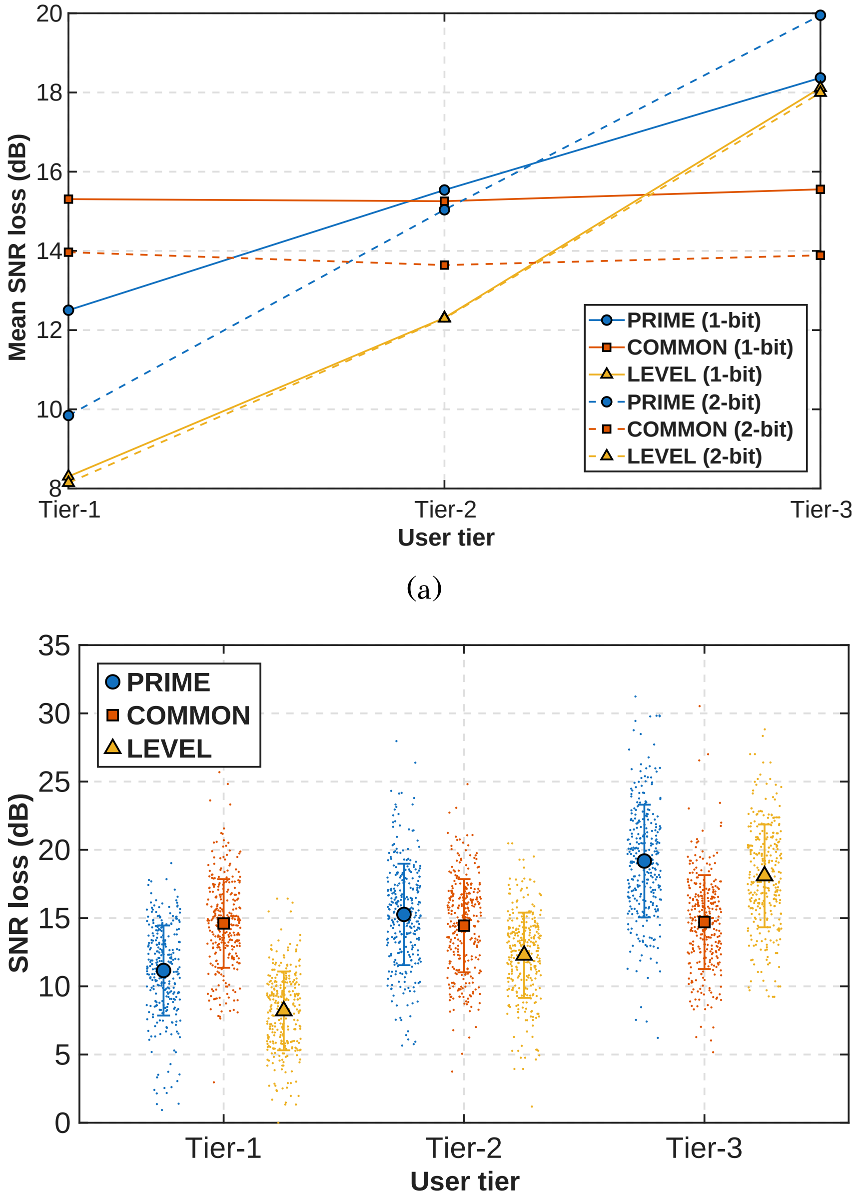

(a)

(b)

Figure 5.19: SNR loss behavior with 1-bit and 2-bit RIS for PRIME, COMMON & LEVEL (a). Detailed results for 50 Monte Carlo trials for PRIME, COMMON & LEVEL (b).

results confirm that codebook-entry multiplexing can support tier-aware resource allocation: both PRIME and LEVEL generally preserve the intended ordering across tiers, while COMMON is only appropriate when all users belong to the same tier (i.e., no differentiation is required). Beyond this, LEVEL consistently improves performance relative to PRIME across tiers and quantization resolutions. In the 1-bit case, tier-1 users under LEVEL and PRIME experience mean SNR losses of 8.3 dB and

$12.5\ dB$, respectively. For tiers 2 and 3, the mean losses are $12.3\ dB$ and $18.11\ dB$ for LEVEL, versus $15.53\ dB$ and $18.36\ dB$ for PRIME. Hence, LEVEL not only enforces tier differentiation more reliably but also reduces loss across all tiers compared to PRIME. The same trend holds for 2-bit quantization: the mean tier-wise losses for LEVEL are $7.5\ dB$, $10.9\ dB$, and $16.3\ dB$, while PRIME yields $8.8\ dB$, $14.1\ dB$, and $19.4\ dB$. As expected, the additional phase resolution of 2-bit RIS, mainly, improves both algorithms, yet LEVEL retains a clear advantage in both fairness and absolute performance.

5.5.3 Discussion

Codebook-entry multiplexing emerges as the key operating-phase primitive for multi-user RIS control because performance is governed primarily by the deviation between each user's stand-alone entry and the deployed common configuration. COMMON and PRIME exploit this link implicitly through entry aggregation, whereas LEVEL makes it explicit: it directly optimizes the per-user deviations Δ_{fk}, yielding controllable and predictable departures from the single-user optimum and thereby operationalizing RIS resource slicing under heterogeneous service tiers.

The evaluation confirms that this design choice translates into both improved tier-consistent fairness and improved absolute performance. Using emulations with Monte Carlo trials, LEVEL exhibits stronger and more stable tier–performance coupling than PRIME, as quantified by the higher correlation between user tier and achieved efficiency across increasing contention. In parallel, LEVEL reduces the per-tier SNR loss relative to PRIME under both 1-bit and 2-bit quantization, indicating that its advantage persists under realistic hardware constraints and does not rely on unrealistically fine phase control. These results support the core claim that explicitly controlling deviation—rather than relying on implicit aggregation—better preserves monotonic tier differentiation and mitigates tier inversions in dense multi-user operation.

A second practical outcome is computational viability. By solving independent scalar problems per RIS element, LEVEL avoids global combinatorial search and maintains millisecond-level runtime for realistic RIS sizes, while scaling primarily with the number of elements and only weakly with the number of users. This makes LEVEL compatible with time-sensitive PWE operation and, importantly, with orchestration over multiple RIS units where repeated reconfiguration is unavoidable.

Overall, LEVEL advances the operational maturity of pricing-aware RIS sharing: it provides an explicit mechanism to map operator policy (tier weights and deviation bounds) into predictable user-level performance, while remaining compatible with codebook-based operation, quantized hardware, and multi-RIS indoor propagation. Together with COMMON and PRIME, the results substantiate the thesis-level claim that multiplexing of precomputed entries can support concurrent service, pricing-based differentiation, and the composition of more complex EM responses from simpler building blocks, without incurring prohibitive online optimization overhead.

CHAPTER 6

A FRAMEWORK FOR FAIR RESOURCE ALLOCATION IN RIS-ENABLED NETWORKS

As noted in Chapter 1, the third challenge for RIS deployment is coordinating multiple RIS units to guarantee fairness across users. Modeling an RIS-enabled network as a graph is foundational, but existing graph-based methods ignore stakeholder pricing policies and do not enforce tiered fairness. These gaps motivate THEMIS[1]: a framework that integrates pricing policies and enforces fairness among multi-tier users.

The workflow (Fig. 6.1) builds on this thesis's logic: information generated during codebook compilation and RIS resource sharing via multiplexing of codebook entries. Each user issues a service request, specifies source and destination RIS units, is tagged by pricing tier, and shares location with THEMIS. Using these inputs, THEMIS queries the Codebook Database, computes the common configuration CC per unit, and determines a path that balances efficiency and fairness across tiers. Finally, it delivers

---

[1] In Greek mythology, Themis—daughter of Uranus and Gaea—is the Titaness of divine law, order, wisdom, and justice, often depicted as "Lady Justice."

configuration data to users and RIS units, yielding a unified, low-latency framework for B5G/6G deployment.

6.1 THEMIS Description

THEMIS workflow is illustrated in Alg. F.1 that is included in Appendix F. THEMIS encompasses two core components: (i) a fairness-aware methodology to determine the CC shared among multiple users with distinct pricing levels, as defined by the network operator, and (ii) an optimal, congestion-aware path selection mechanism for each user. Specifically, the framework initiates by defining the RIS-assisted network structure within a square area characterized by the parameter *areaSize*. RIS units, denoted by the node set $\mathcal{V}$, are strategically positioned in a two-dimensional Cartesian space according to operator-defined deployment strategies. Node connectivity decisions typically rely on LoS conditions but may also factor in criteria such as operational frequencies or RIS functionalities clustering. This process generates an adjacency matrix representing inter-node distances. Edges $\mathcal{E}$ are established only when the inter-node distance is below a specified threshold, $\texttt{LoS}_{\texttt{th}}$. Consequently, the resulting RIS network topology is formally represented as the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$.

The system serves K users, each with source`userSrc`(k) and destination `userDst`(k). THEMIS follows the PRIME workflow to form a common configuration CC at a shared RIS: each user is assigned a pricing factor v_k (integer voting weight); the user-optimal entries CE_k are fetched from the Codebook Database; and a stakeholder-defined efficiency floor η_{min} is in place. Each CE_k is then discretized (when continuous values appear) per stakeholder preferences, and every discretized entry is replicated proportionally to v_k. As in PRIME, this replication amplifies higher-paying users' influence while preserving pricing fairness. The resulting CC is obtained by the elementwise mode over the replicated entries:

$$CC(m,n) = \text{mode}\left(\begin{array}{c} \underbrace{CE_1(m,n), \ldots, CE_1(m,n),}_{v_1 \text{ times}} \\ \vdots \\ \underbrace{CE_K(m,n), \ldots, CE_K(m,n)}_{v_K \text{ times}} \end{array} \right) \tag{6.1}$$

Subsequently, the efficiency achieved for each user k under the allocated common

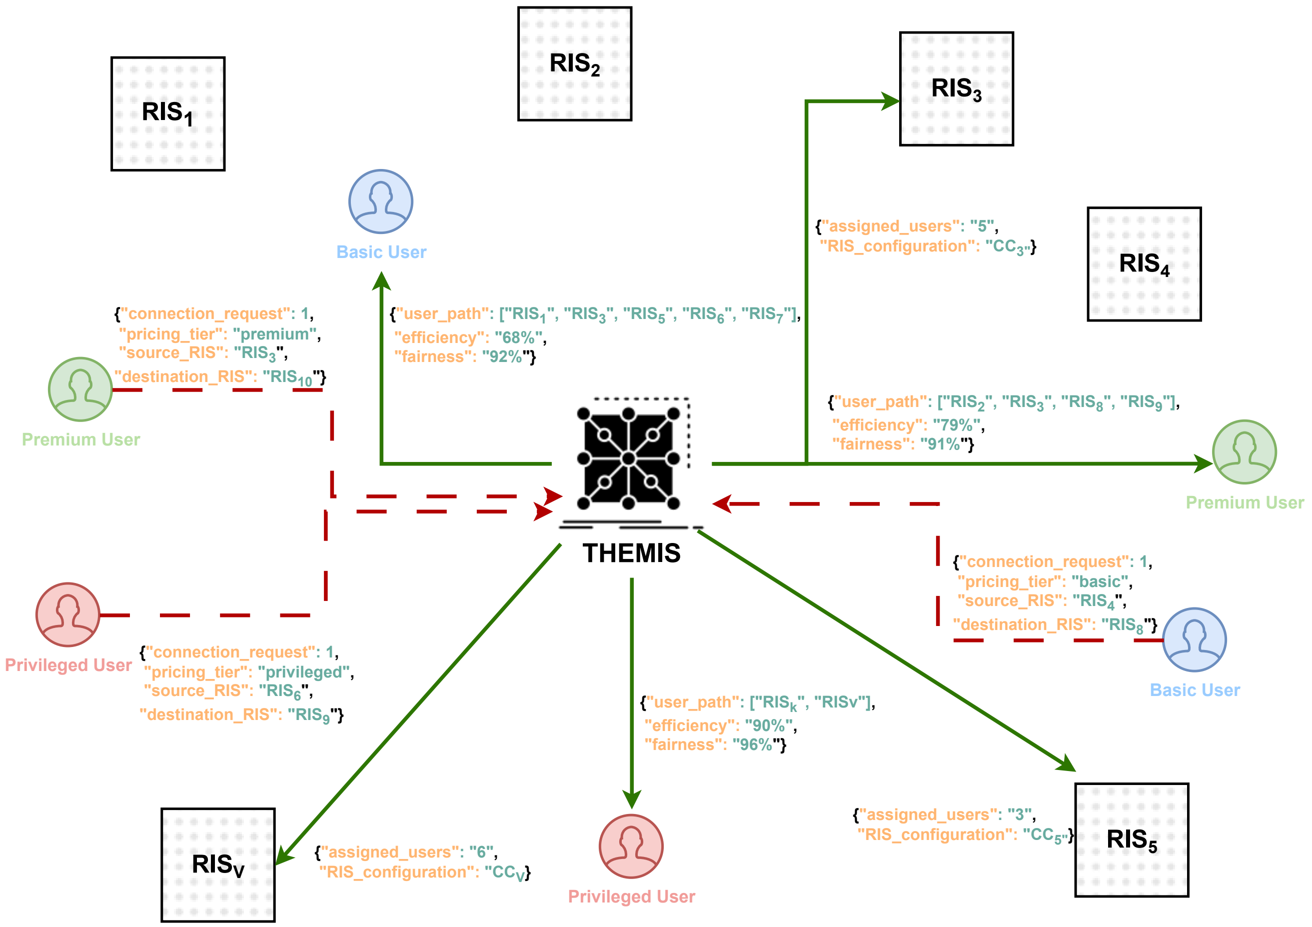

Figure 6.1: The workflow of the novel framework for fair Resource Allocation in RIS-enabled networks. THEMIS composses the Codebook Database and the pricing policy of the network's stakeholder. Receiving connection request, computes the optimal path per user and the CC for each RIS unit.

RIS configuration CC is estimated by comparing it against their ideal configuration CE_k:

$$\frac{\epsilon'_k}{\epsilon_k} = \left(1 - \frac{1}{MN}\sum_{m=1}^{M}\sum_{n=1}^{N}|CC(m,n) - CE_k(m,n)|\right) \times 100 \tag{6.2}$$

The key difference between THEMIS and PRIME in forming CC is an added fairness component that enforces a minimum performance for all users. Specifically, if a user's efficiency $\frac{\epsilon'_k}{\epsilon_k}$ falls below the threshold ηmin, a Nash–bargaining-inspired adjustment is triggered. This mechanism evaluates per-user fairness F_k and a unified network metric UFM, defined as:

$$F_k = \frac{\epsilon'_k/\epsilon_k}{\eta_k^{\text{exp}}/100}, \quad UFM = \frac{1}{K}\sum_{k=1}^{K} F_k \tag{6.3}$$

where F_k is averaged over all RIS units along k user path and the expected efficiency η_k^{exp} for each user is linearly scaled according to their voting factor v_k:

$$\eta_k^{\text{exp}} = \eta_{\text{min}} + \frac{v_k - v_{min}}{v_{max} - v_{min}}(100 - \eta_{\text{min}}) \tag{6.4}$$

where v_{min} and v_{max} are the lowest and the highest, respectively, pricing levels as defined by stakeholder's policy. The Nash-inspired fairness algorithm partitions the RIS panel into candidate sub-panels (`candG`) and iteratively updates the voting weights as follows:

$$\begin{aligned} w_{candG,k}^{(t+1)} &= (1-\alpha)w_{candG,k}^{(t)} \\ &\quad + \alpha\frac{w_{candG,k}^{(t)}}{\sum_{j=1}^{K} w_{candG,j}^{(t)}}\left(\sum_{j=1}^{K} v_j\right) \end{aligned} \tag{6.5}$$

ensuring each user's efficiency surpasses the threshold while maximizing the UFM. Here, $\alpha \in [0,1]$ controls convergence dynamics; a higher α leads to rapid but potentially unstable convergence, whereas lower α ensures smoother, more stable convergence.

For optimal user-specific path selection between nodes `userSrc`(k) and `userDst`(k), THEMIS employs an enhanced Dijkstra algorithm [122]. The path-selection criteria encompass: (i) inter-node physical distances, (ii) current RIS-edge congestion, and (iii) individual user pricing levels, as illustrated in Fig. 6.2. The d_{ij} expresses the

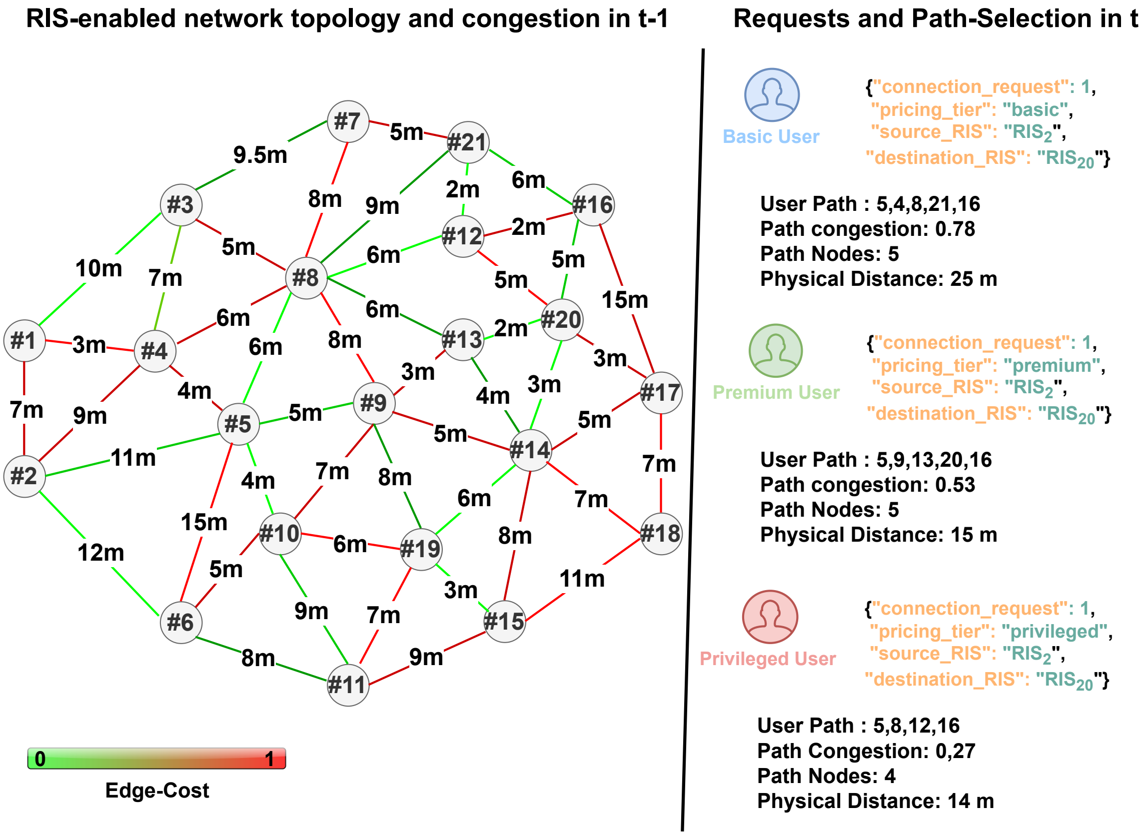

{"connection_request": 1,
"pricing_tier": "basic",
"source_RIS": "RIS$_2$",
"destination_RIS": "RIS$_{20}$"}

User Path : 5,4,8,21,16
Path congestion: 0.78
Path Nodes: 5
Physical Distance: 25 m

{"connection_request": 1,
"pricing_tier": "premium",
"source_RIS": "RIS$_2$",
"destination_RIS": "RIS$_{20}$"}

User Path : 5,9,13,20,16
Path congestion: 0.53
Path Nodes: 5
Physical Distance: 15 m

{"connection_request": 1,
"pricing_tier": "privileged",
"source_RIS": "RIS$_2$",
"destination_RIS": "RIS$_{20}$"}

User Path : 5,8,12,16
Path Congestion: 0,27
Path Nodes: 4
Physical Distance: 14 m

Figure 6.2: Mechanism of THEMIS for path selection in respect of fair Resource Allocation among different pricing-tiers users.

baseline hop cost and it is computed as the Euclidean distance between RIS i and j (Eq. (6.6)) capturing the free-space path-loss penalty of one cascaded reflection.

$$d_{ij} = \begin{cases} \|\mathbf{x}_i - \mathbf{x}_j\|^2, & \text{if } (i,j) \in \mathcal{E}, \\ \infty, & \text{otherwise.} \end{cases} \tag{6.6}$$

The per–edge cost used in the congestion-pricing Dijkstra is

$$c_{ij}^{(t)}(k) \;=\; d_{ij}\big(1 + \rho\, \frac{f_{ij}^{(t-1)}}{C}\big)\frac{v_k}{v_{\min}} \tag{6.7}$$

where $f_{ij}^{(t-1)}$ denotes the instantaneous edge load and is the number of user routes that already employ edge (i,j) at the previous iteration. $C = \lceil K/|\mathcal{V}|\rceil$ is the uniform capacity proxy and is computed as the average number of flows an edge could sustain if traffic were evenly spread across the $|\mathcal{V}|$ RIS nodes. A smaller C makes the congestion term grow sooner, forcing paths to diverge; a larger C tolerates more flows before detouring. Parameter $\rho \in (0,1]$ is the congestion-strength coefficient and is scaled the impact of $\frac{f_{ij}}{C}$. Higher ρ penalizes busy edges aggressively (better load balancing, longer routes); lower ρ keeps routes short but risks bottlenecks. Finally, the $\frac{v_k}{v_{\min}}$ expresses the pricing priority among the multiple-tiers users ensuring the fairness within the network also in the path selection procedure. Premium users see their entire path cost down-scaled, allowing them to keep shorter or less-congested cascaded paths, whereas lower-tier users are steered to cheaper edges. The optimal path for user k at iteration t is obtained by

$$\pi_k^{(t)} = \arg\min_{\pi \in \mathcal{P}(k)} \sum_{(i,j)\in\pi} c_{ij}^{(t)}(k) \tag{6.8}$$

and the edge costs are updated iteratively until $\max_{(i,j)} |c_{ij}^{(t+1)} - c_{ij}^{(t)}| < \varepsilon$ or the iteration limit `maxRouteIter` is reached.

The computational complexity of THEMIS is primarily governed by two core stages. The first is the congestion-pricing routing phase. This phase incurs a complexity of $\mathcal{O}(\texttt{maxRouteIter}K|\mathcal{E}|\log|\mathcal{V}|)$. The second stage is the computation of the CC at the RIS node level. When the Nash bargaining mechanism requires I_{N} iterations, the associated complexity is $\mathcal{O}(I_{\mathrm{N}}KMN)$. Therefore, the total computational complexity of THEMIS scales as $\mathcal{O}(\texttt{maxRouteIter}K|\mathcal{E}|\log|\mathcal{V}| + I_{\mathrm{N}}KMN)$. This indicates that THEMIS operates with polynomial complexity, rendering it suitable for deployment in practical RIS-enabled wireless networks.

6.2 Evaluation of THEMIS

In this section, we evaluate the THEMIS framework. We begin assessing CC computation mechanism efficiency and fairness across various RIS sizes under tests with binary (0/1) configuration—arguably the most challenging case for multiplexing of RIS elements' values. We then validate this mechanism on the testbed. Finally, we assess the complete THEMIS workflow, again with binary RIS configurations, across scenarios with differing RIS unit counts and user densities.

In order to demonstrate the fair definition of the CC that THEMIS proposes, we consider 10 users simultaneously utilizing a single 16×16 RIS unit, each with distinct pricing levels and corresponding voting weights. The efficiency threshold ($\eta_{\min}$) is set to 55%. The fairness mechanism is configured to allow a maximum of 300 iterations, with the fairness adaptation parameter (α) tuned to 0.7. The evaluation outcomes are

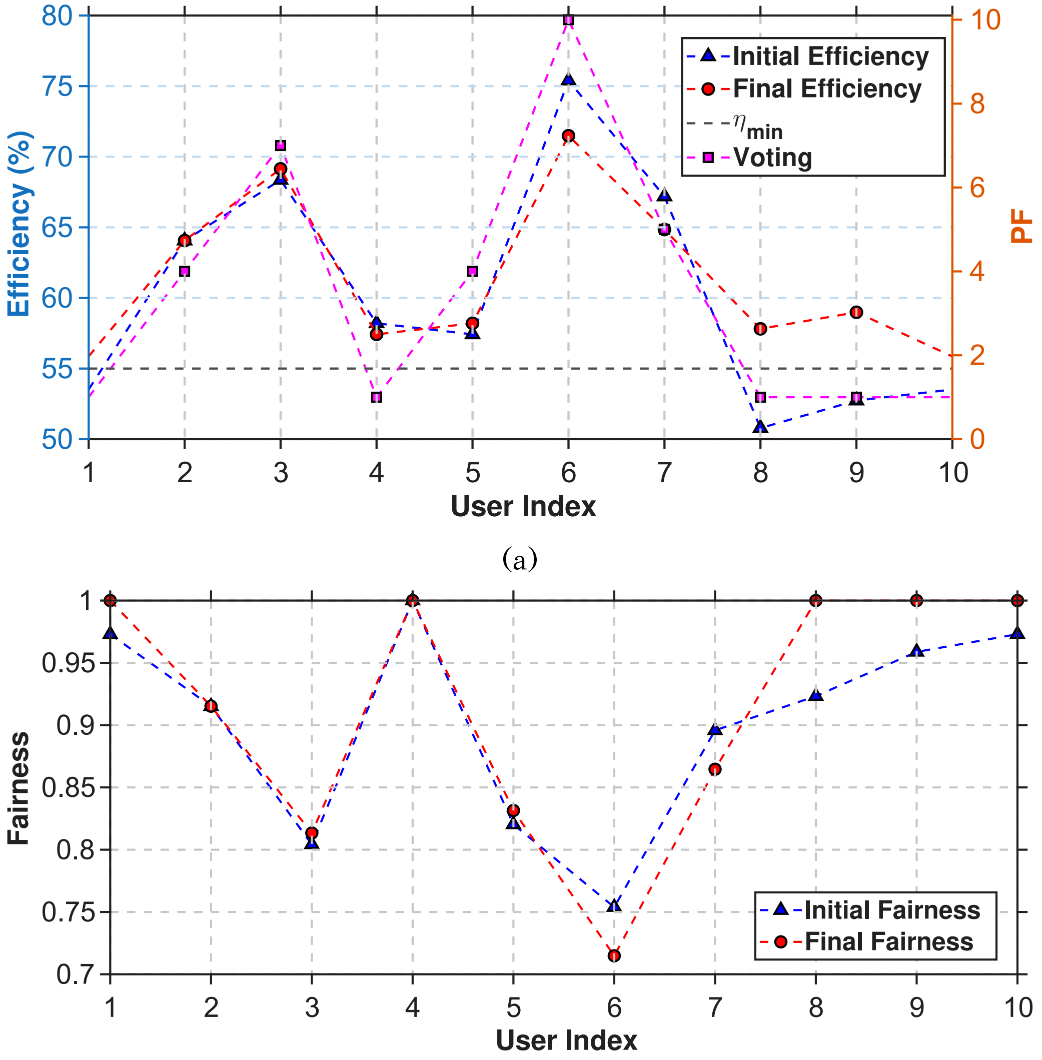

(a)

(b)

Figure 6.3: Efficiency (up) and fairness (down) with and without the Nash-inspired mechanism.

shown in Fig. 6.3a. We compare the achieved efficiency with and without employing the Nash-inspired fairness mechanism. For this mechanism, we test various `candG` subgroup sizes, specifically $N/2$, $N/3$, $N/4$, and $N/5$, to determine the option providing the highest fairness. It is evident from the results that the fairness mechanism guarantees that each user's efficiency, as calculated by Eq. (6.2), meets or exceeds the defined threshold (η_{min}). Furthermore, the efficiency achieved by each user remains proportional to their respective voting factors. This observation is reinforced by the increase in the overall UFM of the network from 0.88 to 0.96, as illustrated by the per-user fairness metrics shown in Fig. 6.3b.

Keeping the same parameters for CC determination as before, we further analyze

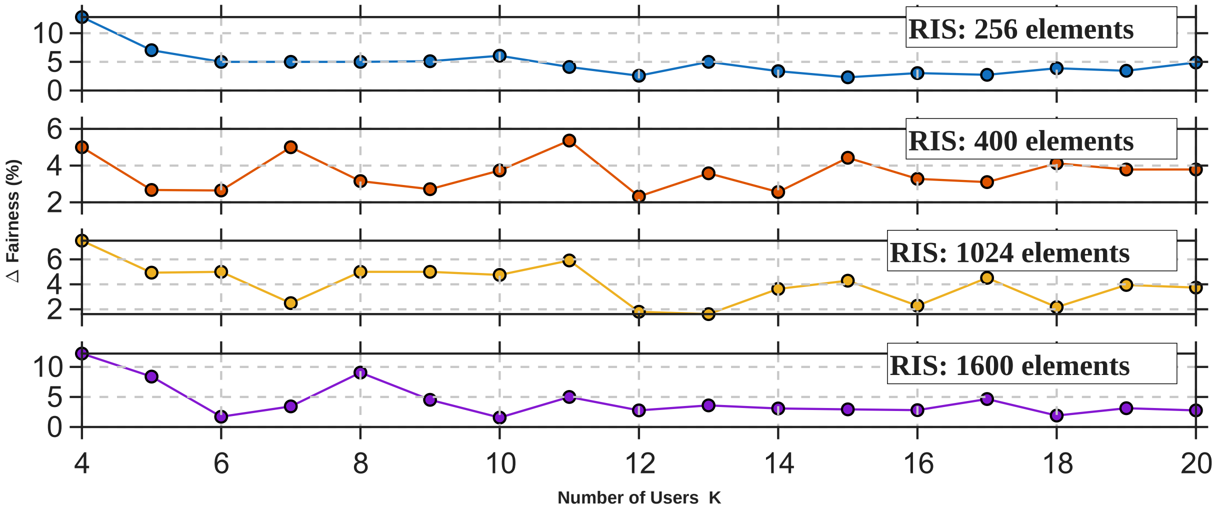

Figure 6.4: Fairness enhancement Δ across different RIS dimensions and number of users K.

the fairness improvements across different RIS dimensions and varying user counts. Results summarized from 100 simulation runs per scenario are depicted in Fig. 6.4. The greatest fairness enhancements (approximately 9–13%) occur independently with RIS size once fewer users exist, attributed to the increased flexibility in resource allocation. However, notable fairness gains (up to 6%) are also achieved in scenarios involving higher user densities even in smaller RIS dimensions, underscoring the robustness of THEMIS across diverse operational contexts.

The fairness mechanism is evaluated on the real testbed of Section 3.3. We consider one transmitter serving four users via a single RIS unit. Using GNU Radio, we extract I,Q samples in standalone and multi-user cases and, for each user, compute G_k^{multi} against G_k^{std}, exactly as in Section 5.4.2. The results are illustrated in Fig. 6.5. Users are identified by ID, with respective voting factors set to 1, 5, 10, and 20, and a minimum efficiency threshold of 55%. The initial standalone user gains are 6.9 dB, 6.97 dB, 6.00 dB, and 6.16 dB. After applying THEMIS for the CC computations, the gains are 3.2 dB, 6.08 dB, 5.5 dB, and 6.15 dB respectively. These gains result in resource fairness factors of 46.37%, 87.23%, 91.66%, and 99.83%. Clearly, the codebook multiplexing component of THEMIS effectively determines CC allocations for fair beam-splitting and resource allocation, with user performance proportional to the assigned voting factor.

For the comprehensive evaluation of THEMIS, we deploy a network within a square physical area measuring $areaSize = 50$ m per side. An edge between two

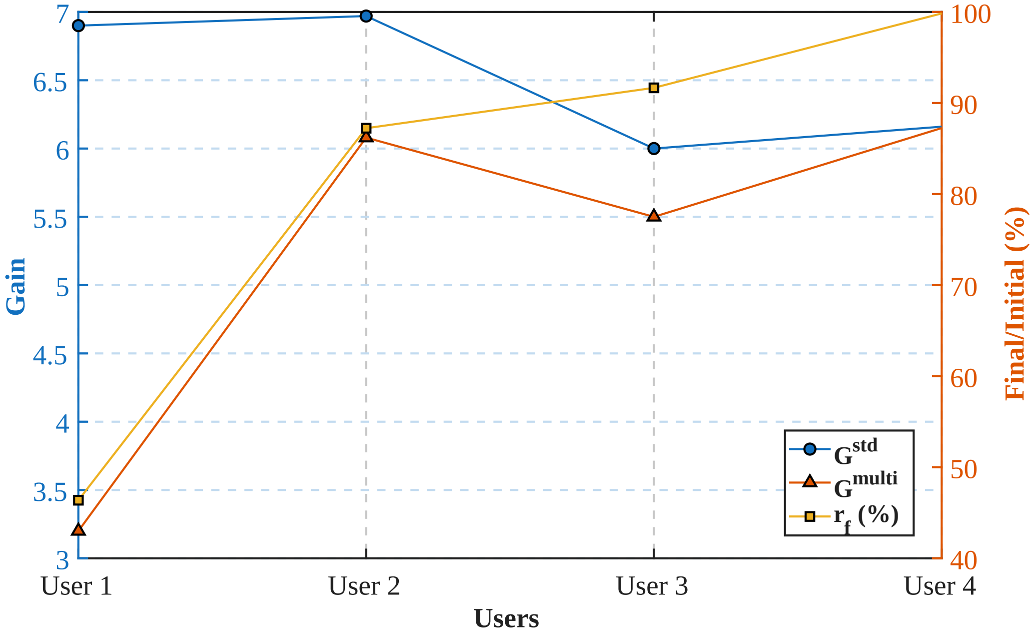

Figure 6.5: THEMIS evaluation results from measurements in RIS testbed.

RIS units is established only when their separation distance is less than or equal to the threshold $\texttt{LoS}_{\texttt{th}} = 15$ m. Each RIS operates in a strictly binary configuration to represent the most demanding scenario, consistent with prior experiments, and each RIS has fixed dimensions of 25×25. All parameters governing the definition of CC remain identical to those used previously to ensure methodological consistency.

To demonstrate THEMIS operation, we instantiate $\mathcal{V} = 10$ RIS nodes serving $K = 20$ users. The congestion-strength parameter is set to $\rho = 0.1$, and the algorithm termination limit is defined as $\texttt{maxRouteIter} = 300$. The resulting cascaded-RIS network topology is illustrated in Fig. 6.6, clearly depicting the coordinates of each RIS node, along with the node pairs that satisfy Eq. (6.6) and thus form valid edges. The figure also highlights the optimal path selected for each $\texttt{userSrc}(k) \rightarrow \texttt{userDst}(k)$ pair. The collective outcomes of the path-selection procedure are summarized in Fig. 6.7. Beyond fairly allocating RIS resources in accordance with the pricing levels of traversing users, THEMIS ensures that higher-level users are routed via shorter paths, while lower-level users are directed through less congested nodes, simultaneously achieving efficiency and fairness. Specifically, premium users' average path length is 20.4 m whereas the respective of the basic users is 35.14 m. The overall network fairness across all users and RIS nodes, reaches approximately 91.1%.

To evaluate THEMIS performance on cascaded-RIS links, we vary the number of RIS nodes from 10 to 100 and the number of users from 20 to 500. Each sce-

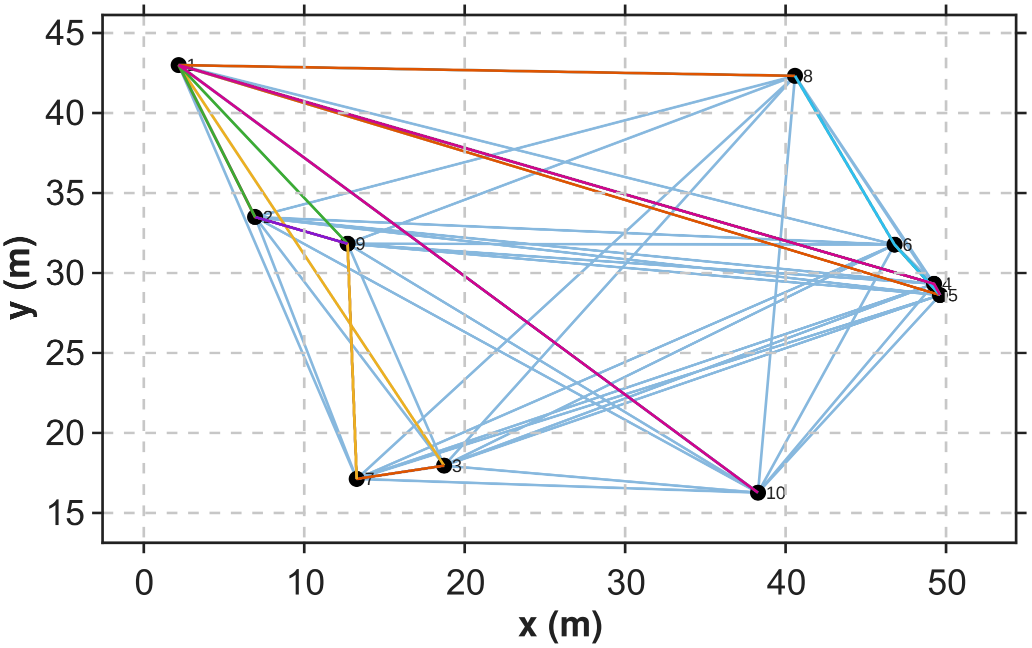

Figure 6.6: Graph-based illustration of RIS-enabled network as defined by THEMIS.

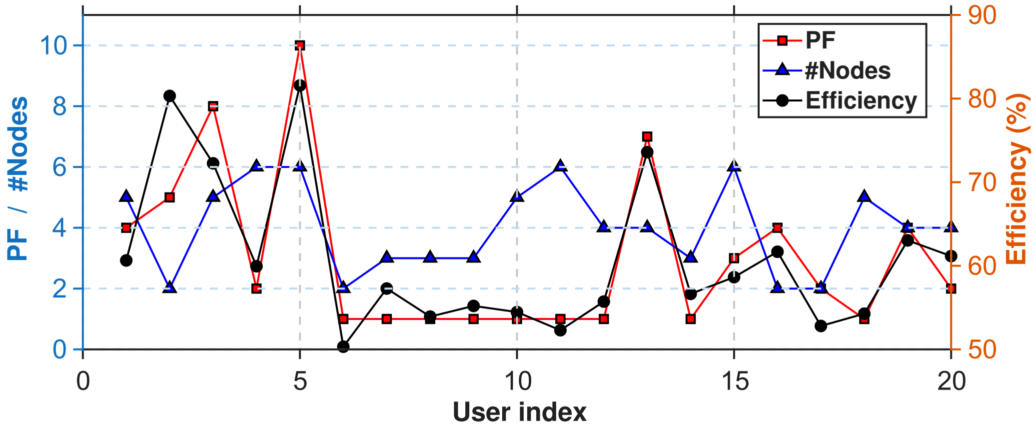

Figure 6.7: Efficiency and nodes per user in RIS-enabled network.

nario is tested for congestion-strength parameters $\rho = 0.1$ and $\rho = 0.3$, and algorithm termination limits (`maxRouteIter`) were set at 100, 300, and 500 iterations. Network fairness and computational overhead are measured over 100 independent simulation runs per scenario, and aggregated statistics are reported. Specifically, Fig. 6.8 illustrates network fairness as a function of the user-to-RIS ratio, displaying only the best-performing curves for each ρ. When the user-to-RIS ratio is below five, fairness surpasses 90% at $\rho = 0.1$. Increasing ρ offers improvements for certain user-to-RIS ratios but leads to an overall slight reduction in fairness (mean fairness decreases from 0.87 to 0.86). This highlights the importance of maintaining low congestion and careful parameter tuning. Nevertheless, even in high-density scenarios (user-to-RIS

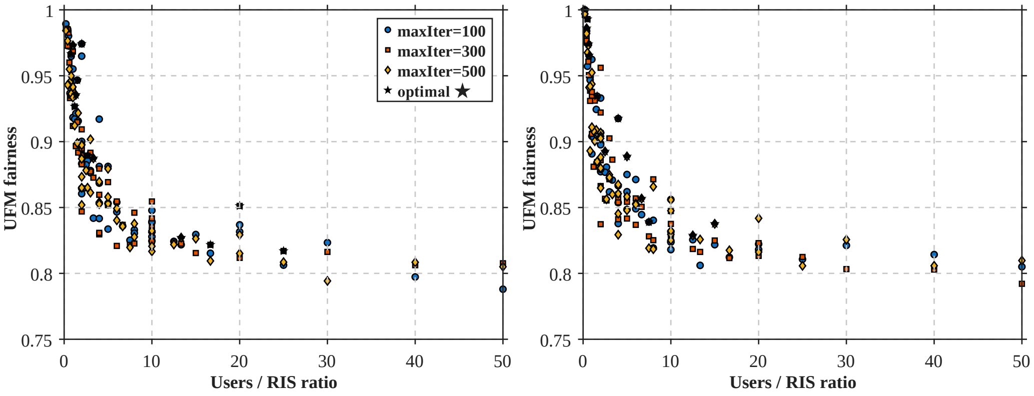

Figure 6.8: Achieved network fairness for $\rho = 0.1$ (left) and $\rho = 0.3$ (right) in different number of users and RIS nodes expressed as ratio.

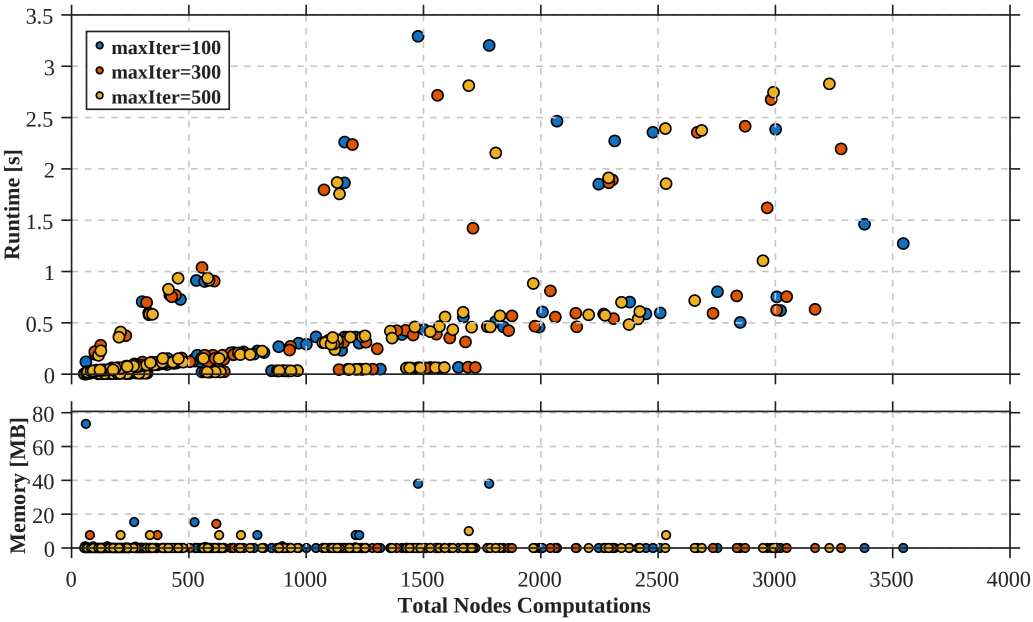

Figure 6.9: Computational Runtime (up) and CPU memory usage (down) of THEMIS usage.

ratios ranging from 20 to 50), fairness remains consistently above 80%, demonstrating THEMIS's robustness in densely populated networks.

Fig. 6.9 presents THEMIS' computational overhead when computing one path per user. Runtime increases linearly with the number of computations per node (including RIS dimensions). For user-to-RIS ratios below 10, runtime remains lower than 0.5 s, increasing proportionally with higher ratios. Even under the most demanding scenario—calculating 500 paths and performing over 4,500 CC evaluations—runtime never exceeds 2.5 s. CPU memory consumption is modest, peaking at approximately 12 MB, confirming THEMIS' suitability for real-time applications. Finally, Fig. 6.10

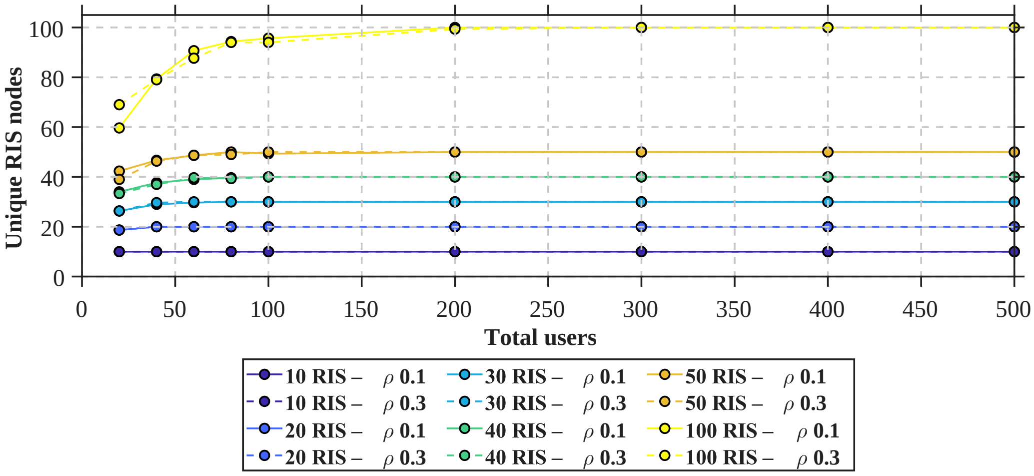

Figure 6.10: Usage of RIS nodes in different scenario based on PathFinding.

reports the number of RIS nodes activated for each network scenario. Under both ρ conditions, nearly all available RIS nodes become active as network capacity increases, confirming that THEMIS efficiently exploits all the available resources. High-priority users consistently receive shorter paths, while specific RIS nodes are strategically allocated to lower-priority users, thus balancing fairness and baseline efficiency across the network.

6.3 Discussion

THEMIS consolidates the thesis pipeline into a single, deployable framework for RIS-enabled B5G/6G networks, with an explicit focus on low-latency operation. It couples (i) codebook-based control, (ii) multi-user resource allocation via multiplexing of codebook entries with pricing-aware fairness guarantees, and (iii) congestion-aware path selection over cascaded RIS graphs. In doing so, THEMIS addresses a key limitation of prior graph-based RIS networking approaches: tiered fairness is not treated as an afterthought but is embedded both locally (per shared RIS node through the CC computation with an efficiency floor) and globally (through pricing-aware congestion-pricing routing).

The evaluation supports the practicality of this design. At the RIS-node level, the Nash-inspired adjustment enforces the minimum-efficiency constraint while preserving proportional differentiation across voting factors, and the testbed results confirm that configuration-deviation-based efficiency estimation remains a useful proxy for real performance under multiplexing. At the network level, THEMIS maintains high fairness across a wide range of user-to-RIS ratios, while routing premium users through shorter paths and steering lower-tier users away from congested resources, achieving a concrete fairness–efficiency trade-off through the congestion parameter ρ. The measured runtime and memory overheads remain modest even in dense scenarios, indicating that the framework can scale to realistic network sizes without requiring per-element online optimization.

Beyond operational use, THEMIS provides a system-level tool for deployment planning: its graph formulation and explicit congestion/fairness mechanisms can be used to study RIS placement, LoS thresholds, and topology design under stakeholder policy constraints. Overall, THEMIS demonstrates that pricing-aware fairness and cascaded-RIS path selection can be integrated coherently within the codebook paradigm, yielding an architecture-consistent route to scalable, multi-RIS programmable wireless environments.

CHAPTER 7

RF FENCING: A NOVEL RIS-BASED SERVICE FOR PROACTIVE COVERT COMMUNICATIONS

Codebook multiplexing, apart from fair resource allocation, can also be used in other missions within RIS-assisted PWEs, such as physical layer security. As illustrated in Fig. 7.1, conventional wireless networks let signals propagate according to natural EM phenomena rather than network intent, exposing them to jammers and eavesdroppers. To counter this, communication environments must precisely control wave propagation and shape trajectories to meet operational objectives. By jointly steering energy toward intended receivers and suppressing radiation toward potential adversaries, PWEs improve communication performance while limiting external exposure, enabling covert and private services through spatially selective control of signal propagation.

Building on this foundation, RF-fencing enforces spatially selective suppression and delivery via coordinated control of multiple RISs. It leverages RIS multi-functionality to receive radiation from diverse directions and flexibly reroute or cancel it per network objectives. To support such versatility, RF-fencing multiplexes precompiled codebook entries, dynamically combining simple configurations to synthesize complex

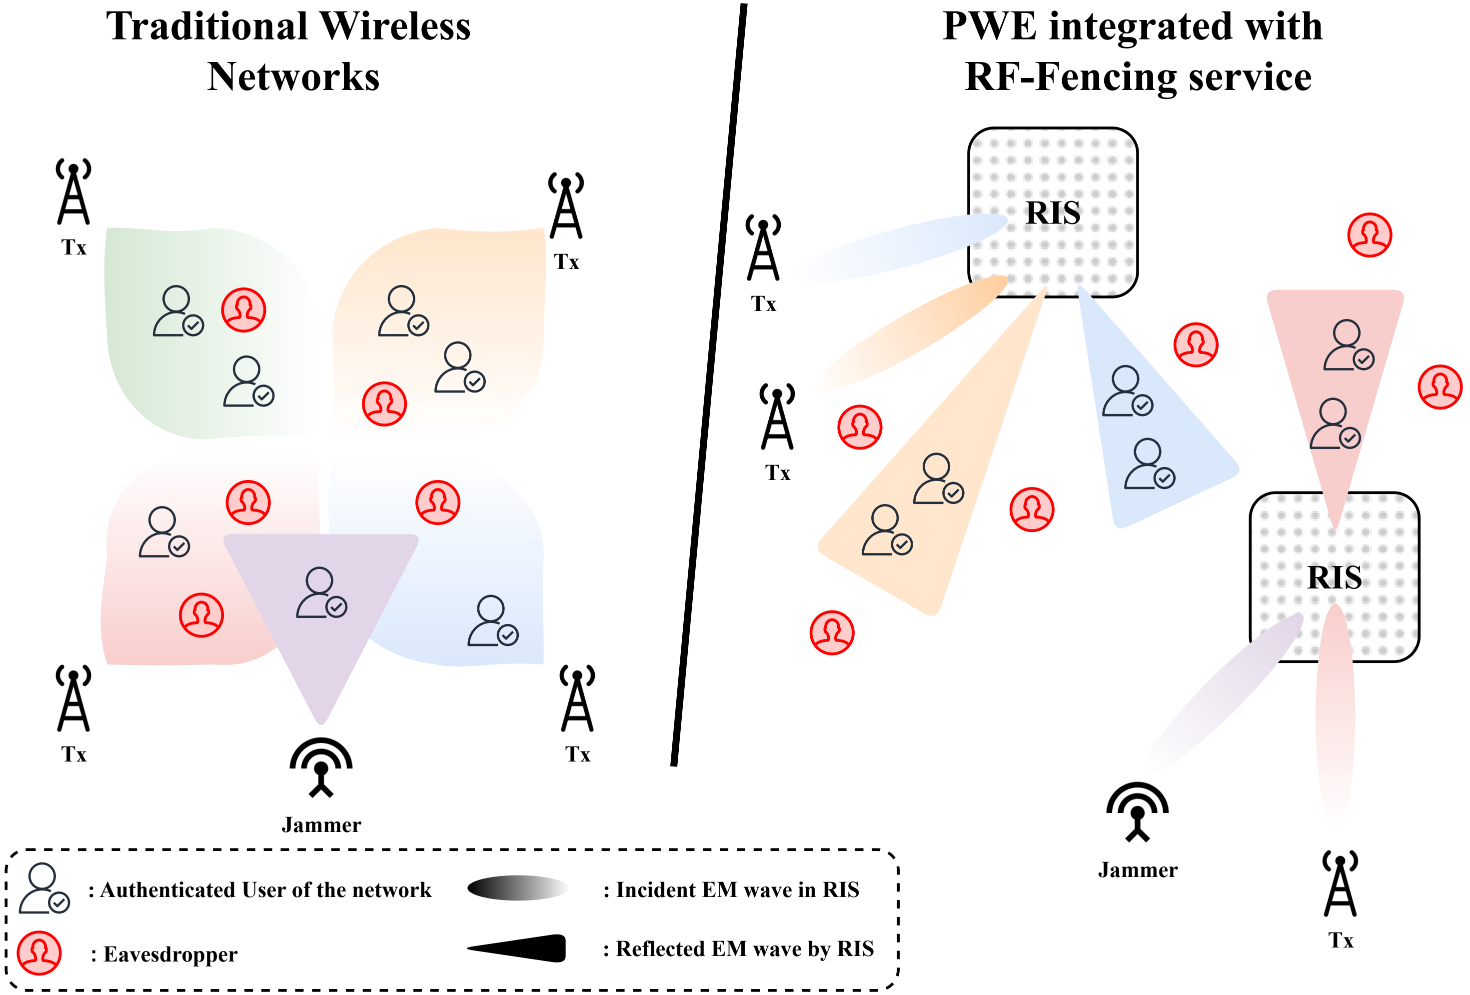

Figure 7.1: Comparison between traditional networks and PWEs in respect of security and covertness.

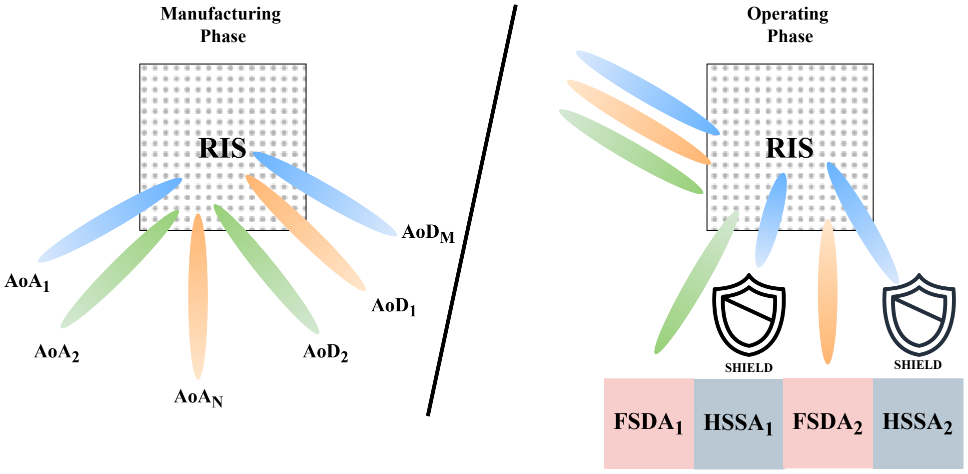

Figure 7.2: Manufacturing (left) and Operating (right) phase using RF-Fencing SHIELD algorithm.

propagation responses and avoid iterative optimization. Because effectiveness hinges on the accuracy of these entries, they must capture RIS physical behavior—especially for spatially informed services like covert communications—hence the codebook must derive from EM-consistent models rather than statistical channel abstractions lacking the needed spatial granularity. Orchestrating multiple RISs requires a codebook-based controller that translates precomputed entries into real-time actions. *SHIELD* is introduced as the first dedicated algorithm for practical RF-fencing within PWEs.

7.1 SHIELD Description

In this section, we present SHIELD algorithm. Its overall workflow—codebook retrieval, field multiplexing, and online phase refinement—appears in Alg.G.1 in Appendix G. SHIELD partitions the network domain into proactively selected hostile signal suppression areas (HSSAs) and friendly signal delivery areas (FSDAs), while also creating fully confined quiet zones dynamically synthesizing the desired EM response by multiplexing precompiled RIS configurations with lightweight optimization. Specifically, as shown in Fig.7.2, SHIELD leverages manufacturing-time codebooks specifying RIS configurations for predefined beam-steering scenarios (AoA/AoD). At runtime, it retrieves and multiplexes the entries, then applies a lightweight refinement to synthesize the configuration matrix $\mathbf{\Phi}$ that suppresses undesired directions while preserving coverage in intended zones, thus avoiding heavy iterative optimization and full channel estimation. Notably, when the jammer's AoA is known, SHIELD mitigates jamming by mapping the legitimate-user region to HSSAs, effectively canceling interference.

The operation of SHIELD begins with the identification of angular regions that correspond to both service provisioning and signal suppression objectives, leveraging precomputed EM field responses E_{field} obtained from the RIS codebook. To this end, the algorithm takes as input the desired number of FSDAs and HSSAs, denoted by d and u, respectively. For each of them, SHIELD analyzes the associated field distribution $|E_{\text{field}}|$ across the angular domain in order to detect dominant scattering components. In this process, angular regions where the field magnitude exceeds the predefined thresholds τ_{FSDA} for FSDAs or τ_{HSSA} for HSSAs are systematically aggregated into direction-specific masks, which can mathematically be expressed as

$$
\begin{aligned}
\mathcal{M}_d &= \bigcup_{k=1}^{d} \{\, \theta, \varphi \;:\; |E_k| > \tau_{\text{FSDA}} \max |E_d(\theta, \varphi)| \} \\
\mathcal{M}_u &= \bigcup_{k=d+1}^{d+u} \{\, \theta, \varphi \;:\; |E_k| > \tau_{\text{HSSA}} \max |E_d(\theta, \varphi)| \}
\end{aligned}
\tag{7.1}
$$

where $\mathcal{M}_d$ and $\mathcal{M}_u$ represent the sets of angular positions associated with the FSDA and HSSA regions, respectively, and $|E_k|$ corresponds to the field magnitude retrieved from the k-th codebook entry. These masks effectively localize the angular sectors relevant for service provisioning and signal suppression.

Building upon the constructed masks, SHIELD proceeds to synthesize a composite EM response $\boldsymbol{E}_{\text{common}}$ that reinforces the dominant fields within the FSDA regions while simultaneously suppressing emissions toward the identified HSSAs. To achieve this, SHIELD computes the maximum of the FSDA field components in areas where there is no overlap with the HSSAs, while in overlapping regions, it applies a compromise factor η to balance the conflicting objectives of suppression and service preservation. Outside these designated regions, the field components are deliberately nulled to minimize unintended radiation. Thus, the composite field can be expressed as

$$\boldsymbol{E}_{\text{common}} = \begin{cases} \max\limits_{1 \leq k \leq d} E_k & \text{in } \mathcal{M}_d \setminus \mathcal{M}_u, \\ \eta \max\limits_{1 \leq k \leq d} E_k & \text{in } \mathcal{M}_d \cap \mathcal{M}_u, \\ 0 & \text{otherwise,} \end{cases} \tag{7.2}$$

where η serves to control the trade-off between maintaining service quality and enhancing suppression in overlapping regions. Following this, the initial phase configuration for all RIS elements, denoted by $\boldsymbol{\Phi}_{\text{common}}$, is extracted from the argument of the composite field, as given by

$$\boldsymbol{\Phi}_{\text{common}} = \arg(\boldsymbol{E}_{\text{common}}), \tag{7.3}$$

where the operation $\arg(\cdot)$ returns the phase angle of the complex field values. This step provides an efficient near-optimal estimation of the $\boldsymbol{\Phi}_{\text{common}}$, similar to methods used in smart antenna arrays and analog beamforming for both far-field and near-field cases [120]. Consequently, a lightweight gradient-based procedure is used in order to fine-tune the initial guess of $\boldsymbol{\Phi}_{\text{common}}$ and correct minor deviations [121]. By this approach, SHIELD workflow remains applicable for real-time RIS control offering the minimal computational overhead. Specifically for the optimization stage, SHIELD employs a cost function that quantifies the discrepancy between the synthesized EM field and the desired target amplitudes within the angular masks $\mathcal{M}_d$ and $\mathcal{M}_u$, corresponding to the FSDAs and HSSAs, respectively. For each region type, this mismatch is computed as

$$\mathcal{J}_{d/u}(\boldsymbol{\Phi}) = \sum_{(\theta,\varphi) \in \mathcal{M}_{d/u}} \left| E(\boldsymbol{\Phi}; \theta, \varphi) - E_{\text{POI},d/u} \right|^2, \tag{7.4}$$

where $E(\boldsymbol{\Phi}; \theta, \varphi)$ represents the EM field evaluated under the current RIS configuration $\boldsymbol{\Phi}$, and $E_{\text{POI},d/u}$ denotes the target field amplitude within the FSDAs and HSSAs, respectively. These region-specific costs are then combined into a unified objective function, formulated as

$$\mathcal{J}(\boldsymbol{\Phi}) = \mathcal{J}_d(\boldsymbol{\Phi}) - w_{\text{opt}} \mathcal{J}_u(\boldsymbol{\Phi}), \tag{7.5}$$

where the weight parameter w_{opt} regulates the balance between maximizing field strength in the FSDAs and suppressing leakage toward the HSSAs. Accordingly, the overall optimization problem is expressed in the standard form as

$$\boldsymbol{\Phi}_{\text{opt}} = \arg\max_{\boldsymbol{\Phi}} \mathcal{J}(\boldsymbol{\Phi}), \tag{7.6}$$

which is efficiently solved via a single-step gradient descent update, expressed as

$$\mathbf{\Phi}_{\text{opt}} = \mathbf{\Phi}_{\text{init}} - \mu \nabla_{\mathbf{\Phi}} \mathcal{J}(\mathbf{\Phi}), \tag{7.7}$$

where μ represents the learning rate (in m^2/V^2), and the initialization $\mathbf{\Phi}_{\text{init}}$ is set equal to $\mathbf{\Phi}_{\text{common}}$. Therefore, the final optimized phase profile $\mathbf{\Phi}_{\text{opt}}$ is then applied across the RIS elements to realize the spatial signal shaping required for the RF-Fencing service.

To quantify the performance of SHIELD, we utilize P_k which is defined as the ratio of the electric field magnitudes at the POIs between the final optimized configuration and the initial beam-steering scenario, which is given as

$$P_k = 20 \log_{10} \left(\frac{E_k^{\text{POI,final}}}{E_k^{\text{POI,initial}}} \right), \tag{7.8}$$

where $E_k^{\text{POI,final}}$ and $E_k^{\text{POI,initial}}$ denote the electric field magnitudes at the POI after applying SHIELD and under the initial, optimal beam-steering configuration, respectively. In the case of FSDAs, values of P_k approaching 0 dB indicate that users experience nearly identical performance compared to conventional beam steering, thus preserving communication quality. Conversely, for HSSAs, the objective is to maximize suppression, leading to strongly negative values of P_k that reflect effective mitigation of unintended emissions. In scenarios involving multiple FSDAs or HSSAs, the overall performance is obtained by computing the individual amplitude ratios for each region using Eq. (7.8) and averaging the resulting field values across all corresponding POIs.

Finally, the computational complexity of SHIELD is determined by the RIS size $N_{\text{el}} = N_r \times N_c$ and the number of iterations N_I required for convergence. Both dominant region identification and multiplexed field computation scale linearly with the number of RIS elements, i.e., $\mathcal{O}(N_{\text{el}})$, while the gradient-based refinement step requires full-array updates at each iteration, resulting in an overall complexity of $\mathcal{O}(N_I \cdot N_{\text{el}})$.

7.2 Evaluation of SHIELD

In this section, we evaluate SHIELD in managing directional beam control and spatial region suppression across distinct propagation regimes and deployment requirements. First, we assess SHIELD's ability to slice the network into FSDAs and HSSAs, steering and suppressing EM emissions under varied conditions using the emulations of Section 3.1.1, with E_{POI} denoting the EM field within FSDAs/HSSAs. We then extend the evaluation to near-field environments (Section 3.1.2), shifting from directional beam management to the creation of spatially transform ZoI into quiet zones in enclosed spaces, ensuring localized EM silence without degrading surrounding performance. In all scenarios, the incident plane wave has reference magnitude $E_0 = 1$ V/m, serving as a baseline for evaluating the scattered field.

7.2.1 PWE Slicing into FSDAs and HSSAs

We first focus on the far-field regime, where the control of RIS-scattered fields is inherently determined by angular distributions, and varying operating frequencies and RIS configurations directly influence both the physical system behavior and the complexity of the optimization process. Specifically, by examining a THz scenario and one mmWave scenario, we demonstrate SHIELD's ability to dynamically manage multiple FSDAs and HSSAs across diverse deployment contexts and propagation conditions.

THz Scenario

The performance of SHIELD is first evaluated within a THz-band communication network, operating at a carrier frequency of 1 THz. The RIS is configured as a 50×50 element array, with an inter-element spacing of $\lambda_w/5$. To model the EM behavior of the system, the far-field expressions, as described in Section 3.1.1, are employed, computing the scattered field E_{field} over an angular resolution of 1° to ensure precise evaluation of spatial field variations. During the manufacturing phase, a total of 500 precomputed, optimal beam-steering configurations, covering combinations of (AoA, AoD), are stored within the codebook for use during operation. The online evaluation focuses on the creation of two FSDAs and one HSSA. Moreover, the input parameters for SHIELD include a convergence tolerance of 10^{-3}, an optimization weight factor of $w_{\text{opt}} = 0.5$, a learning rate of $\mu = 0.02$, and a compromise factor $\eta = 0.75$. Finally, the thresholds for identifying dominant regions are set to $\tau_{\text{FSDA}} = 0.95$ and $\tau_{\text{HSSA}} = 0.96$, respectively.

To illustrate SHIELD's behavior in this scenario, Fig. 7.3 presents the E_{field} distributions before and after optimization. In more detail, the initial beam-steering configurations are shown in the top row, targeting the selected FSDAs at $(30^\circ, 75^\circ)$

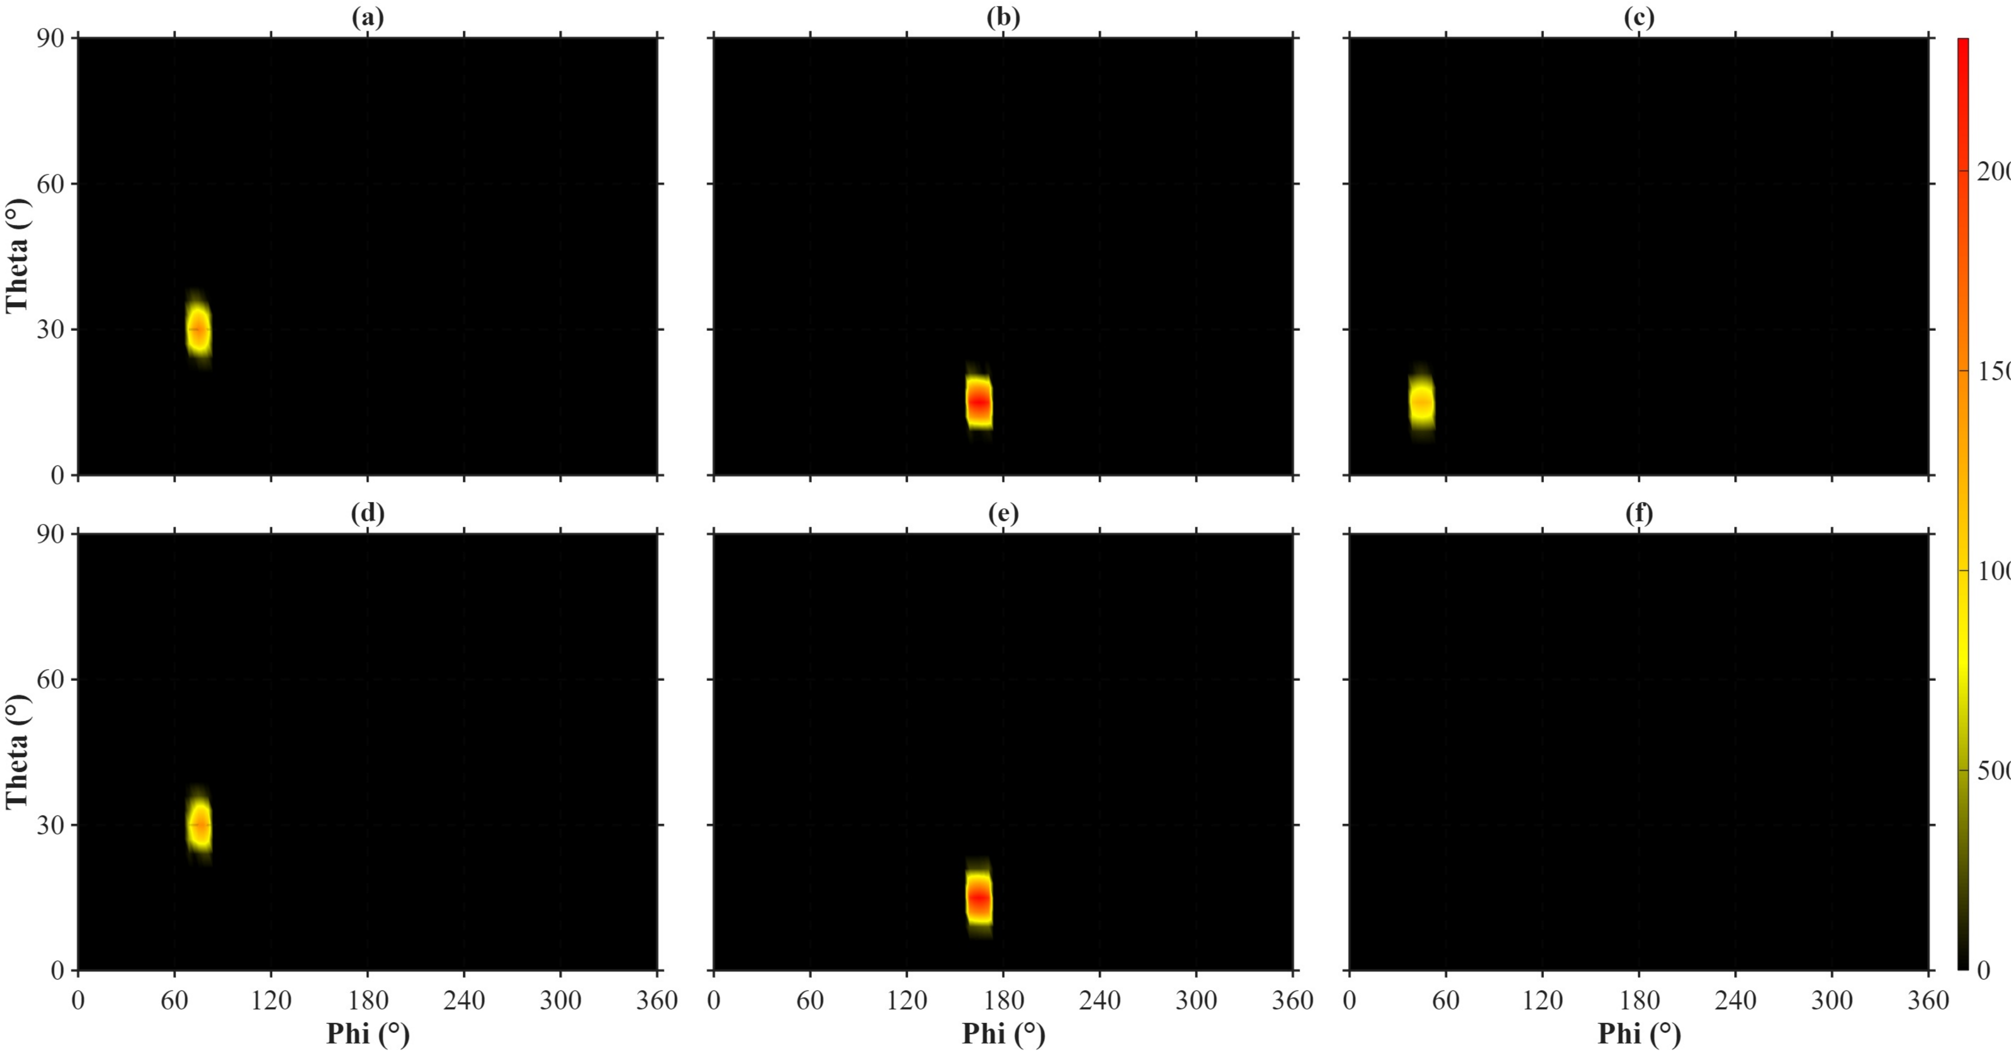

Figure 7.3: E_{field} magnitude before (top) and after (bottom) SHIELD usage for FSDAs $(30°, 75°)$, $(15°, 165°)$ and HSSA $(15°, 45°)$.

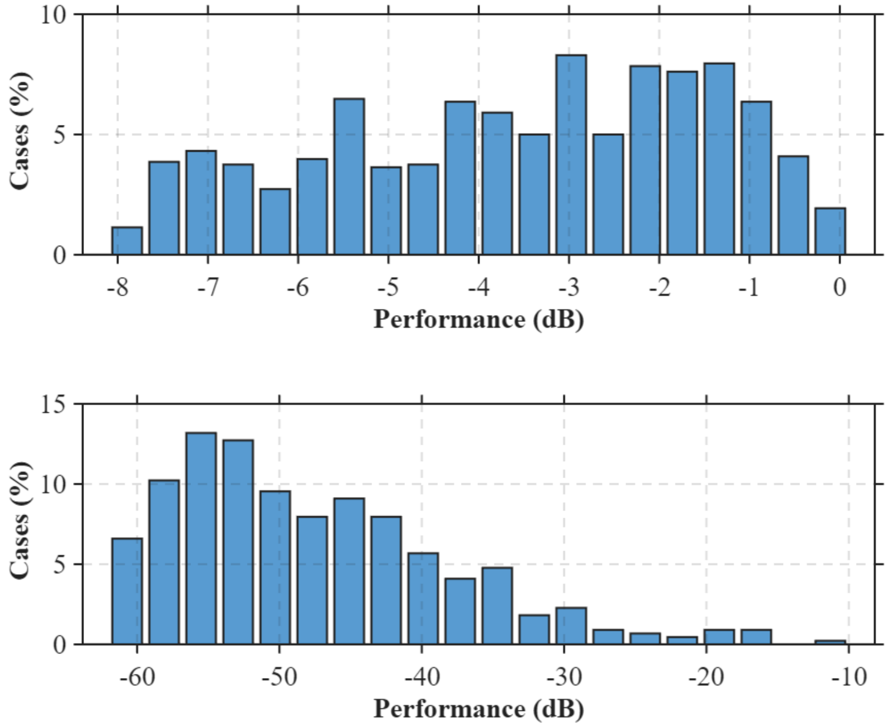

Figure 7.4: Cumulative performance results across multiple cases for FSDAs (up) and HSSAs (down).

and $(15^\circ, 165^\circ)$, alongside the identified HSSA at $(15^\circ, 45^\circ)$. In this initial condition, the E_{POI} values measured at the two FSDAs are 1530.9 V/m and 2332.5 V/m, respectively, whereas the HSSA exhibits a field magnitude of 1207.4 V/m. After applying SHIELD, the HSSA field falls to 0.41 V/m—a -69.25 dB reduction—while the FSDAs still receive 1447 V/m and 2269.4 V/m, corresponding to only -0.49 dB and -0.24 dB losses in respect of their optimal performances. These results collectively illustrate SHIELD's effectiveness in suppressing undesired emissions towards the HSSA while maintaining the signal quality within the intended service regions.

Fig. 7.4 provides a comprehensive summary of SHIELD's performance across 500 evaluated scenarios, highlighting both suppression effectiveness and service continuity. For HSSAs, SHIELD achieves suppression levels exceeding -50 dB in approximately 49.1% of cases, between -50 dB and -20 dB in 49.32%, and greater than -20 dB in only 1.59%. These results confirm SHIELD's consistent capability to mitigate unwanted emissions across a wide range of spatial configurations. In terms of FSDA, the distribution of performance E_{POI} illustrates SHIELD's ability to preserve

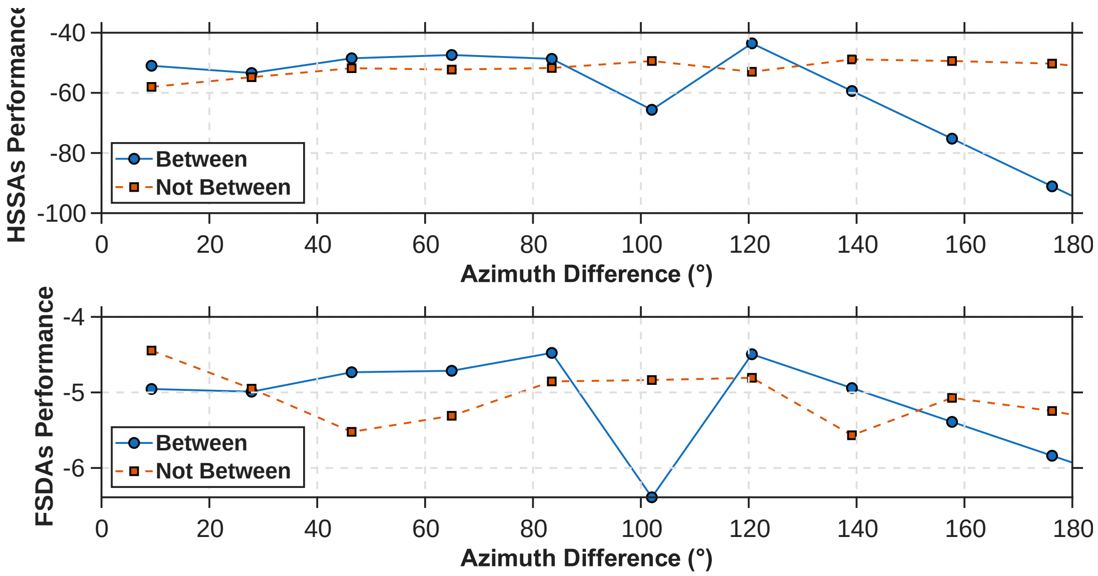

Figure 7.5: Azimuthal difference effect when the single HSSA lies between (blue) or outside (red) the two FSDAs.

communication quality for legitimate users despite the stringent suppression requirements. Specifically, about 29.3% of cases experience less than 2 dB reduction relative to a beam-steering baseline—where the RIS serves each user optimally and independently without simultaneous HSSA suppression. Moreover, 30.7% and 23.9% of cases exhibit moderate reductions in E_{POI} within the ranges $[-4, -2]$ dB and $[-6, -8]$ dB, respectively, while only 16.1% show more notable degradations up to -8 dB. These results underscore SHIELD's robustness in balancing precise spatial suppression with reliable service provision, confirming its suitability for enabling covert communications in complex environments with multiple, spatially distinct performance objectives.

Finally, Figs. 7.5 and 7.6 analyze the influence of angular separation between HSSAs and FSDAs, as well as physical constraints, on SHIELD's performance. Specifically, in Fig. 7.5, performance is plotted against azimuthal separation for two layouts: one with the HSSA outside the angular span of the two FSDAs, and one with the HSSA between them. When the HSSA lies outside the FSDA span, SHIELD consistently achieves suppression of approximately -50 to -60 dB, with only controllable impact on FSDA performance relative to the unconstrained beam-steering baseline. In the more challenging arrangement, with the HSSA between the FSDAs, SHIELD maintains robust suppression even at narrower angles. When the angular separation becomes clearer, suppression is also achieved up to -80 dB. Fig. 7.6 reveals similar trends with elevation angle differences, showing suppression consistently within ap-

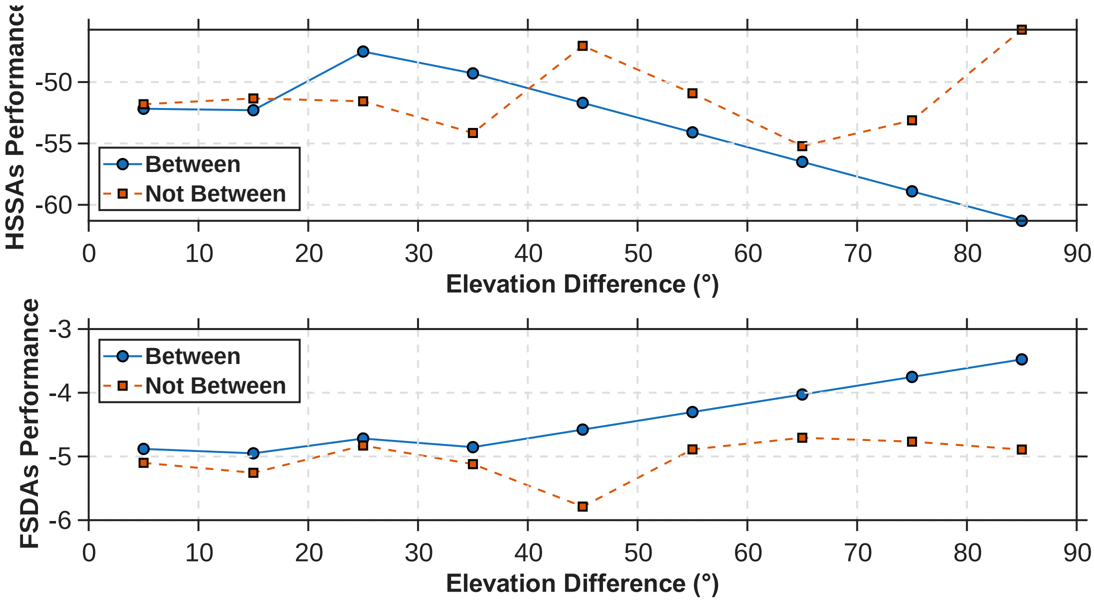

Figure 7.6: Elevation difference impact on SHIELD performance for the same HSSA/FSDA layout categories as in Fig. 7.5.

proximately -45 to -65 dB for the HSSA, while FSDA performance approaches its optimal value as elevation separation increases. Therefore, Figs. 7.5 and 7.6 confirm the effectiveness of SHIELD in ensuring substantial suppression and service continuity, validating its applicability in dense multi-RIS deployments.

mmWave Scenario

To further evaluate the versatility of SHIELD under distinct propagation conditions, we consider an mmWave communication network operating at a carrier frequency of 300 GHz. In this setup, the RIS comprises a 75×75 element array with an inter-element spacing of $\lambda_w/5$. The associated codebook includes 500 precomputed beam-steering configurations, ensuring adequate coverage across the angular domain despite the broader beam profiles inherent at this frequency compared to the THz scenario. The goal here is the slicing of the network into two HSSAs and one FSDA, trying to estimate the SHIELD performance in a more challenging scenario. The evaluation encompasses 500 distinct cases. Moreover, SHIELD is applied using for threshold values $\tau_{\text{FSDA}} = 0.97$ and $\tau_{\text{HSSA}} = 0.99$. The compromise factor η is equal to 0.1 while the optimization weight factor is set to $w_{\text{opt}} = 150$. The learning rate μ is kept equal to 0.02.

Fig. 7.7 presents the cumulative performance distribution for all evaluated mmWave cases, highlighting SHIELD's ability to achieve reliable suppression in undesired di-

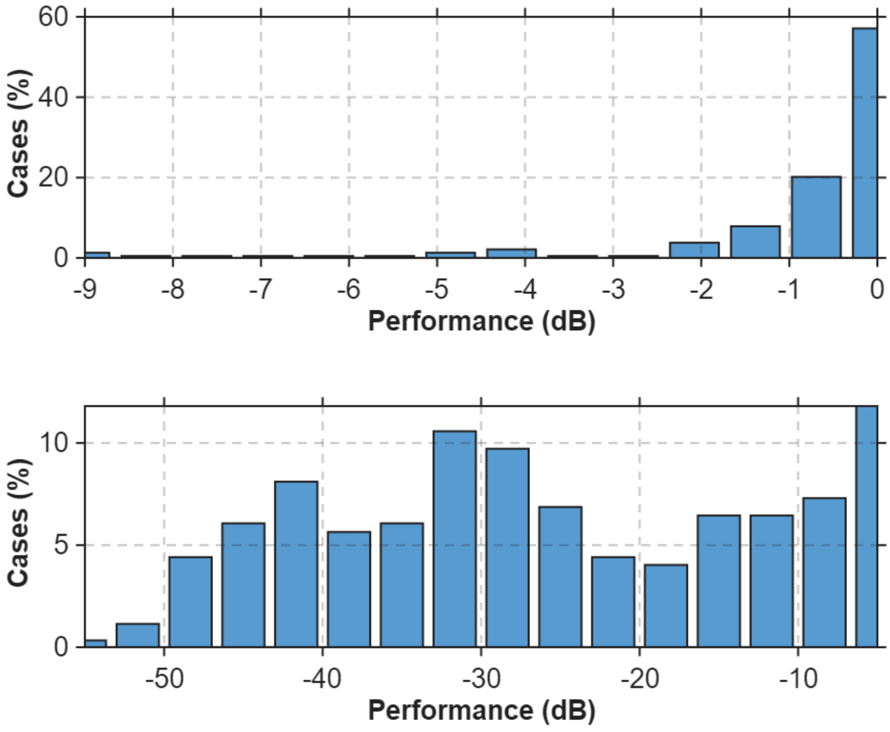

Figure 7.7: Cumulative performance results across multiple cases for FSDAs (up) and HSSAs (down).

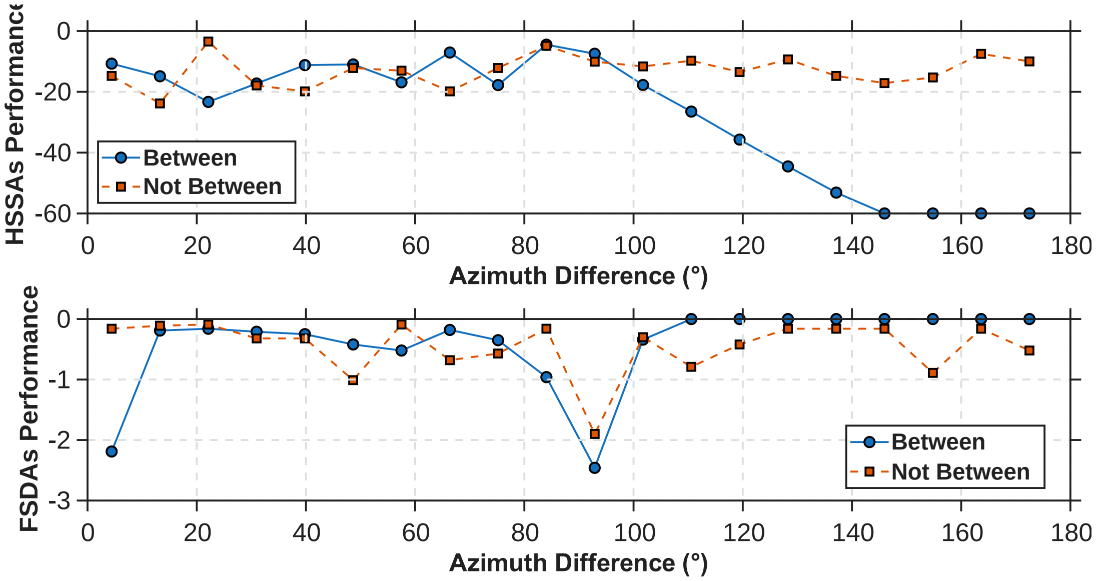

Figure 7.8: Azimuth difference when FSDA is between (blue) and not between (red) the HSSAs.

rections while maintaining high service quality in intended areas. In the HSSAs, suppression exceeds -30 dB in approximately 25% of cases, with a further 50% achieving attenuation between -30 dB and -10 dB, confirming consistent mitigation of unwanted emissions across most scenarios. Among the remaining cases, suppression remains significant, largely within the range of -10 dB to -5 dB. In terms of FSDA performance, 89.43% of cases exhibit negligible impact relative to the optimal beam-steering baseline, with 4.8% and 2.4% experiencing minor reductions within $[-4, -2]$ dB and $[-6, -4]$ dB, respectively. The final 0.81% of cases show more substantial reductions beyond -8 dB, reflecting the inherent trade-offs involved in managing multiple HSSAs simultaneously. These results demonstrate SHIELD's capability to balance strong suppression with service continuity, even under the broader beam patterns and complex propagation conditions characteristic of mmWave environments.

Fig. 7.8 examines the influence of azimuthal separation between HSSAs and FSDAs on the performance of SHIELD in mmWave environments, considering configurations where the FSDA is either located between or outside the angular span of the HSSAs, as before. When the FSDA is placed outside this range, HSSA suppression stays between -10 dB to -20 dB for most separations, and the FSDA achieves optimal performance or at most a -2 dB loss. When the FSDA lies between the HSSAs, HSSA suppression worsens by about -1 dB on average, while the FSDA performance remains near its optimal value. As the azimuthal difference increases, HSSA suppression improves sharply—reaching up to -60 dB, while the FSDA's performance returns to the baseline optimum. Thus, Fig. 7.8 underscores SHIELD's ability to dynamically adapt to spatial constraints, with performance gains becoming more pronounced as angular diversity increases, highlighting its strategic advantage in managing interference and maintaining selective coverage in complex environments.

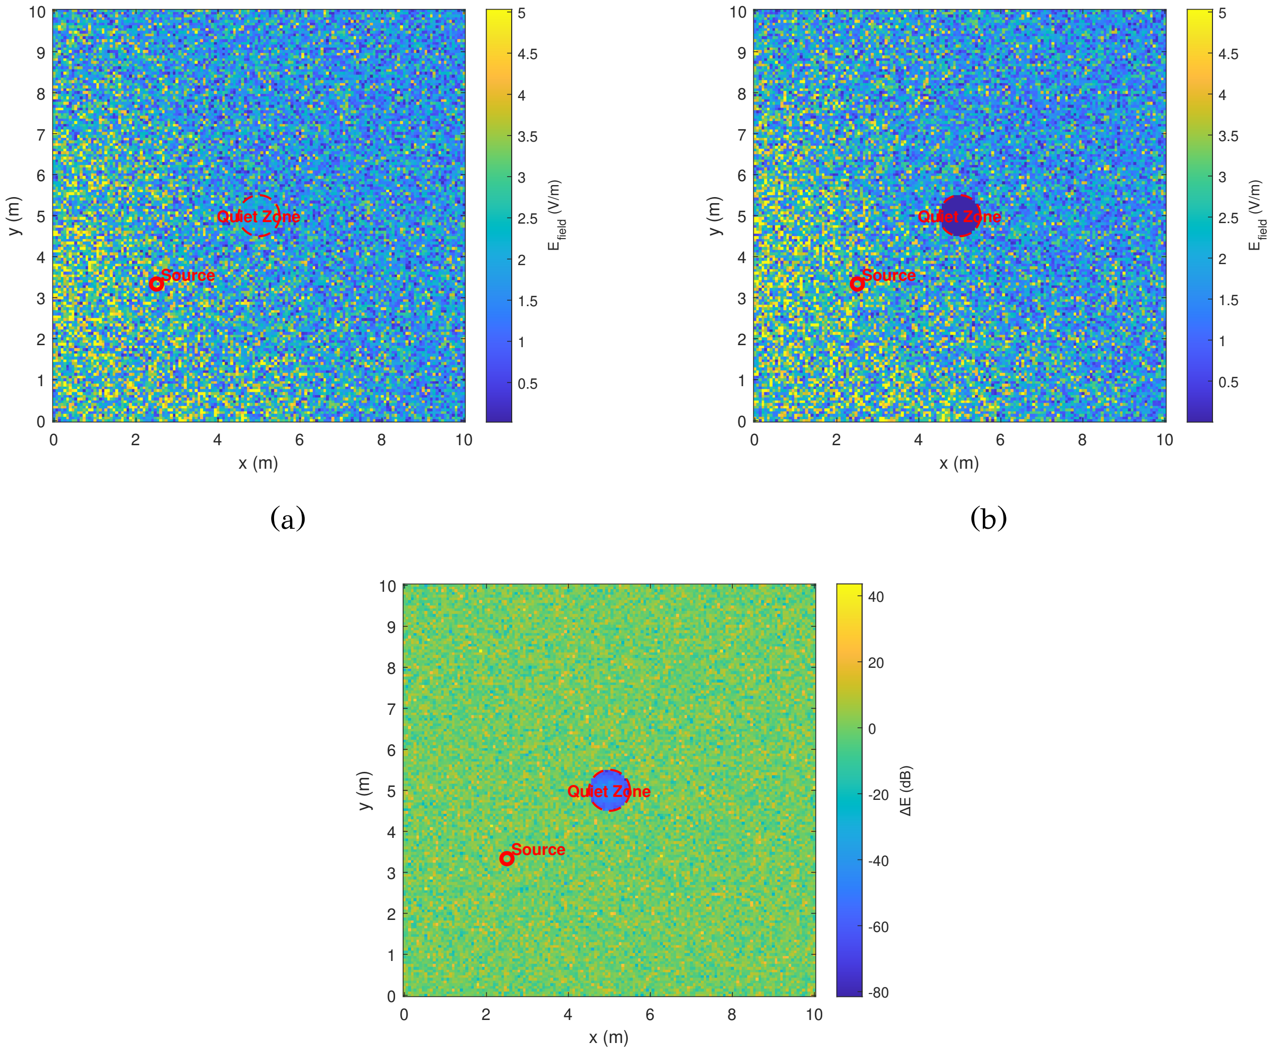

(a)

(b)

(c)

Figure 7.9: Spatial distribution of E_{field} (in V/m) before (left) and after (center) SHIELD, and the corresponding point-wise deviation (right) in dB.

7.2.2 Quiet Zone Creation within PWE

The next step of our evaluation focuses on the creation of a quiet zone within an enclosed environment, aiming to suppress EM leakage within a specific spatial region while preserving reflections elsewhere. Specifically, in dense indoor settings, the objective is to demonstrate fine-grained EM control for localized suppression, where the confined environment amplifies the impact of multipath reflections and the close proximity of RIS elements intensifies the need for highly precise tuning to effectively isolate the quiet zone. Conversely, in expansive outdoor scenarios, the emphasis moves to validating SHIELD's scalability and its ability to sustain effective suppression across broader areas, where the propagation environment becomes increasingly complex and the trade-off between control precision and coverage must be carefully managed. Thus, by exploring both scenarios, we comprehensively assess SHIELD's capability for quiet-zone formation across various realistic deployment conditions.

In this direction, SHIELD is applied to jointly optimize four RIS units positioned along the walls of a square area of side length L, covering both enclosed indoor spaces and larger-scale outdoor environments. The objective is to minimize the RIS-scattered electric field E_{field} within a predefined suppression region of radius R_0 considering it as an HSSA and ensuring proactive covertness against potential wardens in the targeted area. The source transmitter is assumed to be freely positioned at any point within the space, reflecting a realistic deployment scenario and the E_{field} in each region within the room is computed as detailed in Sec. 3.1.2. The process begins by configuring all RIS units as perfect electric conductors, initializing their reflection phases uniformly to $\Phi_n = \pi$. As a first step, SHIELD computes for each RIS the phase shift that minimizes the EM propagation in the quiet zone while keeping all other elements fixed. This per-element solution serves as the initial guess for the upcoming optimization procedure, thereby accelerating convergence and avoiding very poor local minima. Subsequently, a binary spatial mask is introduced to explicitly define the quiet zone, assigning a logical value of one to grid points within the spherical region satisfying Eq. (3.14), and zero elsewhere, thereby isolating the optimization domain and enhancing computational efficiency. During each iteration of the gradient-based optimization routine, individual RIS elements sequentially test incremental phase variations $\Phi_n \pm \Delta\phi$, starting from an initial step size $\Delta\phi$. For every candidate phase configuration $\mathbf{\Phi}$, the average scattered field magnitude $E_{\text{field}}^{\text{avg}}$ within the quiet zone is recalculated, and each RIS element is updated according to:

$$\Phi_n \leftarrow \arg \min_{\Phi_n \pm \Delta\phi} E_{\text{field}}^{\text{avg}}. \tag{7.9}$$

After completing updates for all RIS elements, if no further improvement is detected, the step size $\Delta\phi$ is halved to refine the search. The procedure iteratively continues until the power level within the quiet zone falls below a predefined threshold or the maximum number of iterations is reached, ultimately resulting in a minimized residual scattered field within the suppression region.

Emulation Setup for Indoor Scenario

In this scenario, we consider a cubic room with side length $L = 10$ m. The operating frequency is set to 28 GHz, reflecting typical sub-6 GHz indoor deployments scaled to mmWave bands. The objective is to establish as HSSA a spherical quiet zone at the center of the room, defined by coordinates $(L/2, L/2, 0)$ and radius $R_0 = 0.5$m, while minimizing the impact on the FSDA that consists of all the surrounding regions. The environment outside the quiet zone is discretized into a three-dimensional grid with $N_x = N_y = N_z = 251$ points whereas within it the denser grid consists of about $3,400$ points for more accurate computations. Regarding the optimization procedure, the maximum number of iterations is fixed at 150, with a convergence tolerance of 10^{-9}.

Fig. 7.9a illustrates the initial $|E_{\text{field}}|$ distribution within the room under the default RIS configuration, where both the source and the designated quiet zone are marked with red circles, and the field magnitude within the quiet zone is nearly identical to that of the surrounding space. This uniformity highlights the absence of any suppression mechanism, as the quiet zone initially experiences an average field magnitude of 2.01V/m, closely matching the 2.14 V/m observed elsewhere. Following the application of SHIELD, the optimized distribution in Fig. 7.9b reveals a marked attenuation within the quiet zone, now visually and quantitatively distinguished from the rest of the environment. Specifically, the average field magnitude within the quiet zone is reduced to 4.5×10^{-3}V/m, while the surrounding area remains effectively unchanged at 2.13 V/m, confirming SHIELD's ability to deliver strong, localized suppression without disrupting the broader EM behavior. The extent and precision of this suppression are further quantified in Fig. 7.9c, which shows the point-wise deviation ΔE across the domain. Within the quiet zone, suppression reaches approximately -49 dB, while outside, deviations remain minimal, ranging from -1 dB in adjacent points to $+1.5$ dB at greater distances. This outcome directly reflects the spatial mask

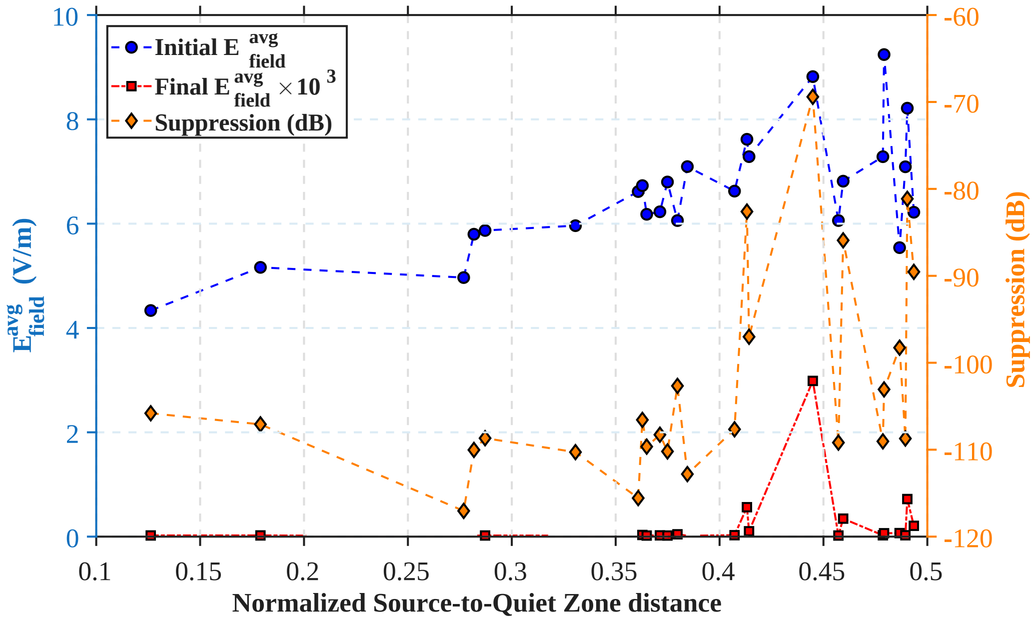

Figure 7.10: Initial and optimized values of $E^{\text{avg}}_{\text{field}}$, along with the corresponding suppression levels.

and optimization process, which are specifically tuned to isolate EM suppression to the intended region while preserving normal field propagation elsewhere. Moreover, the ability to achieve such deep suppression in a confined area, with negligible impact beyond its boundaries, demonstrates the potential of SHIELD as a practical tool for spatially controlled EM quiet zones.

Fig. 7.10 illustrates the initial and optimized values of the average electric field $E^{\text{avg}}_{\text{field}}$ within the quiet zone, together with the corresponding suppression levels in dB, plotted against the normalized distance between the source and the quiet zone relative to the room dimension L. The analysis considers multiple source positions along a circular trajectory centered at $(L/2, L/2, 0)$, ranging from a radius of 1 meter—coinciding with the quiet zone boundary—up to 5 meters, thereby encompassing diverse spatial configurations. Throughout this evaluation, the suppression region and computational settings remain fixed, ensuring that the results solely reflect the influence of source placement. Under default RIS conditions, the average $E^{\text{avg}}_{\text{field}}$ within the quiet zone remains consistently between 4.3 and 9.2 V/m across all positions, highlighting the absence of any intrinsic spatial shielding. After optimization, the quiet zone field magnitude is reduced uniformly to approximately 2.2×10^{-4} V/m, with a standard deviation of 6×10^{-4} V/m, yielding suppression levels between -69 and -117 dB. Therefore, Fig. 7.10 confirms the capability of SHIELD to achieve stable and effective suppression, irrespective of source location, thereby reinforcing its suitability for consistent quiet-zone formation in dynamic scenarios.

Emulation Setup for Outdoor Scenario

In this scenario, we consider the creation of a quiet zone within an open-area scenario with side length $L = 100$ meters. In this setup, RIS units are positioned on the surrounding buildings to ensure comprehensive spatial coverage. The quiet zone is defined with a radius of $R_0 = 3$ meters, centered at coordinates $(L/4, L/3, 0)$, intentionally offset from the geometric center of the area to emulate realistic deployment scenarios. The operating frequency is maintained at 28 GHz. The discretization of the computational grid without the quiet zone is set to $N_x = N_y = N_z = 351$ whereas within it the grid consists of $22,400$ points for more precise results. The optimization process runs for 50 iterations, using a convergence tolerance of 10^{-9}.

Figs. 7.11a and 7.11b illustrate the spatial distributions of E_{field} before and after the application of SHIELD, with the source located at $\left(\frac{9L}{10}, \frac{9L}{10}, 0\right)$. Initially, the average scattered field magnitude within the quiet zone is approximately 5.1×10^{-2} V/m, slightly lower than the 6.6×10^{-2} V/m observed outside, primarily due to the greater distance from the source. After SHIELD is applied, the field within the quiet zone is suppressed to 6.5×10^{-9} V/m, while the field outside remains stable at 6.3×10^{-9}V/m, demonstrating that suppression is confined to the designated region. This effect is further corroborated by Fig. 7.11c, where the deviation reaches approximately -164 dB within the quiet zone, and remains limited to the range of -0.2 to $+0.25$ dB elsewhere. These findings confirm SHIELD's scalability and its ability to achieve highly localized suppression in outdoor scenarios, effectively managing RIS-induced reflections.

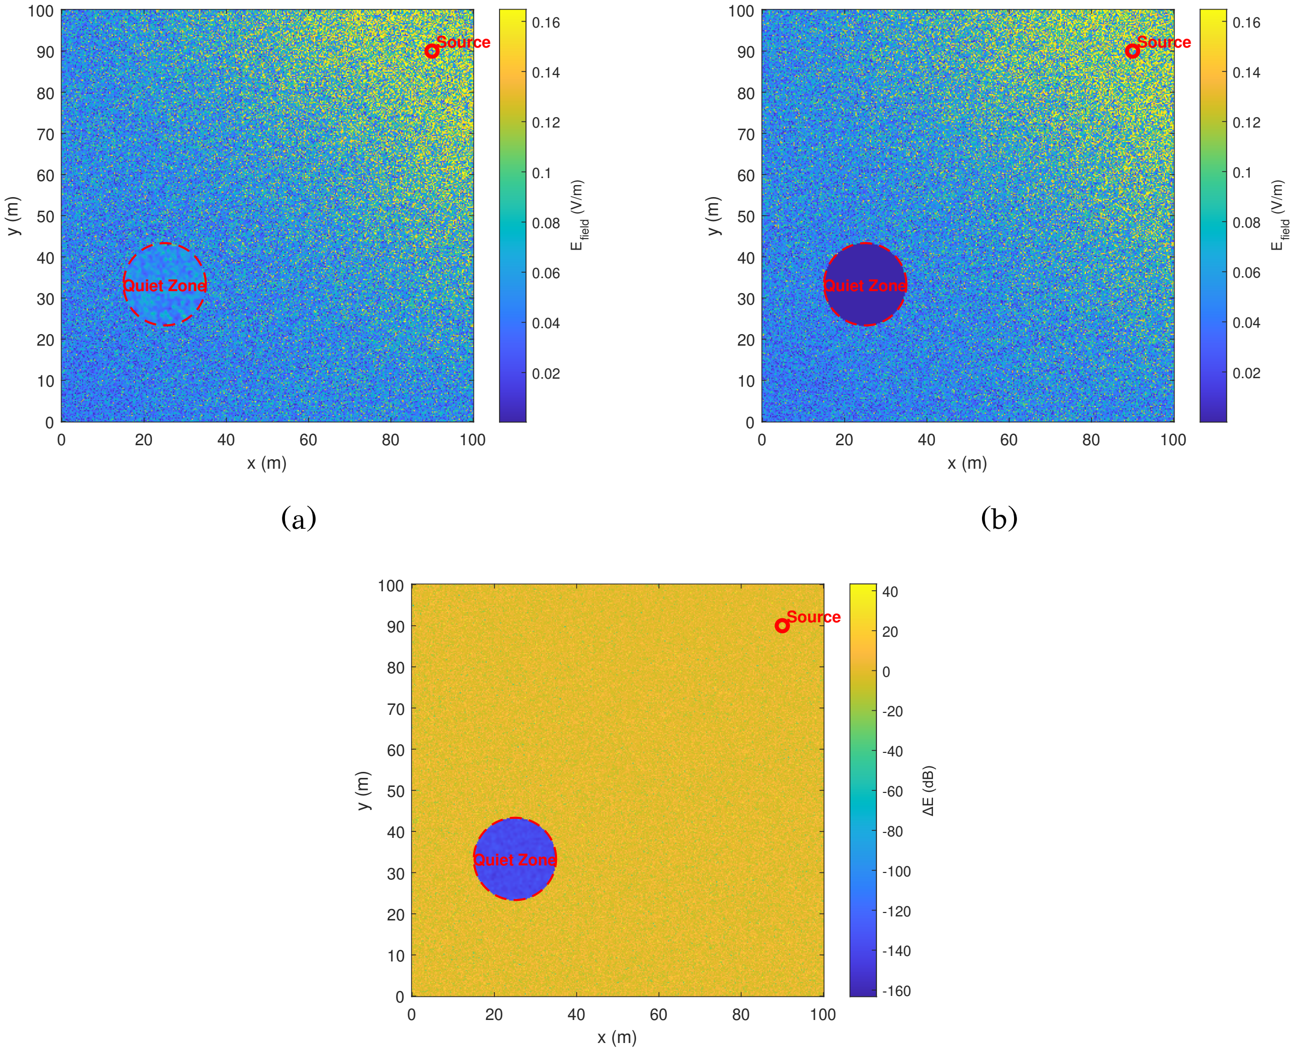

(a)

(b)

(c)

Figure 7.11: Spatial distribution of E_{field} (in V/m) before (left) and after (center) SHIELD, and the corresponding point-wise deviation (right) in dB.

7.3 Discussion

RF-Fencing extends the codebook-multiplexing paradigm beyond multi-user resource allocation to proactive covert communications in RIS-assisted PWEs. By treating the propagation medium as an entity that can be instructed to both *deliver* and *withhold* energy, RF-Fencing operationalizes spatial selectivity as a service primitive: strong field levels are preserved within designated FSDAs, while radiation is actively suppressed toward proactively selected HSSAs and fully confined quiet zones.

SHIELD provides a practical controller for this service. It combines manufacturing-time, EM-consistent codebooks (beam-steering entries over (AoA, AoD) with operating-time multiplexing and a lightweight refinement step to synthesize the phase matrix $\mathbf{\Phi}$ without resorting to heavy iterative optimization or full channel estimation. The evaluations confirm the core design premise. In far-field slicing (THz and mmWave), SHIELD achieves substantial suppression in HSSAs while maintaining near-baseline performance in FSDAs, and the observed trends with azimuth/elevation separation expose the expected geometric trade-offs: suppression becomes easier as angular diversity increases and becomes most constrained when suppression regions lie between service lobes. In near-field operation, SHIELD generalizes from angular masks to spatial masks and forms highly localized quiet zones in both indoor and outdoor settings, demonstrating that the same codebook-driven logic can be extended from directional control to spatial confinement under rich multipath and multi-RIS interaction.

These results also clarify what RIS uniquely contribute to future PWEs. While LoS restoration is a compelling baseline use case, it can often be approximated through densification (additional access points) as cell sizes shrink and deployments become denser at higher frequencies. In contrast, services that require *intentional shaping of exposure*—such as proactive covertness, physical-layer security, and controlled quiet zones—depend directly on software-defined manipulation of the propagation medium and benefit materially from coordinated multi-RIS control. In this sense, RF-Fencing and SHIELD exemplify a class of PWE-native capabilities that are difficult to replicate by densification alone, and they provide an architecture-consistent, computationally efficient route to privacy-aware propagation control for B5G/6G networks.

CHAPTER 8

CONCLUSIONS

This thesis advanced the operational maturity of Reconfigurable Intelligent Surfaces (RISs) as infrastructure for Programmable Wireless Environments (PWEs) by addressing three intertwined challenges: (i) physics-consistent and time-efficient codebook compilation, (ii) low-latency multi-user sharing under operator pricing policies, and (iii) fairness-aware coordination of multiple RIS units at the network level. Across these layers, the central theme is that RIS control must be grounded in the coupling between microscopic element states and macroscopic electromagnetic (EM) responses, while remaining compatible with stringent B5G/6G latency and scalability requirements. The developed methods follow a consistent lifecycle view—manufacturing (offline) computation of configurations and operating-phase (online) selection, multiplexing, and orchestration—so that sophisticated propagation control becomes deployable without per-element online optimization.

8.1 Summary of Contributions

- **Physics-based codebook compilation.** We introduced a sequence of physics-informed tools that progressively elevate compilation fidelity and practicality. PICCS demonstrated that physical-layer insights can drastically reduce optimization complexity by exploiting element, response correlations and GO-based initialization, enabling near-optimal configurations with orders-of-magnitude fewer EM evaluations. Building on this, MATCH introduced a staged optimization synergy that explicitly incorporates mutual coupling and specular reflections in near-field focusing and, critically, treats focusing as a holistic energy-redistribution problem: maximizing energy inside the target region while suppressing it elsewhere. Finally, OSIRIS addressed the scalability bottleneck of 'optimal-per-location' compilation by separating physics-intensive optimization from real-time operation through sparse 3D sampling, phase-domain compression, and spatial anchoring, achieving near-optimal SNR with predictable, bounded online complexity and multi-order runtime reduction relative to exhaustive optimization.

- **RIS as a network resource and codebook-entry multiplexing.** A foundational contribution is the formalization of RIS as a shareable network resource hosted by macroscopic EM functions, together with a resource-slice notion that connects multi-user efficiency to deviation from single-user optimal configurations. This abstraction enabled a practical operating-phase mechanism, *codebook-entry multiplexing*, in which multi-user service is realized by combining precompiled entries rather than performing repeated channel estimation and re-optimization. Within this paradigm, COMMON provided the baseline equal-sharing mechanism, PRIME introduced pricing-aware multiplexing via voting/payment factors with millisecond-level latency, and LEVEL advanced policy enforcement by directly controlling per-user configuration deviation, yielding more predictable tier differentiation and improved fairness under dense contention and quantized hardware.

- **System-level orchestration and fairness in RIS-enabled networks.** To coordinate multiple RIS units, we introduced THEMIS, a framework that integrates pricing-aware fairness into both node-level configuration formation and end-to-end path selection over cascaded RIS graphs. THEMIS extends graph-based

RIS networking by embedding stakeholder policy directly into (i) the computation of shared configurations under an efficiency floor and (ii) congestion-aware routing whose cost reflects physical distance, edge load, and user tier. The resulting framework unifies codebook-based control, multiplexing, and routing into a deployable workflow with modest computational and memory overheads across dense network scenarios.

- **RF-Fencing as a proactive covert-communications service.** Beyond performance-oriented services, this thesis showed that codebook-driven control enables PWE-native missions that are difficult to replicate by densification alone. RF-Fencing was introduced as a paradigm for spatially selective signal delivery and suppression, and SHIELD as the first dedicated controller that synthesizes friendly-service delivery areas and hostile-suppression areas (and fully confined quiet zones) by multiplexing manufacturing-time, EM-consistent codebook entries with lightweight refinement. Emulations across THz/mmWave and near-field indoor/outdoor quiet-zone settings validated that substantial suppression can be achieved while maintaining service continuity, with clear geometric trade-offs driven by angular and spatial separation.

8.2 Key Findings

- **Physics awareness is not optional in codebook compilation.** Accurate EM effects (finite-size diffraction, coupling, reflections) shape the configuration landscape; incorporating them early improves both convergence and interpretability.
- **The deviation between shared and single-user configurations is a robust control handle.** It provides a fast proxy for performance degradation and a practical lever for enforcing fairness and pricing policies with minimal online latency.
- **Tiered fairness must be enforced at both the RIS-node and network levels.** Local sharing alone is insufficient when multiple RIS units and congested paths interact; integrating fairness into routing and node-level configuration formation yields stable behavior under density.

- **PWEs unlock capabilities beyond LoS restoration.** While densification can address some coverage problems, services that require intentional exposure control—covert communications, quiet zones, and proactive suppression—benefit directly from coordinated multi-RIS programmability.

8.3 Future Directions

- **Tile-centric control is the scalability bottleneck.** Moving from per-element actuation to tile-level abstractions, while still exploiting EM structure (coupling, edge dominance, sensitivity patterns), is key to realizing large, networked metasurface coatings with bounded control overhead.
- **Passive RIS must remain the default, but hardware innovation is accelerating.** Energy-efficient passive operation should be preserved, yet emerging *novel passive RIS schemes* and large-area manufacturing routes (e.g., large-area electronics and hybrid integration) should be evaluated as practical enablers for deployment at scale.
- **Codebooks need physics-constrained structure, not just size reduction.** Compression and selection should preserve EM consistency and interpretability, enabling predictable, bounded-complexity inference over continuous user locations and robust behavior under moderate topology perturbations.
- **Multi-user sharing requires explicit feasibility management.** Pricing- and policy-aware resource allocation should be extended beyond discrete tiers, complemented by admission control, and accompanied by quantization-aware guarantees that limit tier inversion and minimum-efficiency violations.
- **Compilation-time influence information should drive runtime orchestration.** Element/tile influence scores and sensitivity matrices should be carried forward to prioritize the degrees of freedom that matter most, improving both efficiency and stability in configuration selection, sharing, and network coordination.
- **Security, sensing, and networking should be co-designed.** Spatial exposure control (suppression zones, quiet regions) should be integrated with resource

management and routing to deliver unified multi-mission programmable environments under low-latency control.

- **Network-scale testbeds are the next credibility milestone.** Multi-RIS experimental platforms with mobility and real-time loops are required to quantify overheads, robustness, and interoperability with emerging 6G architectures under realistic traffic dynamics.

8.4 Closing Remark

Overall, this thesis delivers an end-to-end, physics-consistent control stack for RIS-assisted PWEs—spanning codebook compilation, multi-user multiplexing with operator policy, network-level orchestration, and proactive security services. By replacing per-element online optimization with structured offline computation and lightweight online control, the proposed methods offer a practical route to scalable, fair, and mission-flexible programmable propagation in B5G/6G networks.

BIBLIOGRAPHY

[1] S. Tsimpoukis *et al.*, "Realizing rf wavefront copying with ris for future extended reality applications," in *2024 7th International Balkan Conference on Communications and Networking (BalkanCom)*. IEEE, 2024, pp. 37–42.

[2] M. Segata *et al.*, "Cooperis: A framework for the simulation of reconfigurable intelligent surfaces in cooperative driving environments," *Computer Networks*, p. 110443, 2024.

[3] C. De Lima, D. Belot, R. Berkvens, A. Bourdoux, D. Dardari, M. Guillaud, M. Isomursu, E.-S. Lohan, Y. Miao, A. N. Barreto *et al.*, "Convergent communication, sensing and localization in 6g systems: An overview of technologies, opportunities and challenges," *IEEE Access*, vol. 9, pp. 26 902–26 925, 2021.

[4] Z. Zhang *et al.*, "Wireless power transfer—an overview," *IEEE transactions on industrial electronics*, vol. 66, no. 2, pp. 1044–1058, 2018.

[5] Z. Zhang and et al., "6G wireless networks: Vision, requirements, architecture, and key technologies," *IEEE Veh. Technol. Mag.*, vol. 14, no. 3, pp. 28–41, 2019.

[6] C. Liaskos *et al.*, "A tutorial on controlling metasurfaces from the network perspective," 2026. [Online]. Available: https://arxiv.org/abs/2601.12118

[7] F. A. P. De Figueiredo, "An overview of massive mimo for 5g and 6g," *IEEE Latin America Transactions*, vol. 20, no. 6, pp. 931–940, 2022.

[8] D. d. S. Brilhante, J. C. Manjarres, R. Moreira, L. de Oliveira Veiga, J. F. de Rezende, F. Müller, A. Klautau, L. Leonel Mendes, and F. A. P. de Figueiredo, "A literature survey on ai-aided beamforming and beam management for 5g and 6g systems," *Sensors*, vol. 23, no. 9, p. 4359, 2023.

[9] B. Clerckx, Y. Mao, Z. Yang, M. Chen, A. Alkhateeb, L. Liu, M. Qiu, J. Yuan, V. W. Wong, and J. Montojo, "Multiple access techniques for intelligent and multifunctional 6g: Tutorial, survey, and outlook," *Proceedings of the IEEE*, vol. 112, no. 7, pp. 832–879, 2024.

[10] M. Rowshan, M. Qiu, Y. Xie, X. Gu, and J. Yuan, "Channel coding toward 6g: Technical overview and outlook," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 2585–2685, 2024.

[11] W. K. Alsaedi, H. Ahmadi, Z. Khan, and D. Grace, "Spectrum options and allocations for 6g: A regulatory and standardization review," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 1787–1812, 2023.

[12] P. Lohan, B. Kantarci, M. A. Ferrag, N. Tihanyi, and Y. Shi, "From 5g to 6g networks: A survey on ai-based jamming and interference detection and mitigation," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 3920–3974, 2024.

[13] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, and I. Akyildiz, "Realizing wireless communication through software-defined hypersurface environments," in *2018 IEEE 19th International Symposium on" A World of Wireless, Mobile and Multimedia Networks"(WoWMoM)*. IEEE, 2018, pp. 14–15.

[14] D. Moltchanov, E. Sopin, V. Begishev, A. Samuylov, Y. Koucheryavy, and K. Samouylov, "A tutorial on mathematical modeling of 5g/6g millimeter wave and terahertz cellular systems," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 1072–1116, 2022.

[15] Q. Xue, C. Ji, S. Ma, J. Guo, Y. Xu, Q. Chen, and W. Zhang, "A survey of beam management for mmwave and thz communications towards 6g," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 3, pp. 1520–1559, 2024.

[16] Y. Hong, T. Thaj, and E. Viterbo, *Delay-Doppler communications: principles and applications*. Academic Press, 2022.

[17] D. Serghiou, M. Khalily, T. W. Brown, and R. Tafazolli, "Terahertz channel propagation phenomena, measurement techniques and modeling for 6g wireless communication applications: A survey, open challenges and future research

directions," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 1957–1996, 2022.

[18] C. Liaskos *et al.*, "Modeling, simulating and configuring programmable wireless environments for multi-user multi-objective networking," *arXiv preprint arXiv:1812.11429*, 2018.

[19] ——, "Software-defined reconfigurable intelligent surfaces: From theory to end-to-end implementation," *Proceedings of the IEEE*, vol. 110, no. 9, pp. 1466–1493, 2022.

[20] S. E. Trevlakis, A.-A. A. Boulogeorgos, D. Pliatsios, J. Querol, K. Ntontin, P. Sarigiannidis, S. Chatzinotas, and M. Di Renzo, "Localization as a key enabler of 6g wireless systems: A comprehensive survey and an outlook," *IEEE open journal of the Communications Society*, vol. 4, pp. 2733–2801, 2023.

[21] C. Psomas, K. Ntougias, N. Shanin, D. Xu, K. Mayer, N. M. Tran, L. Cottatellucci, K. W. Choi, D. I. Kim, R. Schober *et al.*, "Wireless information and energy transfer in the era of 6g communications," *Proceedings of the IEEE*, vol. 112, no. 7, pp. 764–804, 2024.

[22] M. Mitev, A. Chorti, H. V. Poor, and G. P. Fettweis, "What physical layer security can do for 6g security," *IEEE Open Journal of Vehicular Technology*, vol. 4, pp. 375–388, 2023.

[23] A. Lalas *et al.*, "Programmable terahertz metamaterials through v-beam electrothermal devices," *Applied Physics A*, vol. 117, pp. 433–438, 2014.

[24] O. Tsilipakos *et al.*, "Toward intelligent metasurfaces: the progress from globally tunable metasurfaces to software-defined metasurfaces with an embedded network of controllers," *Advanced Optical Materials*, vol. 8, no. 17, p. 2000783, 2020.

[25] A. M. Wong and G. V. Eleftheriades, "Perfect anomalous reflection with a bipartite huygens' metasurface," *Physical Review X*, vol. 8, no. 1, p. 011036, 2018.

[26] M. Jang, Y. Horie, A. Shibukawa, J. Brake, Y. Liu, S. M. Kamali, A. Arbabi, H. Ruan, A. Faraon, and C. Yang, “Wavefront shaping with disorder-engineered metasurfaces,” *Nature photonics*, vol. 12, no. 2, pp. 84–90, 2018.

[27] H.-X. Xu, S. Tang, G.-M. Wang, T. Cai, W. Huang, Q. He, S. Sun, and L. Zhou, “Multifunctional microstrip array combining a linear polarizer and focusing metasurface,” *IEEE Transactions on Antennas and Propagation*, vol. 64, no. 8, pp. 3676–3682, 2016.

[28] Z. Wei, Y. Cao, X. Su, Z. Gong, Y. Long, and H. Li, “Highly efficient beam steering with a transparent metasurface,” *Optics express*, vol. 21, no. 9, pp. 10 739–10 745, 2013.

[29] M. Khorasaninejad, W. Zhu, and K. Crozier, “Efficient polarization beam splitter pixels based on a dielectric metasurface,” *Optica*, vol. 2, no. 4, pp. 376–382, 2015.

[30] Y. Hu, X. Wang, X. Luo, X. Ou, L. Li, Y. Chen, P. Yang, S. Wang, and H. Duan, “All-dielectric metasurfaces for polarization manipulation: principles and emerging applications,” *Nanophotonics*, vol. 9, no. 12, pp. 3755–3780, 2020.

[31] J. Li, W. Wang, Y. Xie, B.-I. Popa, and S. A. Cummer, “A sound absorbing metasurface with coupled resonators,” *Applied Physics Letters*, vol. 109, no. 9, 2016.

[32] M. Pu, Y. Guo, X. Ma, X. Li, and X. Luo, “Methodologies for on-demand dispersion engineering of waves in metasurfaces,” *Advanced Optical Materials*, vol. 7, no. 14, p. 1801376, 2019.

[33] A. Li *et al.*, “Metasurfaces and their applications,” *Nanophotonics*, vol. 7, no. 6, pp. 989–1011, 2018.

[34] A. M. Shaltout, A. V. Kildishev, and V. M. Shalaev, “Evolution of photonic metasurfaces: from static to dynamic,” *Journal of the Optical Society of America B*, vol. 33, no. 3, pp. 501–510, 2016.

[35] A. Pitilakis *et al.*, “Multifunctional metasurface architecture for amplitude, polarization and wave-front control,” *Physical Review Applied*, vol. 17, no. 6, p. 064060, 2022.

[36] Y. Wang, P. Landreman, D. Schoen, K. Okabe, A. Marshall, U. Celano, H.-S. P. Wong, J. Park, and M. L. Brongersma, "Electrical tuning of phase-change antennas and metasurfaces," *Nature Nanotechnology*, vol. 16, no. 6, pp. 667–672, 2021.

[37] A. I. Kuznetsov, M. L. Brongersma, J. Yao, M. K. Chen, U. Levy, D. P. Tsai, N. I. Zheludev, A. Faraon, A. Arbabi, N. Yu *et al.*, "Roadmap for optical metasurfaces," *ACS photonics*, vol. 11, no. 3, pp. 816–865, 2024.

[38] A. C. Overvig, S. A. Mann, and A. Alù, "Thermal metasurfaces: complete emission control by combining local and nonlocal light-matter interactions," *Physical Review X*, vol. 11, no. 2, p. 021050, 2021.

[39] D. Brizi and A. Monorchio, "Magnetic metasurfaces properties in the near field regions," *Scientific Reports*, vol. 12, no. 1, p. 3258, 2022.

[40] Y. Zhao, Z. Liu, C. Li, W. Jiao, S. Jiang, X. Li, J. Duan, and J. Li, "Mechanically reconfigurable metasurfaces: fabrications and applications," *npj Nanophotonics*, vol. 1, no. 1, p. 16, 2024.

[41] C. Liaskos *et al.*, "Using any surface to realize a new paradigm for wireless communications," *Commun. ACM*, vol. 61, pp. 30–33, 2018.

[42] Á. Palomares-Caballero, M. Pérez-Escribano, C. Molero, P. Padilla, M. García-Vigueras, and R. Gillard, "Broadband 1-bit reconfigurable intelligent surface at millimeter waves: Overcoming pin-diode degradation with slotline ring topology," *IEEE Transactions on Antennas and Propagation*, 2025.

[43] L. G. Da Silva, Z. Chu, P. Xiao, and A. Cerqueira S Jr, "A varactor-based 1024-element ris design for mm-waves," *Frontiers in communications and networks*, vol. 4, p. 1086011, 2023.

[44] F. Yang, P. Pitchappa, and N. Wang, "Terahertz reconfigurable intelligent surfaces (riss) for 6g communication links," *Micromachines*, vol. 13, no. 2, p. 285, 2022.

[45] S. Dash, C. Psomas, I. Krikidis, I. F. Akyildiz, and A. Pitsillides, "Active control of thz waves in wireless environments using graphene-based ris," *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 10, pp. 8785–8797, 2022.

[46] S. Aboagye, A. R. Ndjiongue, T. M. Ngatched, and O. A. Dobre, “Design and optimization of liquid crystal ris-based visible light communication receivers,” *IEEE Photonics Journal*, vol. 14, no. 6, pp. 1–7, 2022.

[47] F. Saggese, V. Croisfelt, R. Kotaba, K. Stylianopoulos, G. C. Alexandropoulos, and P. Popovski, “On the impact of control signaling in ris-empowered wireless communications,” *IEEE Open Journal of the Communications Society*, vol. 5, pp. 4383–4399, 2024.

[48] C. Ozturk, M. F. Keskin, H. Wymeersch, and S. Gezici, “Ris-aided near-field localization under phase-dependent amplitude variations,” *IEEE Transactions on Wireless Communications*, vol. 22, no. 8, pp. 5550–5566, 2023.

[49] J.-B. Gros, V. Popov, M. A. Odit, V. Lenets, and G. Lerosey, “A reconfigurable intelligent surface at mmwave based on a binary phase tunable metasurface,” *IEEE Open Journal of the Communications Society*, vol. 2, pp. 1055–1064, 2021.

[50] R. Fara, P. Ratajczak, D.-T. Phan-Huy, A. Ourir, M. Di Renzo, and J. De Rosny, “A prototype of reconfigurable intelligent surface with continuous control of the reflection phase,” *IEEE Wireless Communications*, vol. 29, no. 1, pp. 70–77, 2022.

[51] D. F. Dinges, M. T. Orne, and E. C. Orne, “Assessing performance upon abrupt awakening from naps during quasi-continuous operations,” *Behavior Research Methods, Instruments, & Computers*, vol. 17, no. 1, pp. 37–45, 1985.

[52] C. Liaskos *et al.*, “Integrating software-defined metasurfaces into wireless communication systems: Design and prototype evaluation,” in *2022 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 2022, pp. 1–6.

[53] A. Papadopoulos *et al.*, “On modeling the ris as a resource: Multi-user allocation and efficiency-proportional pricing,” *IEEE Transactions on Network and Service Management*, pp. 1–1, 2025.

[54] S. I. Raptis *et al.*, “Synthesis of polarization-independent perfect anomalous reflectors via modulated metasurfaces and an analytical design model,” *IEEE Transactions on Microwave Theory and Techniques*, 2023.

[55] A. I. Papadopoulos *et al.*, "Physics-informed metaheuristics for fast ris codebook compilation," *IEEE Communications Magazine*, vol. 62, no. 11, pp. 152–158, 2024.

[56] X. Wei *et al.*, "Codebook design and beam training for extremely large-scale ris: Far-field or near-field?" *China Communications*, vol. 19, no. 6, pp. 193–204, 2022.

[57] M. Nerini *et al.*, "Discrete-value group and fully connected architectures for beyond diagonal reconfigurable intelligent surfaces," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 12, pp. 16 354–16 368, 2023.

[58] A. Papadopoulos *et al.*, "Ris as a network resource: User multiplexing and pricing algorithms," in *2024 3rd International Conference on 6G Networking (6GNet)*, 2024, pp. 144–152.

[59] Q. Wu *et al.*, "Intelligent reflecting surface enhanced wireless network: Joint active and passive beamforming design," in *2018 IEEE global communications conference (GLOBECOM)*. IEEE, 2018, pp. 1–6.

[60] H. Alwazani *et al.*, "Intelligent reflecting surface-assisted multi-user miso communication: Channel estimation and beamforming design," *IEEE Open Journal of the Communications Society*, vol. 1, pp. 661–680, 2020.

[61] C. You, B. Zheng, and R. Zhang, "Fast beam training for irs-assisted multiuser communications," *IEEE Wireless Communications Letters*, vol. 9, no. 11, pp. 1845–1849, 2020.

[62] P. Wang, J. Fang, W. Zhang, and H. Li, "Fast beam training and alignment for irs-assisted millimeter wave/terahertz systems," *IEEE Transactions on Wireless Communications*, vol. 21, no. 4, pp. 2710–2724, 2022.

[63] X. Liu, Q. Wu, D. Hu, R. Wang, and J. Wu, "Hierarchical codebook design and analytical beamforming solution for irs-assisted communication," *IEEE Transactions on Wireless Communications*, vol. 23, no. 8, pp. 8924–8938, 2024.

[64] S. Lv *et al.*, "RIS-aided near-field MIMO communications: Codebook and beam training design," *IEEE Transactions on Wireless Communications*, vol. 23, no. 9, pp. 12 531–12 546, 2024.

[65] P. Zheng *et al.*, “Mutual coupling in RIS-aided communication: Model training and experimental validation,” *IEEE Transactions on Wireless Communications*, vol. 23, no. 11, pp. 17 174–17 188, 2024.

[66] S. Lv *et al.*, “Ris-aided near-field mimo communications: Codebook and beam training design,” *IEEE Transactions on Wireless Communications*, vol. 23, no. 9, pp. 12 531–12 546, 2024.

[67] A. Abdallah *et al.*, “Deep reinforcement learning based beamforming codebook design for ris-aided mmwave systems,” in *2023 IEEE 20th Consumer Communications & Networking Conference (CCNC)*, 2023, pp. 1020–1026.

[68] H. Yang *et al.*, “Implicit neural representations for codebook configuration in ris-aided communication systems,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 12, pp. 2829–2843, 2026.

[69] J. Chen *et al.*, “A practical design of ris with subsurfaces using codebook-based configuration method,” *IEEE Wireless Communications Letters*, vol. 13, no. 4, pp. 1098–1102, 2024.

[70] S. Abeywickrama *et al.*, “Intelligent reflecting surface: Practical phase shift model and beamforming optimization,” *IEEE Transactions on Communications*, vol. 68, no. 9, pp. 5849–5863, 2020.

[71] S. Luke, *Essentials of Metaheuristics*, 2nd ed. Lulu, 2013.

[72] B. Sheen *et al.*, “A deep learning based modeling of reconfigurable intelligent surface assisted wireless communications for phase shift configuration,” *IEEE Open Journal of the Communications Society*, vol. 2, pp. 262–272, 2021.

[73] Z. Yu *et al.*, “Environment-aware codebook design for ris-assisted mu-miso communications: Implementation and performance analysis,” 2024.

[74] C. Avin, K. Mondal, and S. Schmid, “Demand-aware network design with minimal congestion and route lengths,” *IEEE/ACM Transactions on Networking*, vol. 30, no. 4, pp. 1838–1848, 2022.

[75] A. Petrosino *et al.*, “Dynamic management of forwarding rules in a t-sdn architecture with energy and bandwidth constraints,” in *Ad-Hoc, Mobile, and*

Wireless Networks: 19th International Conference on Ad-Hoc Networks and Wireless, ADHOC-NOW 2020, Bari, Italy, October 19–21, 2020, Proceedings 19. Springer, 2020, pp. 3–15.

[76] Y. Abiko, T. Saito, D. Ikeda, K. Ohta, T. Mizuno, and H. Mineno, "Flexible resource block allocation to multiple slices for radio access network slicing using deep reinforcement learning," *IEEE Access*, vol. 8, pp. 68 183–68 198, 2020.

[77] F. Fossati, S. Moretti, P. Perny, and S. Secci, "Multi-resource allocation for network slicing," *IEEE/ACM Transactions on Networking*, vol. 28, no. 3, pp. 1311–1324, 2020.

[78] I. Chang, T. Ji, R. Zhu, Z. Wu, C. Li, and Y. Jiang, "Towards an efficient and dynamic allocation of radio access network slicing resources for 5g era," *IEEE Access*, 2023.

[79] M. Leconte, G. S. Paschos, P. Mertikopoulos, and U. C. Kozat, "A resource allocation framework for network slicing," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 2177–2185.

[80] Z. Jahedi and T. Kunz, "Fast and cost-efficient virtualized network function placement algorithm in wireless multi-hop networks," in *Ad-Hoc, Mobile, and Wireless Networks: 19th International Conference on Ad-Hoc Networks and Wireless, ADHOC-NOW 2020, Bari, Italy, October 19–21, 2020, Proceedings 19*. Springer, 2020, pp. 23–36.

[81] C. Liaskos *et al.*, "Resource management for programmable metasurfaces: Concept, prospects and challenges," *IEEE Communications Magazine*, vol. 61, no. 11, pp. 208–214, 2023.

[82] Y. Zhang *et al.*, "Optimal resource allocation for reconfigurable intelligent surface assisted dynamic wireless network via online reinforcement learning," in *2022 IEEE International Conference on Sensing, Communication, and Networking (SECON Workshops)*. IEEE, 2022, pp. 13–18.

[83] P. S. Aung *et al.*, "Deep reinforcement learning based joint spectrum allocation and configuration design for star-ris-assisted v2x communications," *IEEE Internet of Things Journal*, 2023.

[84] Y. Sun *et al.*, "Coalition formation game for resource allocation in d2d uplink underlaying cellular networks," *IEEE Communications Letters*, vol. 23, no. 5, pp. 888–891, 2019.

[85] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, and I. Akyildiz, "A new wireless communication paradigm through software-controlled metasurfaces," *IEEE Communications Magazine*, vol. 56, no. 9, pp. 162–169, 2018.

[86] Y. Zhang and H. Xu, "Two-stage online reinforcement learning based distributed optimal resource allocation for multiple ris-assisted mobile ad-hoc network," in *2023 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 2023, pp. 563–567.

[87] S. Ao *et al.*, "Resource allocation for ris-assisted device-to-device communications in heterogeneous cellular networks," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 9, pp. 11 741–11 755, 2023.

[88] M. Zeng *et al.*, "Resource allocation for ris-empowered wireless communications: Low-complexity and robust designs," *IEEE Wireless Communications*, vol. 31, no. 5, pp. 68–75, 2024.

[89] A. Khaleel *et al.*, "Reconfigurable intelligent surface-empowered mimo systems," *IEEE Systems Journal*, vol. 15, no. 3, pp. 4358–4366, 2021.

[90] J. Huang *et al.*, "Reconfigurable intelligent surfaces: Channel characterization and modeling," *Proceedings of the IEEE*, vol. 110, no. 9, pp. 1290–1311, 2022.

[91] C. Liaskos *et al.*, "A scheduling framework for performing resource slicing with guarantees in 6g ris-enabled smart radio environments," *ITUJournal*, vol. 4, no. 1, pp. 33–49, 2023.

[92] S. Lin *et al.*, "Outage probability analysis of ris-assisted wireless powered multiuser communications," in *2021 13th International Conference on Wireless Communications and Signal Processing (WCSP)*, 2021, pp. 1–5.

[93] J. Chen *et al.*, "Channel estimation for reconfigurable intelligent surface aided multi-user mmwave mimo systems," *IEEE Transactions on Wireless Communications*, vol. 22, no. 10, pp. 6853–6869, 2023.

[94] D. Tyrovolas *et al.*, "Performance analysis of cascaded reconfigurable intelligent surface networks," *IEEE Wireless Communications Letters*, vol. 11, no. 9, pp. 1855–1859, 2022.

[95] A. Huang *et al.*, "Latency-aware computation offloading in multi-ris-assisted edge networks," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 1204–1221, 2024.

[96] A. Taneja *et al.*, "Leveraging reconfigurable intelligent surfaces for task offloading in edge iot networks," *IEEE Internet of Things Journal*, vol. 12, no. 3, pp. 2422–2429, 2025.

[97] F. Naeem and et al., "Security and privacy for reconfigurable intelligent surface in 6G: A review of prospective applications and challenges," *IEEE Open J. Commun. Soc.*, vol. 4, pp. 1196–1217, 2023.

[98] X. Lu and et al., "Intelligent reflecting surface enabled covert communications in wireless networks," *IEEE Netw.*, vol. 34, no. 5, pp. 148–155, 2020.

[99] U. Altun and E. Basar, "RIS-enabled secure communication with covert constraint," in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 685–689.

[100] X. Chen and et al., "Enhancing MIMO covert communications via intelligent reflecting surface," *IEEE Wireless Commun. Lett.*, vol. 11, no. 1, pp. 33–37, 2022.

[101] S. Pejoski and et al., "Full-duplex covert communications assisted by intelligent reflective surfaces," *IEEE Commun. Lett.*, vol. 26, no. 12, pp. 2846–2850, 2022.

[102] L. Lv and et al., "Covert communication in intelligent reflecting surface-assisted NOMA systems: Design, analysis, and optimization," *IEEE Trans. Wireless Commun.*, vol. 21, no. 3, pp. 1735–1750, 2022.

[103] C. Liaskos and et al., "Software-defined reconfigurable intelligent surfaces: From theory to end-to-end implementation," *Proc. IEEE*, vol. 110, no. 9, pp. 1466–1493, 2022.

[104] ——, “A novel communication paradigm for high capacity and security via programmable indoor wireless environments in next generation wireless systems,” *Ad Hoc Netw.*, vol. 87, pp. 1–16, 2019.

[105] H. Xiao and et al., “Simultaneously transmitting and reflecting RIS (STAR-RIS) assisted multi-antenna covert communications: Analysis and optimization,” *IEEE Trans. Wireless Commun.*, vol. 22, no. 12, pp. 10 674–10 689, 2023.

[106] X. Li and et al., “Covert communication of STAR-RIS aided NOMA networks,” *IEEE Trans. Veh. Technol.*, vol. 73, no. 6, pp. 9055–9060, 2024.

[107] M. Zhu and et al., “Active-RIS-aided covert communications in NOMA-inspired ISAC wireless systems,” *IEEE Commun. Lett.*, vol. 28, no. 3, pp. 679–683, 2024.

[108] X. Kuai and et al., “Covert communications in active-RIS-aided NOMA systems,” *Electronics*, vol. 13, no. 14, p. 2882, 2024.

[109] Y. Gao and et al., “Multi-RIS aided covert communications: A multi-agent reinforcement learning approach,” *IEEE Trans. Cogn. Commun. Netw.*, 2024, early Access.

[110] A. Pitilakis and et al., “On the mobility effect in UAV-mounted absorbing metasurfaces: A theoretical and experimental study,” *IEEE Access*, vol. 11, pp. 79 777–79 792, 2023.

[111] J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed. London, U.K.: McGraw-Hill, 1996.

[112] Z. Wan *et al.*, “Terahertz massive mimo with holographic reconfigurable intelligent surfaces,” *IEEE Transactions on Communications*, vol. 69, no. 7, pp. 4732–4750, 2021.

[113] P. Mursia, S. Phang, V. Sciancalepore, G. Gradoni, and M. Di Renzo, “Saris: Scattering aware reconfigurable intelligent surface model and optimization for complex propagation channels,” *IEEE Wireless Communications Letters*, vol. 12, no. 11, pp. 1921–1925, 2023.

[114] T. Liebig, A. Rennings, S. Held, and D. Erni, “openems–a free and open source equivalent-circuit (ec) fdtd simulation platform supporting cylindrical coordinates suitable for the analysis of traveling wave mri applications,” *International*

Journal of Numerical Modelling: Electronic Networks, Devices and Fields, vol. 26, no. 6, pp. 680–696, 2013.

[115] R. Liu, M. Li, Y. Liu, Q. Wu, and Q. Liu, “Joint transmit waveform and passive beamforming design for ris-aided dfrc systems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 5, pp. 995–1010, 2022.

[116] Z. Wang, X. Mu, and Y. Liu, “Stars enabled integrated sensing and communications,” *IEEE Transactions on Wireless Communications*, vol. 22, no. 10, pp. 6750–6765, 2023.

[117] A. Magbool, V. Kumar, Q. Wu, M. Di Renzo, and M. F. Flanagan, “A survey on integrated sensing and communication with intelligent metasurfaces: Trends, challenges, and opportunities,” *IEEE Open Journal of the Communications Society*, 2025.

[118] S. Zhang, R. Liu, M. Li, W. Wang, and Q. Liu, “Target detection in isac systems with active riss: A multi-perspective observation approach,” *IEEE Transactions on Cognitive Communications and Networking*, pp. 1–1, 2025.

[119] V. Loscrí, C. Rizza, A. Benslimane, A. M. Vegni, E. Innocenti, and R. Giuliano, “BEST-RIM: a mmWave beam steering approach based on computer vision-enhanced reconfigurable intelligent metasurfaces,” *IEEE Transactions on Vehicular Technology*, 2023.

[120] C. A. Balanis, *Antenna Theory: Analysis and Design*. John Wiley & Sons, 2016.

[121] C. Liaskos and et al., “The internet of metamaterial things and their software enablers,” *Int. Telecommun. Union J.*, vol. 1, no. 1, pp. 55–77, 2020.

[122] E. W. Dijkstra *et al.*, *A discipline of programming*. prentice-hall Englewood Cliffs, 1976, vol. 613924118.

APPENDIX

APPENDIX A

WORKFLOW OF THE MATCH ALGORITHM

Algorithm A.1 (MATCH: Physics-based RIS Codebook Compilation Algorithm). *Consider Tx position* $\mathbf{t}$, *focus points* $\{\mathbf{c}^{(\ell)}\}_{\ell=1}^{L}$, *RIS with* $N = N_d \times N_d$ *elements and spacing* d, *and element coordinates* $\{\mathbf{p}_n\}_{n=1}^{N}$. *The goal is to construct the codebook* $\mathcal{C} = \{\Gamma^{\star}(\ell)\}_{\ell=1}^{L}$.

1. *Set hyper-parameters: frequency* f, *wavenumber* $k = 2\pi/\lambda$, *phase step* $\delta\phi$, *local/final tolerances* $\varepsilon_{\text{loc}}, \varepsilon_{\text{fin}}$, *influence threshold* τ, *number of candidates* G, *sphere radius* r, *sample counts* $N_{\mathcal{F}}, N_{\mathcal{O}}$.
2. *Initialize* $\mathcal{C} \leftarrow \emptyset$.
3. *For each focus index* $\ell = 1, \dots, L$:
    (a) *For each element* $n = 1, \dots, N$, *set* $\Gamma_n \leftarrow \exp(-\mathrm{i}k\|\mathbf{p}_n - \mathbf{c}^{(\ell)}\|)$.
    (b) *Apply local refinement:* $(\Gamma, C) \leftarrow \textsc{LocalRefine}(\Gamma, \varepsilon_{\text{loc}})$.
    (c) *Apply global search:* $\Gamma \leftarrow \textsc{GlobalSearch}(\Gamma, C, \tau, G)$.
    (d) *Apply final local refinement:* $(\Gamma^{\star}, -) \leftarrow \textsc{LocalRefine}(\Gamma, \varepsilon_{\text{fin}})$.
    (e) *Update codebook:* $\mathcal{C} \leftarrow \mathcal{C} \cup \{\Gamma^{\star}\}$.
4. *Output* $\mathcal{C}$.

Algorithm A.2 (MATCH: Auxiliary Functions). **LocalRefine**(Γ, ε): *Repeat*

1. *For each* $n = 1, \dots, N$:
    (a) *Update* $\Gamma_n \leftarrow \Gamma_n e^{\mathrm{i}\delta\phi}$.

```
            (b) Compute $(\bar{E}_{\mathcal{F}}^{\text{new}}, -) \leftarrow \textsc{Evaluate}(\Gamma)$.
            (c) If $\bar{E}_{\mathcal{F}}^{\text{new}} < \bar{E}_{\mathcal{F}}$, revert $\Gamma_n$; otherwise set $\bar{E}_{\mathcal{F}} \leftarrow \bar{E}_{\mathcal{F}}^{\text{new}}$.
            (d) Set $C_n \leftarrow \dfrac{\partial \bar{E}_{\mathcal{F}}}{\partial \phi_n}$.
        until $|\Delta \bar{E}_{\mathcal{F}}| \leq \varepsilon$, then return $(\Gamma, C)$.

GlobalSearch$(\Gamma, C, \tau, G)$:    1. Define $\mathcal{I}_{\text{freeze}} \leftarrow \{n : C_n < \tau\}$.
        2. Generate $G$ candidates $\{\Gamma^{(g)}\}_{g=1}^{G}$ by mutating phases of $n \notin \mathcal{I}_{\text{freeze}}$.
        3. For each $g = 1, \ldots, G$, compute $(\bar{E}_{\mathcal{F}}^{(g)}, \bar{E}_{\mathcal{O}}^{(g)}) \leftarrow \textsc{Evaluate}(\Gamma^{(g)})$.
        4. Return the selected $\Gamma$.

Evaluate$(\Gamma)$:    1. Compute $E(\mathbf{r}; \Gamma)$.
        2. Sample $\mathcal{F}$ (inside distance $r$ of each $\mathbf{c}^{(\ell)}$) and $\mathcal{O}$ (outside all such spheres).
        3. Set $\bar{E}_{\mathcal{F}} \leftarrow \dfrac{1}{N_{\mathcal{F}}} \sum_{j=1}^{N_{\mathcal{F}}} |E(\mathbf{r}_j)|$.
        4. Set $\bar{E}_{\mathcal{O}} \leftarrow \dfrac{1}{N_{\mathcal{O}}} \sum_{j=1}^{N_{\mathcal{O}}} |E(\mathbf{r}'_j)|$.
        5. Return $(\bar{E}_{\mathcal{F}}, \bar{E}_{\mathcal{O}})$.
```

Appendix B

Workflow of the OSIRIS Strategy

Algorithm B.1 (Optimized Strategy Instantiating RIS Codebook Entries (OSIRIS)). *Consider two height levels* z_1, z_2 *with dense training grids on each plane; an RIS of size* $N_r \times N_c$ *(with* $N_{\text{RIS}} = N_r N_c$*); number of clusters* K*; shortlist size* K_{top}*; user positions* $\{\mathbf{x}_\ell\}_{\ell=1}^{L}$*; and* $\varepsilon > 0$*. The goal is to compute an offline codebook* $\{(\mathbf{C}_k, \boldsymbol{\mu}_k)\}_{k=1}^{K}$ *and, online, a per-user selected configuration* $\boldsymbol{\Gamma}_{\text{MAP}}(\mathbf{x}_\ell)$ *for each user.*

1. ***Phase 1 (Offline) – MATCH and Database Creation****:*
    - *(a) Define a dense set of training points* $\{\mathbf{x}_n\}$ *on the two* z*-planes covering the region of interest.*
    - *(b) For each training point* $\mathbf{x}_n$*:*
        - *i. Compute via MATCH the optimal RIS coefficients* $\boldsymbol{\Gamma}_n^{\star} \in \mathbb{C}^{N_{\text{RIS}}}$ *such that* $|\boldsymbol{\Gamma}_{n,i}^{\star}| = 1$ *for all* i*.*
    - *(c) Store the database*

    $$\mathcal{D} \leftarrow \{(\mathbf{x}_n, \boldsymbol{\Gamma}_n^{\star})\}.$$

2. ***Phase 2 (Offline) – Clustering and Codebook Construction****:*
    - *(a) For each database entry* $(\mathbf{x}_n, \boldsymbol{\Gamma}_n^{\star}) \in \mathcal{D}$*:*
        - *i. Extract phases* $\phi_n \leftarrow \angle \boldsymbol{\Gamma}_n^{\star}$*.*
        - *ii. Map to unit phasors* $\mathbf{u}_n \leftarrow \exp(j\phi_n) \in \mathbb{C}^{N_{\text{RIS}}}$*.*
    - *(b) Apply* K*-means to* $\{\mathbf{u}_n\}$ *using distance* $d(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|_2$*, yielding cluster memberships* $\{\mathcal{C}_k\}_{k=1}^{K}$ *and centroids* $\{\mathbf{m}_k\}_{k=1}^{K}$*.*

(c) *For each cluster* $k = 1, \dots, K$*:*

   i. *Form the codebook entry by unit-modulus normalization:*

$$\mathbf{C}_k \leftarrow \mathbf{m}_k \oslash \big(|\mathbf{m}_k| + \varepsilon\big).$$

   ii. *Compute the anchor (cluster mean location):*

$$\boldsymbol{\mu}_k \leftarrow \frac{1}{|\mathcal{C}_k|} \sum_{n \in \mathcal{C}_k} \mathbf{x}_n.$$

(d) *Store the offline codebook*

$$\{(\mathbf{C}_k, \boldsymbol{\mu}_k)\}_{k=1}^{K}.$$

3. ***Online Inference – Configuration Selection for*** L ***Users****:*

(a) *For each user* $\ell = 1, \dots, L$ *with position* $\mathbf{x} \leftarrow \mathbf{x}_\ell$*:*

   i. *For each codebook entry* $k = 1, \dots, K$*, compute the anchor distance*

$$d_k \leftarrow \|\mathbf{x} - \boldsymbol{\mu}_k\|.$$

   ii. *Select the shortlist* $\mathcal{K}_{\text{top}}(\mathbf{x})$ *containing the* K_{top} *smallest values of* $\{d_k\}$*.*

   iii. *For each* $k \in \mathcal{K}_{\text{top}}(\mathbf{x})$*, evaluate the received SNR using* $\mathbf{C}_k$*, i.e.,* $\mathrm{SNR}(\mathbf{x}; \mathbf{C}_k)$*.*

   iv. *Select the maximum-a-posteriori configuration index and output*

$$\boldsymbol{\Gamma}_{\text{MAP}}(\mathbf{x}_\ell) \leftarrow \arg \max_{k \in \mathcal{K}_{\text{top}}(\mathbf{x})} \mathrm{SNR}(\mathbf{x}; \mathbf{C}_k).$$

(b) *Output the set of selected configurations*

$$\mathbf{C}_L \leftarrow \{\boldsymbol{\Gamma}_{\text{MAP}}(\mathbf{x}_\ell)\}_{\ell=1}^{L}.$$

Appendix C

Workflow of the COMMON Algorithm

Algorithm C.1 (Codebook Multiplexing Algorithm (COMMON)). *Consider an RIS with dimensions* (M, N)*, number of users* K*, discretization parameter* N_d*, and per-user codebook entries* CE_k*. The goal is to compute a common RIS configuration* $CC(M, N)$*.*

1. *For each user* $k = 1, \ldots, K$*:*
    (a) *Compute the discretized entries* $DE_k \leftarrow f_{\text{dis}}(CE_k, N_d)$*.*
2. *Initialize the common configuration* $CC(M, N)$*.*
3. *For each element* (m, n) *of* CC*:*
    (a) *Initialize an empty list of values.*
    (b) *For each user* $k = 1, \ldots, K$*:*
        i. *Append* $DE_k(m, n)$ *to the list of values.*
    (c) *Set the element value as the most frequent one in the list:*

$$CC(m, n) \leftarrow \text{ComputeMostFrequent}(\textit{values}).$$

4. *Output the common configuration* CC*.*

APPENDIX D

WORKFLOW OF THE PRIME ALGORITHM

Algorithm D.1 (PRIcing-based Multiplexing of codebook Entries (PRIME)). *Consider an RIS with dimensions (M, N), number of users K, discretization parameter N_d, per-user codebook entries CE_k, and payment factors PF_k. The goal is to compute a common RIS configuration $CC(M, N)$.*

1. ***Discretization**: For each user $k = 1, \ldots, K$:*
    - *(a) Compute the discretized entries $DE_k \leftarrow f_{\text{dis}}(CE_k, N_d)$.*
2. ***Replica creation based on payment factors**:*
    - *(a) Initialize an empty collection of replicas.*
    - *(b) For each user $k = 1, \ldots, K$:*
        - *i. Append DE_k to the replicas PF_k times.*
3. *Initialize the common configuration $CC(M, N)$.*
4. ***Element-wise decision**: For each element (m, n) of CC:*
    - *(a) Collect all values taken at position (m, n) across the replicas.*
    - *(b) Set*

$$CC(m,n) \leftarrow \text{ComputeMostFrequent}(\textit{values at position } (m,n)).$$

5. *Output the common configuration CC.*

APPENDIX E

WORKFLOW OF THE LEVEL ALGORITHM

Algorithm E.1 (LEVEL: MuLtiplexing codebook Entries with controllable performance deviation among multiple users' leVELs). *Consider user counts per level* $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$*, an RIS of size* $M \times N$*, codebook matrices* $\mathbf{CE}$*, and level weights* w_1, w_2, w_3 *with* $w_1 > w_2 > w_3$*. Optionally, consider a quantization operator* $\mathcal{Q}(\cdot)$*. The goal is to compute a final common RIS configuration* CC*, a resource slice per user* $\{r_k\}$*, and average resource slices per level.*

1. ***Initialization***
   (a) *Combine the codebook matrices into a single matrix* $\mathbf{CE} \in \mathbb{R}^{MN \times K}$.
   (b) *Assign a weight vector* $\mathbf{w} \in \mathbb{R}^K$ *to the* K *users according to their levels, using the level weights* w_1, w_2, w_3.
2. ***Parallel pool initialization (implementation-dependent)***
   (a) *Initialize a parallel computing pool for element-wise optimization if available.*
3. ***Element-wise optimization (can run in parallel)***
   (a) *For each RIS element index* $i = 1, \ldots, MN$ *(in parallel if supported):*
      i. *Extract the* i*-th row vector* $\mathbf{CE}_i \in \mathbb{R}^K$ *from* $\mathbf{CE}$.
      ii. *Define the objective function*

$$f_i(\sigma) = \frac{\sum_{k=1}^{K} w_k \left|\sigma - \mathbf{CE}_i(k)\right|}{\sum_{k=1}^{K} w_k}.$$

iii. *Set the feasible bounds for* σ *based on the entries of* $\mathbf{CE}_i$ *(or, equivalently, the* i*-th row of* $\mathbf{CE}$*).*

iv. *Optionally choose an initial guess* $\sigma^{(0)}$*.*

v. *Compute the optimizer*

$$\sigma_i^\star = \arg\min_\sigma f_i(\sigma).$$

vi. *Update the optimized configuration entry* $\mathbf{CC}_{\text{opt}}(i) \leftarrow \sigma_i^\star$*.*

4. ***Optional quantization***

(a) *If* $\mathcal{Q}(\cdot)$ *is enabled, quantize the optimized configuration:*

$$\mathbf{CC}_{\text{opt}} \leftarrow \mathcal{Q}\big(\mathbf{CC}_{\text{opt}}\big).$$

(b) *Otherwise, keep* $\mathbf{CC}_{\text{opt}}$ *unchanged.*

5. ***Resource-slice computation***

(a) *For each user* $k = 1, \ldots, K$*:*

i. *Compute the deviation*

$$\Delta_{fk} \leftarrow \frac{1}{MN} \sum_{i=1}^{MN} \big|\mathbf{CC}_{\text{opt}}(i) - \mathbf{CE}(i,k)\big|.$$

ii. *Set the resource slice*

$$r_k \leftarrow 1 - \Delta_{fk}.$$

(b) *From the set* $\{r_k\}_{k=1}^K$*, compute the average resource slice for each user level (using the membership given by* $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$*).*

6. ***Output***

(a) *Set* $CC \leftarrow \mathbf{CC}_{\text{opt}}$*.*

(b) *Output the final common RIS configuration* CC*, the per-user resource slices* $\{r_k\}_{k=1}^K$*, and the level-wise average resource slices.*

APPENDIX F

WORKFLOW OF THE THEMIS FRAMEWORK

Algorithm F.1 (THEMIS: A Novel Framework for Fair Resource Allocation in RIS-enabled Networks). *The THEMIS framework operates on a RIS-enabled network graph with multiple RIS nodes and user pairs, and returns common RIS configurations, user routes, and fairness metrics.*

Input

- *Set of RIS nodes* $\mathcal{V}$.
- *User pairs* $\{(\mathrm{userSrc}(k), \mathrm{userDst}(k))\}_{k=1}^{K}$.
- *Codebook entries* CE_k *for each user* k.
- *Thresholds* $\eta_{\min}$ *and* $\mathrm{LoS}_{\mathrm{th}}$.
- *Pricing policy used to determine* v_k *for all users* k.

Output

- *Common RIS configuration(s) CC for all nodes.*
- *User paths* $\{\pi_k\}_{k=1}^{K}$.
- *Per-user and overall network fairness metrics.*

1. ***Initialization***. *Define the RIS-enabled network graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ *based on* (6.6) *and the line-of-sight threshold* $\mathrm{LoS}_{\mathrm{th}}$. *Set the (uniform) capacity proxy*

$$C \leftarrow \left\lceil \frac{K}{|\mathcal{V}|} \right\rceil .$$

2. ***Congestion-aware user routing.*** *Perform iterative routing for* $t = 1, \ldots,$ `maxRouteIter`*:*
    - *For each user* k*, compute the routing cost according to* (6.7)*.*
    - *Determine the path* π_k *via* (6.8)*.*
    - *Update the edge-usage (or congestion) matrix using the chosen path* π_k*.*
    - *If a convergence criterion is satisfied (for example, paths no longer change or costs stabilize), terminate the routing loop early.*
3. ***RIS configuration and adaptation.*** *For each RIS node in* $\mathcal{V}$*:*
    - *Aggregate the user-specific RIS configurations* CE_k *that traverse this node and replicate them according to the pricing weights* v_k*.*
    - *Derive the common configuration* CC *for the node using* (6.1)*.*
    - *Compute the user efficiencies* ϵ'_k/ϵ_k *for the users served by this node.*
    - *If at least one user* k *violates the minimum efficiency constraint, that is* $\epsilon'_k/\epsilon_k < \eta_{\min}$*, then:*
        - *For each candidate sub-panel in the candidate set* `candG`*, update the weight vector using*

          $$w^{(t+1)} \leftarrow (1-\alpha)\, w^{(t)} + \alpha \cdot \left(\frac{w^{(t)} v_k}{\sum_j w_j^{(t)} v_j} \right).$$
        - *Recompute the common configuration* CC *and the associated efficiencies.*
        - *If all users now satisfy* $\epsilon'_k/\epsilon_k \geq \eta_{\min}$*, stop exploring further candidate sub-panels for this node.*
4. ***Final metrics and output.*** *Compute the unified fairness metric* UFM *according to* (6.3)*. Output the final common configurations* CC*, the user paths* $\{\pi_k\}_{k=1}^{K}$*, and the metric* UFM *(along with any per-user fairness measures, if required).*

APPENDIX G

WORKFLOW OF THE SHIELD ALGORITHM

Algorithm G.1 (SHIELD: RF-Fencing for Signal Suppression and Signal Delivery Area Creation). *The SHIELD algorithm takes as input a codebook database and design parameters, and outputs an optimized RIS phase profile and performance metrics for Friendly Signal Delivery Areas (FSDAs) and Hostile Signal Suppression Areas (HSSAs).*

Input

- *Codebook database* $\{E_k\}$ *for supported angles of arrival/departure (AoA/AoD).*
- *Numbers of desired FSDAs and HSSAs (denoted by* d *and* u*).*
- *Tolerance (learning rate)* μ *and weight parameter* w_{opt}.
- *FSDA and HSSA thresholds* $\tau_{\text{FSDA}}, \tau_{\text{HSSA}}$.
- *Compromise factor* η.

Output

- *Final phase profile* $\mathbf{\Phi}_{\text{opt}}$.
- *Performance metrics for FSDAs and HSSAs.*

1. ***Dominant field identification***
   (a) *For each case* k *in the codebook database:*
      i. *Identify the dominant field values in* E_k.

  (b) *Form the FSDA and HSSA masks* $\mathcal{M}_d$ *and* $\mathcal{M}_u$ *using the procedure in* (7.1).

2. ***Common field computation***

   (a) *Compute the common field* $\boldsymbol{E}_{\text{common}}$ *according to* (7.2).

   (b) *Initialize the phase profile via*

$$\boldsymbol{\Phi}_{\text{init}} \leftarrow \boldsymbol{\Phi}_{\text{common}}$$

   as defined in (7.3).

3. ***Online optimization***

   (a) *Define the cost function* $\mathcal{J}(\boldsymbol{\Phi})$ *using the FSDA/HSSA terms in* (7.4) *and the total cost in* (7.5).

   (b) *Update the phase profile by gradient descent:*

$$\boldsymbol{\Phi}_{\text{opt}} = \boldsymbol{\Phi}_{\text{init}} - \mu \, \nabla_{\boldsymbol{\Phi}} \mathcal{J}(\boldsymbol{\Phi}).$$

4. ***Final output***

   (a) *Set the final phase profile* $\boldsymbol{\Phi}_{\text{opt}}$ *as obtained from the optimization step.*

   (b) *Compute the FSDA and HSSA performance metrics using* (7.8).

   (c) *Output* $\boldsymbol{\Phi}_{\text{opt}}$ *and the corresponding performance metrics.*

AUTHOR'S PUBLICATIONS

Parts of this thesis have been published in various journals and conference proceedings; we present these works in detail in the following.

Conferences

1. A. Papadopoulos et al., ”An Open Platform for Simulating the Physical Layer of 6G Communication Systems with Multiple Intelligent Surfaces,” 2022 18th International Conference on Network and Service Management (CNSM), Thessaloniki, Greece, 2022, pp. 359-363, doi: 10.23919/CNSM55787.2022.9965098

2. A. Papadopoulos, A. Lalas, K. Votis, S. Schmid, K. Katsalis and C. Liaskos, ”RIS as a Network Resource: User Multiplexing and Pricing Algorithms,” 2024 3rd International Conference on 6G Networking (6GNet), Paris, France, 2024, pp. 144-152, doi: 10.1109/6GNet63182.2024.10765689.

3. Papadopoulos, A. I., Tyrovolas, D., Pitilakis, A., Diamantoulakis, P. D., Lalas, A., Votis, K., ... & Liaskos, C. (2025, September). SHIELD: A Codebook-Based Methodology for RIS-Based Covert Communications. In 2025 IEEE 36th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC) (pp. 1-6). IEEE.

4. Papadopoulos, A. I., Pistela, M. A., Tyrovolas, D., Lalas, A., Votis, K., Ioannidis, S., ... & Liaskos, C. (2026). Physics-Aware RIS Codebook Compilation for Near-Field Beam Focusing under Mutual Coupling and Specular Reflections. arXiv preprint arXiv:2601.12982. Accepted in IEEE ICC 2026.

5. Papadopoulos, A. I., Tyrovolas, D., Lalas, A., Votis, K.,Karagiannidis, G. & Liaskos, C. (2026). OSIRIS: A Strategy for Online Codebook-based RIS Configuration. Accepted in IEEE EuCNC 2026.

6. Papadopoulos, A. I., Lalas, A., Votis, K., Tassiulas, L. & Liaskos, C. (2026). A Novel Framework for Fair Resource Allocation in RIS-enabled Networks. Accepted as full paper in IEEE WoWMoM 2026.

7. Papadopoulos, A. I., Tyrovolas, D., Tsimpoukis, S., Lalas, A., Votis, K., Ioannidis, S., Karagiannidis, G. & Liaskos, C. K. (2026). Experimental Evaluation of RIS Configuration Strategies for Multi-user Communications. Accepted in ICT 2026.

Journals

1. A. Papadopoulos et al., "On Modeling the RIS as a Resource: Multi-User Allocation and Efficiency-Proportional Pricing," in IEEE Transactions on Network and Service Management, doi: 10.1109/TNSM.2025.3576038

2. A. I. Papadopoulos et al., "Physics-Informed Metaheuristics for Fast RIS Codebook Compilation," in IEEE Communications Magazine, vol. 62, no. 11, pp. 152-158, November 2024, doi: 10.1109/MCOM.001.2300582.

3. Papadopoulos, A. I., Lalas, A., Votis, K., Tassiulas, L., & Liaskos, C. K. (2026). Modeling RIS for Multi-User Resource Allocation under Pricing-Aware Policies. IEEE Access.

Moreover, additional research contributions by the author include:

1. S. I. Raptis et al., "An Accurate Semi-Analytical Model for Periodic Tunable Metasurfaces Electromagnetic Response," 2024 18th European Conference on Antennas and Propagation (EuCAP), Glasgow, United Kingdom, 2024, pp. 1-5, doi: 10.23919/EuCAP60739.2024.10501453.

2. P. -V. Mekikis et al., "Dynamic Programmable Wireless Environment with UAV-mounted Static Metasurfaces," 2022 IEEE Conference on Standards for Communications and Networking (CSCN), Thessaloniki, Greece, 2022, pp. 101-104, doi: 10.1109/CSCN57023.2022.10051011.

3. A. Pitilakis et al., "On the Mobility Effect in UAV-Mounted Absorbing Metasurfaces: A Theoretical and Experimental Study," in IEEE Access, vol. 11, pp. 79777-79792, 2023, doi: 10.1109/ACCESS.2023.3299379

4. Segata, M., Casari, P., Lestas, M., Papadopoulos, A., Tyrovolas, D., Saeed, T., ... & Liaskos, C. (2024). Cooperis: A framework for the simulation of reconfigurable

intelligent surfaces in cooperative driving environments. Computer Networks, 248, 110443.

Finally, parts of this thesis have been submitted for publication in various journals and conference proceedings; we present the works that are in peer review in detail in the following.

Conferences

1. Papadopoulos, A. I., Ntougias, K., Lalas, A., Votis, K., Krikidis, I., & Liaskos, C. K. (2026). On the Physics-Aware Active Element Selection for Hybrid RIS Configuration. Submitted in GLOBECOM 2026.

Journals

1. Papadopoulos, A. I., Kopsinos, S., Tyrovolas, D., Lalas, A., Votis, K., Karagiannidis, G. & Liaskos, C. K. (2026). On RIS-assisted Networks Orchestration: A Physics-aware Framework for Resource Allocation, Energy Efficiency & Admission Control. Submitted in IEEE TNSM.

2. Papadopoulos, A. I., Tyrovolas, D., Lalas, A., Votis, K., Ioannidis, S., Karagiannidis, G. & Liaskos, C. K. (2026). Towards Real-World Multifunctional Programmable Metasurfaces: Physics-Driven Compilation and Orchestration. Submitted in IEEE JSAC.

3. Papadopoulos, A. I., Tyrovolas, D.,Pitilakis, A., Diamantoulakis, D., Lalas, A., Votis, K.,Kantartzis, N., Ioannidis, S. & Liaskos, C. K. (2026). RF-Fencing: A Novel RIS-Based Service for Signal Suppression via Codebook Multiplexing. Submitted in IEEE OJCOMS.

SHORT BIOGRAPHY

Alexandros Papadopoulos received the M.Eng. Diploma in Electrical and Computer Engineering from the Aristotle University of Thessaloniki (AUTH), Thessaloniki, Greece, in 2019. He is currently pursuing the Ph.D. degree in the Department of Computer Science and Engineering at the University of Ioannina, Ioannina, Greece, focusing on Reconfigurable Intelligent Surfaces and programmable wireless environments. Since 2020, he has been a Research Assistant and Wireless Communication Researcher at the Information Technologies Institute of the Centre for Research and Technology Hellas (CERTH-ITI), Thessaloniki, Greece. His research interests include 5G/6G communications, RIS-assisted networks, near-field electromagnetic modeling and codebook compilation, multi-user resource allocation under fairness and pricing policies, covert communications, and jamming detection/mitigation.