Design and implementation of multi-user MIMO precoding algorithms

Maitane Barrenechea Carrasco

Supervisor:

Mikel Mendicute Errasti

MONDRAGON
UNIBERTSITATEA

A thesis submitted for the degree of

Doctor por Mondragon Unibertsitatea

Department of Electronics and Computer Science

Mondragon Goi Eskola Politeknikoa

Mondragon Unibertsitatea

November 2011

Aita, ama, Mikel [2] *eta Jonentzat*

The most exciting phrase to hear in science,
the one that heralds new discoveries,
is not 'Eureka!' but rather
'hmmm... that's funny...'.

Isaac Asimov

Agradecimientos

Hace cuatro largos años desde que me inicié en este loco mundo de las communicaciones inalámbricas. En todo este tiempo, he tenido el placer de conocer a un gran número de investigadores, los cuales han contribuído, cada uno a su manera, al desarrollo de este trabajo de investigación.

En primer lugar, me gustaría agradecer a mi director de tesis **Mikel Mendicute** el apoyo recibido durante este largo camino así como su gran labor de asesoramiento y orientación gracias a las cuales he podido dar forma a este trabajo de investigación. También quisiera agradecer a **Mondragon Goi Eskola Politeknikoa** por la beca Etortek concedida a través de la Fundación Centros Tecnológicos-Iñaki Goenaga.

Quisiera dar mi más profundo agradecimiento a **John Thompson** del *Institute for Digital Communications* (IDCOM) de la *University of Edinburgh* por brindarme la oportunidad de colaborar con él cuando todavía era una recién llegada al mundo de las comunicaciones multiantena. Asimismo, me gustaría agradecer a todos los miembros del IDCOM su cordial trato, en especial a **Mostafa Afgani**, **Zubin Bharucha**, **Nick Johnson**, **Sharad Nagappa**, **Nikola Serafimovski**, **Sinan Sinanović** y **Ki Won Sung** por hacerme sentir como en casa en las frías tierras escocesas.

Debo expresar mi más cálida gratitud a **Luis Barbero** del *Department of Electronics, Communications and Information Technology* de la *Queen's University Belfast* por enseñarme los entresijos del mundo de la implementación en *hardware* y por su inestimable apoyo en estos últimos años.

Por otro lado, me gustaría agradecer a **Michael Joham** del *Fachgebiet Methoden der Signalverarbeitung* de la *Teschnische Universität München* por compartir sus amplios conocimientos técnicos conmigo. Gracias también a mis compañeros de departamento durante mi estancia en Múnich, en especial a **Leonardo Baltar**, **Mario Castañeda**, **Lennart Gerdes**, **Christian Guthy**, **Raphael Hunger** y **David Schmidt** por amenizar los duros días de trabajo con un agradable café en la Glyptothek.

Asimismo, quiero dar las gracias a **Andreas Burg** del *Telecommunications Circuits Laboratory* (TCL) de la *École Polytechnique Fédérale de Lausanne* (EPFL) por su labor de

orientador y por los interesantes debates técnicos de los que tanto he aprendido. Mi más cálida gratitud también a los chicos del TCL, especialmente a **Filippo Borlenghi**, **Jeremy Constantin**, **Nicholas Preyss** y **Christian Senning**. El alegre y distendido ambiente de trabajo de nuestra oficina ha supuesto una gran motivación en los últimos y estresantes meses del desarrollo de esta tesis.

Quisiera dar las gracias también a mis compañeros y amigos doctorandos de la Universidad de Mondragon por su apoyo constante y por las siempre agradables, y a veces interminables, pausas de café en las que tantas y tan divertidas historias hemos contado y escuchado. Gracias por lo tanto a **Joxe Aizpurua**, **Maite Beamurgia**, **Lorea Belategi**, **Iñaki Garitano**, **Idoia Jiménez**, **Aritz Legarda**, **Lorena Martínez**, **Pello Otxandiano** y al gran **Iker Sobrón** por hacer más llevadero el día a día de mis años de doctoranda.

Finalmente, me gustaría expresar mi más sentida gratitud a mis **padres**, **hermanos** y **amigas de la cuadrilla**, cuyos ánimos me han dado fuerzas en los momentos de flaqueza. Y por supuesto, gracias a **Mikel** por su apoyo incondicional y por alentarme a seguir siempre adelante.

Acknowledgments

It has already been four years since I entered this crazy world of wireless communications. During all this time, I have had the pleasure of meeting many researchers which have, each on its own way, contributed to the development of this research work.

First of all, I would like to thank my thesis supervisor **Mikel Mendicute** for his support during this long journey as well as for the invaluable counseling and guidance that have enabled me to carry out this research work. Thanks also to **Mondragon Goi Eskola Politeknikoa** for funding my research through a *Fundación Centros Tecnológicos-Iñaki Goenaga* Etortek grant.

I would like to express my deepest gratitude to **John Thompson** from the Institute for Digital Communications (IDCOM) at the University of Edinburgh for giving me the chance of collaborating with him even when I was still a newcomer to the world of multi-antenna communications. I would also like to thank all the IDCOM members for their hospitality, specially **Mostafa Afgani**, **Zubin Bharucha**, **Nick Johnson**, **Sharad Nagappa**, **Nikola Serafimovski**, **Sinan Sinanović** and **Ki Won Sung** for making me feel at home in the cold Scottish land.

A very special thanks goes to **Luis Barbero** from the department of Electronics, Communications and Information Technology at the Queen's University Belfast for showing me the ins and outs of the world of hardware implementation and also for his continuous and invaluable support.

I also gratefully acknowledge **Michael Joham** from the *Fachgebiet Methoden der Signalverarbeitung* at the *Technische Universität München* (TUM) for sharing his wide technological knowledge with me. Thanks also to my PhD student colleagues at the TUM, specially to **Leonardo Baltar**, **Mario Castañeda**, **Lennart Gerdes**, **Christian Guthy**, **Raphael Hunger** and **David Schmidt** for making the hard working days more amenable with a nice coffee break at the Glyptothek.

I would also like to thank **Andreas Burg** from the Telecommunications Circuits Laboratory (TCL) at the *École Polytechnique Fédérale de Lausanne* (EPFL) for his guidance and for the great collaboration that we had. My most heartfelt thanks also to the guys

at TCL, specially to **Filippo Borlenghi**, **Jeremy Constantin**, **Nicholas Preyss** and **Christian Senning**. The warm and relaxed working atmosphere has greatly contributed to my motivation in the stressful last months of the development of this thesis.

I would like to express my gratitude also to my fellow colleagues and friends at *Mondragon Goi Eskola Politeknikoa* for their constant support and for the always enjoyable, and sometimes endless, coffee breaks in which so many funny stories we have shared. Thanks therefore to **Joxe Aizpurua**, **Maite Beamurgia**, **Lorea Belategi**, **Iñaki Garitano**, **Idoia Jiménez**, **Aritz Legarda**, **Lorena Martínez**, **Pello Otxandiano** and to the great **Iker Sobrón** for making the working days nicer and easier to take.

Finally, on a more personal note, I would like to express my deepest gratitude to my **parents**, **brothers** and **closest friends** whose encouragement has given me strength in the moments of weakness. And of course, thanks to **Mikel** for his unconditional support and understanding.

Abstract

The demand for high-speed communications required by cutting-edge applications has put a strain on the already saturated wireless spectrum. The incorporation of antenna arrays at both ends of the communication link has provided improved spectral efficiency and link reliability to the inherently complex wireless environment, thus allowing for the thriving of high data-rate applications without the cost of extra bandwidth consumption. As a consequence to this, multiple-input multiple-output (MIMO) systems have become the key technology for wideband communication standards both in single-user and multi-user setups.

The main difficulty in single-user MIMO systems stems from the signal detection stage at the receiver, whereas multi-user downlink systems struggle with the challenge of enabling non-cooperative signal acquisition at the user terminals. In this respect, precoding techniques perform a pre-equalization stage at the base station so that the signal at each receiver can be interpreted independently and without the knowledge of the overall channel state. Vector precoding (VP) has been recently proposed for non-cooperative signal acquisition in the multi-user broadcast channel. The performance advantage with respect to the more straightforward linear precoding algorithms is the result of an added perturbation vector which enhances the properties of the precoded signal. Nevertheless, the computation of the perturbation signal entails a search for the closest point in an infinite lattice, which is known to be in the class of non-deterministic polynomial-time hard (NP-hard) problems.

This thesis addresses the difficulties that stem from the perturbation process in VP systems from both theoretical and practical perspectives. On one hand, the asymptotic performance of VP is analyzed assuming optimal decoding. Since the perturbation process hinders the analytical assessment of the VP performance, lower and upper bounds on the expected data rate are reviewed and proposed. Based on these bounds, VP is compared to linear precoding with respect to the performance after a weighted sum rate optimization, the power resulting from a quality of service (QoS) formulation, and the performance when balancing the user rates.

On the other hand, the intricacies of performing an efficient computation of the perturbation vector are analyzed. This study is focused on tree-search techniques that, by means

of an strategic node pruning policy, reduce the complexity derived from an exhaustive search and yield a close-to-optimum performance. To that respect, three tree-search algorithms are proposed. The fixed-sphere encoder (FSE) features a constant data path and a non-iterative architecture that enable the parallel processing of the set of vector hypotheses and thus, allow for high-data processing rates. The sequential best-node expansion (SBE) algorithm applies a distance control policy to reduce the amount of metric computations performed during the tree traversal. Finally, the low-complexity SBE (LC-SBE) aims at reducing the complexity and latency of the aforementioned algorithm by combining an approximate distance computation model and a novel approach of variable run-time constraints.

Furthermore, the hardware implementation of non-recursive tree-search algorithms for the precoding scenario is also addressed in this thesis. More specifically, the hardware architecture design and resource occupation of the FSE and K-Best fixed-complexity tree-search techniques are presented. The determination of the ordered sequence of complex-valued nodes, also known as the Schnorr-Euchner enumeration, is required in order to select the nodes to be evaluated during the tree traversal. With the aim of minimizing the hardware resource demand of such a computationally-expensive task, a novel non-sequential and low-complexity enumeration algorithm is presented, which enables the independent selection of the nodes within the ordered sequence. The incorporation of the proposed enumeration technique along with a fully-pipelined architecture of the FSE and K-Best approaches, allow for data processing throughputs of up to 5 Gbps in a 4×4 antenna setup.

Resumen

La demanda de comunicaciones de alta velocidad requeridas por las aplicaciones más vanguardistas ha impuesto una presión sobre el actualmente saturado espectro inalámbrico. La incorporación de *arrays* de antenas en ambos extremos del enlace de comunicación ha proporcionado una mayor eficiencia espectral y fiabilidad al inherentemente complejo entorno inalámbrico, permitiendo así el desarrollo de aplicaciones de alta velocidad de transmisión sin un consumo adicional de ancho de banda. Consecuentemente, los sistemas *multiple-input multiple output* (MIMO) se han convertido en la tecnología clave para los estándares de comunicación de banda ancha, tanto en las configuraciones de usuario único como en los entornos multiusuario.

La principal dificultad presente en los sistemas MIMO de usuario único reside en la etapa de detección de la señal en el extremo receptor, mientras que los sistemas multiusuario en el canal de bajada se enfrentan al reto de habilitar la adquisición de datos no cooperativa en los terminales receptores. A tal efecto, las técnicas de precodificación realizan una etapa de pre-ecualización en la estación base de tal manera que la señal en cada receptor se pueda interpretar independientemente y sin el conocimiento del estado general del canal. La precodificación vectorial (VP, del inglés *vector precoding*) se ha propuesto recientemente para la adquisición no cooperativa de la señal en el canal de difusión multiusuario. La principal ventaja de la incorporación de un vector de perturbación es una considerable mejora en el rendimiento con respecto a los métodos de precodificación lineales. Sin embargo, la adquisición de la señal de perturbación implica la búsqueda del punto más cercano en un reticulado infinito. Este problema se considera de complejidad no determinística en tiempo polinomial o *NP-complejo*.

Esta tesis aborda las dificultades que se derivan del proceso de perturbación en sistemas VP desde una perspectiva tanto teórica como práctica. Por un lado, se analiza el rendimiento de VP asumiendo una decodificación óptima en escenarios de alta relación señal a ruido. Debido a que el proceso de perturbación dificulta la evaluación analítica del rendimiento de los sistemas de VP, se proponen y revisan diversas cotas superiores e inferiores en la tasa esperada de transmisión de estos sistemas. En base a estas cotas, se realiza una comparación

de VP con respecto a la precodificación lineal en el ámbito de la capacidad suma ponderada, la potencia resultante de una formulación de calidad de servicio y el rendimiento obtenido al equilibrar las tasas de transmisión de los usuarios.

Por otro lado, se han propuesto nuevos procedimientos para un cómputo eficiente del vector de perturbación. Estos métodos se basan en técnicas de búsqueda en árbol que, por medio de diferentes políticas de podado, reducen la complejidad derivada de una búsqueda exhaustiva y obtienen un rendimiento cercano al óptimo. A este respecto, se proponen tres algoritmos de búsqueda en árbol. El *fixed-sphere encoder* (FSE) cuenta con una complejidad constante y una arquitectura no iterativa, lo que permite el procesamiento paralelo de varios vectores candidatos, lo que a su vez deriva en grandes velocidades de procesamiento de datos. El algoritmo iterativo denominado *sequential best-node expansion* (SBE) aplica una política de control de distancias para reducir la cantidad de cómputo de métricas realizadas durante la búsqueda en árbol. Por último, el *low-complexity SBE* (LC-SBE) tiene por objetivo reducir la complejidad y latencia del algoritmo anterior mediante la combinación de un modelo de cálculo aproximado de distancias y una estrategia novedosa de restricción variable del tiempo de ejecución.

Adicionalmente, se analiza la implementación en *hardware* de algoritmos de búsqueda en árbol no iterativos para los escenarios de precodificación. Más específicamente, se presentan el diseño de la arquitectura y la ocupación de recursos de *hardware* de las técnicas de complejidad fija FSE y K-Best. La determinación de la secuencia ordenada de nodos de naturaleza compleja, también conocida como la enumeración de Schnorr-Euchner, es vital para seleccionar los nodos evaluados durante la búsqueda en árbol. Con la intención de reducir al mínimo la demanda de recursos de *hardware* de esta tarea de alta carga computacional, se presenta un novedoso algoritmo no secuencial de baja complejidad que permite la selección independiente de los nodos dentro de la secuencia ordenada. La incorporación de la técnica de enumeración no secuencial junto con la arquitectura *fully-pipeline* de los algoritmos FSE y K-Best, permite alcanzar velocidades de procesamiento de datos de hasta 5 Gbps para un sistema de 4 antenas receptoras.

Laburpena

Aplikazio abangoardistek beharrezko duten abiadura handiko komunikazioen eskaerak presio handia ezarri du dagoeneko saturatuta dagoen haririk gabeko espektruan. Komunikazio loturaren bi muturretan antena *array*-en erabilerak eraginkortasun espektral eta fidagarritasun handiagoez hornitu du berez konplexua den haririk gabeko ingurunea, modu honetan banda zabalera gehigarririk gabeko abiadura handiko aplikazioen garapena ahalbidetuz. Honen ondorioz, *multiple-input multiple output* (MIMO) sistemak banda zabaleko komunikazio estandarren funtsezko teknologia bihurtu dira, erabiltzaile bakarreko ezarpenetan hala nola erabiltzaile anitzeko inguruneetan.

Erabiltzaile bakarreko MIMO sistemen zailtasun garrantzitsuena hartzailean ematen den seinalearen detekzio fasean datza. Erabiltzaile anitzeko sistemetan, aldiz, erronka nagusiena datu jasotze ez kooperatiboa bermatzea da. Prekodifikazio teknikek hartzaile bakoitzaren seinalea kanalaren egoera orokorraren ezagutzarik gabe eta modu independiente baten interpretatzea ahalbidetzen dute estazio nagusian seinalearen pre-ekualizazio fase bat inposatuz. Azken aldian, prekodifikazio bektoriala (VP, ingelesez *vector precoding*) proposatu da erabiltzaile anitzeko igorpen kanalean seinalearen eskuratze ez kooperatiboa ahalbidetzeko. Perturbazio seinale baten erabilerak, prekodifikatutako seinalearen ezaugarriak hobetzeaz gain, errendimenduaren hobekuntza nabarmen bat lortzen du prekodifikazio linearreko teknikekiko. Hala ere, perturbazio seinalearen kalkuluak sare infinitu baten puntu hurbilenaren bilaketa suposatzen du. Problema honen ebazpenaren konplexutasuna denbora polinomialean ez deterministikoa dela jakina da.

Doktoretza tesi honen helburu nagusia VP sistemetan perturbazio prozesuaren ondorioz ematen diren zailtasun teoriko eta praktikoei irtenbide egoki bat ematea da. Alde batetik, seinale/zarata ratio handiko ingurunetan VP sistemen errendimendua aztertzen da, beti ere deskodetze optimoa ematen dela suposatuz. Perturbazio prozesuak VP sistemen errendimenduaren azterketa analitikoa oztopatzen duenez, data transmisio tasaren hainbat goi eta behe borne proposatu eta berrikusi dira. Borne hauetan oinarrituz, VP eta prekodifikazio linealaren arteko errendimendu desberdintasuna neurtu da hainbat aplikazio ezberdinen eremuan. Konkretuki, kanalaren ahalmen ponderatua, zerbitzu kalitatearen formulazio baten

ondorioz esleitzen den seinale potentzia eta erabiltzaileen datu transmisio tasa orekatzean lortzen den errendimenduaren azterketa burutu dira.

Beste alde batetik, perturbazio bektorearen kalkulu eraginkorra lortzeko metodoak ere aztertu dira. Analisi hau zuhaitz-bilaketa tekniketan oinarritzen da, non egitura sinple baten bitartez errendimendu ia optimoa lortzen den. Ildo horretan, hiru zuhaitz-bilaketa algoritmo proposatu dira. Alde batetik, *Fixed-sphere encoder*-aren (FSE) konplexutasun konstateak eta arkitektura ez errekurtsiboak datu prozesaketa abiadura handiak lortzea ahalbidetzen dute. *Sequential best-node expansion* (SBE) delako algoritmo iteratiboak ordea, distantzia kontrol politika baten bitartez metrika kalkuluen kopurua murriztea lortzen du. Azkenik, *low-complexity SBE* (LC-SBE) algoritmoak SBE metodoaren latentzia eta konplexutasuna murriztea lortzen du ordezko distantzien kalkuluari eta exekuzio iraupenean ezarritako muga aldakorreko metodo berri bati esker.

Honetaz gain, prekodifikazio sistementzako zuhaitz-bilaketa algoritmo ez errekurtsiboen *hardware* inplementazioa garatu da. Zehazki, konplexutasun finkoko FSE eta K-Best algoritmoen arkitektura diseinua eta *hardware* baliabideen erabilera landu dira. Balio konplexuko nodoen sekuentzia ordenatua, Schnorr-Euchner zerrendapena bezala ezagutua, funtsezkoa da zuhaitz bilaketan erabiliko diren nodoen aukeraketa egiteko. Prozesu honek beharrezkoak dituen *hardware* baliabideen eskaera murrizteko, konplexutasun bajuko algoritmo ez sekuentzial bat proposatzen da. Metodo honen bitartez, sekuentzia ordenatuko edozein nodoren aukeraketa independenteki egin ahal da. Proposatutako zerrendapen metodoa eta estruktura *fully-pipeline* baten bitartez, 5 Gbps-ko datu prozesaketa abiadura lortu daiteke FSE eta K-Best delako algoritmoen inplementazioan.

Declaration of Originality

I hereby declare that the research recorded in this thesis and the thesis itself were developed entirely by myself at the Signal Theory and Communications Area, Department of Electronics and Computer Science, at the University of Mondragon, with the following exceptions:

- The resolution of the optimization problems in Chapter 3 has been the result of the collaboration work between Dr. Michael Joham and myself.
- The hardware structure of the FSE algorithm is based on an original model provided by Dr. Luis Barbero for the MIMO detection process.

Maitane Barrenechea Carrasco
Department of Electronics and Computer Science
Mondragon Goi Eskola Politeknikoa
Mondragon Unibertsitatea
November, 2011

Contents

Acknowledgments . . . iii
Abstract . . . vii
Declaration of Originality . . . xiii
Contents . . . xiv
List of Figures . . . xvii
List of Tables . . . xxi
List of Symbols . . . xxv

1 Introducción **1**
1.1 Motivación y Objetivos . . . 2
1.2 Contribuciones de la Tesis . . . 3
1.3 Estructura de la Tesis . . . 3

2 Background and Related Work **6**
2.1 Introduction . . . 6
2.2 MIMO . . . 6
2.2.1 Channel Capacity . . . 7
2.3 Multi-user MIMO . . . 8
2.3.1 Capacity Region of the Multiple-Access Channel . . . 10
2.3.2 Capacity Region of the Broadcast Channel . . . 11
2.4 Precoding Techniques . . . 12
2.4.1 Linear Precoding . . . 13
2.4.1.1 Zero-Forcing Precoding . . . 13
2.4.1.2 Wiener-Filter Precoding . . . 14
2.4.2 Non-linear Precoding . . . 14
2.4.2.1 Tomlinson-Harashima Precoding . . . 15
2.4.2.2 Vector Precoding . . . 18
2.5 Chapter Summary . . . 22

3 Asymptotic Sum Rate Analysis and Rate Optimizations for VP **23**
3.1 Introduction . . . 23
3.2 Concepts of Lattice Theory . . . 24
3.3 Bounds on the Unnormalized Transmit Power . . . 26
3.3.1 Lower Bounds . . . 28
3.3.2 Upper Bounds . . . 29
3.4 Bounded Sum Rate in VP Systems . . . 30
3.5 Optimization Problems for Asymptotic VP . . . 33
3.5.1 Weighted Sum Rate . . . 33
3.5.1.1 Linear Precoding . . . 34
3.5.1.2 Vector Precoding . . . 35
3.5.1.3 Ergodic Performance Gap . . . 36
3.5.1.4 Simulation Results . . . 37
3.5.2 Quality of Service . . . 38
3.5.2.1 Linear Precoding . . . 39
3.5.2.2 Vector Precoding . . . 39
3.5.2.3 Ratio of Ergodic Powers . . . 39
3.5.2.4 Simulation Results . . . 40
3.5.3 User Balancing . . . 41
3.5.3.1 Linear Precoding . . . 41
3.5.3.2 Vector Precoding . . . 42
3.5.3.3 Ergodic Performance Gap . . . 42
3.5.3.4 Simulation Results . . . 42
3.6 Chapter Summary . . . 43

4 The Fixed-complexity Sphere Encoder **45**
4.1 Introduction . . . 45
4.2 Tree-search Algorithms for VP . . . 46
4.2.1 Sphere Encoder for VP Systems . . . 48
4.2.2 K-Best Tree Traversal for VP . . . 50
4.3 The Fixed-complexity Sphere Encoder . . . 51
4.3.1 Matrix Preprocessing for the FSE Tree Search . . . 51
4.3.1.1 Ordering Strategy for THP (*Best-last* Rule) . . . 53
4.3.1.2 Performance of the FSE with Different Ordering Strategies . 53
4.3.2 FSE Tree Configuration Vector . . . 54
4.4 Simulation Results . . . 59
4.4.1 Computational Complexity . . . 59
4.4.1.1 Number of Evaluated Nodes . . . 59

4.4.1.2 Number of Operations 60
4.4.2 BER Performance 62
4.5 Chapter Summary 64

5 Complex-plane Enumeration for Precoding Systems 65
5.1 Introduction 65
5.2 Schnorr-Euchner Enumeration 66
5.2.1 Enumeration by Identification of the Admissible Set 67
5.2.2 Enumeration by Node Arrangement in Circular Subsets 67
5.2.3 Enumeration by Node Arrangement in Unidimensional Subsets 69
5.2.4 Enumeration by Neighbor Expansion 70
5.2.5 Other Suboptimum Enumeration Techniques 70
5.3 The Puzzle Enumerator 72
5.3.1 Initial Considerations 72
5.3.2 Low-complexity Enumeration Procedure 73
5.3.3 Unordered Puzzle Enumerator 74
5.4 Implementation of Complex-plane Enumerators 76
5.4.1 Efficient Distance Computation in State-of-the-art Enumerators . . . 79
5.4.2 Implementation of the Puzzle Enumerator 80
5.5 Comparative Analysis 82
5.5.1 Average Number of PED Computations 83
5.5.2 PED Cache Memory Size in Depth-first Systems 83
5.5.3 Latency of the Enumeration Process 84
5.5.4 Resource Occupation 85
5.6 Chapter Summary 86

6 Implementation of Fixed-complexity Algorithms for VP 88
6.1 Introduction 88
6.2 General Architecture Overview 89
6.3 DPU for the K-Best 90
6.3.1 Structure of the Sorting Stage 92
6.3.1.1 The Winner Path Extension Algorithm 92
6.3.2 Structure of the K-Best DPU 93
6.4 DPU for the FSE 94
6.5 Design Considerations 95
6.5.1 Choice of the Design Parameters 95
6.5.2 Implementation of an Approximate Norm 97
6.6 Implementation Results 100
6.7 Chapter Summary 100

7 Low-complexity Sequential Tree-search Algorithms for VP **102**
7.1 Introduction . . . 102
7.2 The SBE Algorithm . . . 103
7.2.1 SBE Tree-search Architecture . . . 103
7.2.1.1 Implementation Considerations . . . 104
7.2.2 Matrix Preprocessing for the SBE . . . 105
7.2.3 Performance Results . . . 106
7.3 The Low-complexity SBE Algorithm . . . 107
7.3.1 SBE with an Approximate Norm . . . 107
7.3.2 Flexible Run-time Constraint . . . 110
7.3.2.1 Performance Degradation in the $\ell^{\tilde{1}}$-SBE . . . 110
7.3.2.2 Flexible Run-time Constraint for the $\ell^{\tilde{1}}$-SBE . . . 111
7.3.3 Performance Results . . . 113
7.3.4 Implementation Considerations . . . 113
7.4 Simulation Results . . . 113
7.4.1 Computational Complexity . . . 114
7.4.1.1 Number of Evaluated Nodes . . . 114
7.4.1.2 Number of Operations . . . 115
7.4.2 BER Performance . . . 116
7.4.2.1 BER Performance under an Overall Run-time Constraint . . . 117
7.5 Chapter Summary . . . 118

8 Summary and Conclusions **120**
8.1 Thesis Contributions . . . 122
8.2 Suggestions for Further Research . . . 122

A Publications **124**

References **126**

List of Figures

2.1 Block diagram of an $N \times M$ MIMO system where M stands for the number of transmit antennas and N represents the number of receive antennas. . . . 8

2.2 Multi-user MIMO broadcast channel with M transmit antennas and N single-antenna user terminals. . . . 9

2.3 Capacity region of the two-user MAC. . . . 10

2.4 Dirty paper broadcast region for multi-user setups with single-antenna users. 11

2.5 Block diagram of a generic precoding system with M transmit antennas and N single-antenna users. . . . 12

2.6 Block diagram of a linear precoding system with M transmit antennas and N single-antenna users. . . . 13

2.7 BER performance of the ZF and WF linear precoding techniques in the 4×4 and 4×6 antenna setups with 16-QAM modulation. . . . 15

2.8 Block diagram of a THP system. . . . 16

2.9 Block diagram of a VP system. . . . 18

2.10 Illustration of the infinite repetition of the modulation constellation before precoding (a) and after (b). Figure (b) shows that it is possible to transmit the same symbol with a smaller power if a perturbation vector is added to the user data stream. . . . 19

2.11 BER performance of non-linear and linear precoding schemes. . . . 20

3.1 Congruent Voronoi regions and covering radius ϱ of an arbitrary lattice. . . . 25

3.2 Block diagram of the VP system used for the study of the unnormalized transmit power. . . . 27

3.3 Dimensionless second moment of several well-known lattices along with the *sphere bound* and the G_{CS} lower bound. . . . 29

3.4 Unnormalized transmit power E_{SE} along with the upper bounds (G_{TB} and G_{Rog}) and lower bounds (*sphere bound* and G_{CS}) derived for this parameter. The provided results are given as a function of the number of single-antenna users N. 31
3.5 Sum rate vs. SNR for asymptotic DPC, VP, and bounded asymptotic VP for an 8×8 antenna system. 32
3.6 Linear precoding system with power loading matrix $\boldsymbol{B}^{-\frac{1}{2}}$. 34
3.7 Ergodic performance gap on the sum rate for systems with $M = N$ and $M = N + 1$. 38
3.8 Ratio of ergodic powers of the linear and VP precoding approaches. 41
3.9 Ergodic performance gap for the user balancing optimization. 43

4.1 Incidence of the lattice values in the optimum solution vector for different SNR regimes. 47
4.2 Effect of employing a reduced search set on the BER performance of VP systems. 48
4.3 SE tree traversal example in a system with $N = 3$ and $|\mathcal{L}| = 3$. 49
4.4 K-Best tree traversal example in a system with $N = 3$, $|\mathcal{L}| = 3$ and $K = 4$. . 50
4.5 FSE tree traversal example in a system with $N = 3$, $|\mathcal{L}| = 3$ and tree configuration vector $\boldsymbol{n} = [1, 2, 2]$. 52
4.6 BER performance of the proposed FSE algorithm with $\boldsymbol{n} = [1, 1, 2, 5]$ and ordering strategies $\mathfrak{O}_{\mathrm{SQR}}$ and $\mathfrak{O}_{\mathrm{BL}}$. Additionally, the unordered case is included for completion. 55
4.7 BER performance of the FSE with different tree configuration vectors in a 6×6 (a) and an 8×8 (b) antenna setup. 56
4.8 Number of PED computations (C_{PED}) and z_i calculations (C_{IP}) required for several tree configuration vectors with $n_T = 24$. Additionally, the total amount of operations is depicted for each one of the FSE tree-search structures. 57
4.9 Amount of required multipliers for different tree configuration vectors in an FSE tree search. 58
4.10 90-percentile of the number of node expansions for the SE, FSE and K-Best tree-search algorithms in an 8×8 antenna system. The amount of expanded nodes per level is given in (a), whereas the total amount of evaluated nodes vs. SNR is considered in (b). 60
4.11 Total number of operations for the SE, K-Best and FSE tree-search approaches in an $N = 8$ user system. 61
4.12 BER performance curves of the FSE applied to VP systems in 4×4, 6×6 and 8×8 antenna setups. The total amount of evaluated branches has been set to 10, 12 and 24 for the aforementioned antenna configurations, respectively. 62

4.13 BER performance of the SE, the FSE with $n_T = 12$ and $n_T = 24$ and the K-Best with $K = 7$ and $K = 8$ in an 8×8 VP system. 63

5.1 Schnorr-Euchner ordered sequence of nodes in a real-valued model (a) and in the complex plane (b). 66
5.2 Identification of the admissible set of candidates by means of the sphere constraint. 68
5.3 Boundaries and node arrangement for the Schnorr-Euchner enumeration by circular subsets. 69
5.4 Node arrangement for the Schnorr-Euchner enumeration by unidimensional subsets. 70
5.5 Schnorr-Euchner enumeration by neighbor expansion. 71
5.6 Distribution of the intermediate points for several SNR values. 73
5.7 Fundaments of the proposed puzzle enumerator for (a) $\rho = 2$, (b) $\rho = 3$,(c) $\rho = 4$,(d) $\rho = 5$,(e) $\rho = 6$,(f) $\rho = 7$, (g) $\rho = 8$, (h) $\rho = 9$, (i) $\rho = 10$, (j) $\rho = 11$, (k) $\rho = 12$ and (l) $\rho = 13$. 75
5.8 Fundaments of the proposed unordered puzzle enumerator to be used along with the FSE for (a) $n_i = 4$, (b)$n_i = 5$,(c) $n_i = 6$,(d) $n_i = 7$,(e) $n_i = 8$,(f) $n_i = 9$, (g) $n_i = 10$ and (h) $n_i = 11$. 77
5.9 Block diagram of the full-sort enumerator. 78
5.10 Block diagram of the enumeration strategies where the sequencing of the nodes is based on the management of a candidate set $\mathcal{S}$. 79
5.11 Hardware implementation model of the proposed complex enumerator for the first eight enumerated values. 81
5.12 Values of the $\Xi_{\Re}$ and $\Xi_{\Im}$ increments for different regions of the $\tilde{\chi}_{\text{ext}}$ axis for the 2^{nd}, 3^{rd} and 4^{th} closest nodes in the real axis. 82
5.13 Average number of distance computations required by different enumeration techniques in an 8×8 SE tree traversal. 84

6.1 General hardware architecture of the fixed-complexity tree-search techniques for an $N = 4$ user system. 90
6.2 Block diagram of the K-Best and FSE DPU. 91
6.3 Illustrative example of the WPE distributed sorting algorithm for a system with $K = 3$ and $B = 4$. 93
6.4 Puzzle enumeration unit with selectable sequence index ϱ. 94
6.5 Number of required multipliers for the K-Best and FSE tree-search techniques as a function of the number of candidate branches. 97
6.6 BER performance of the implemented FSE and K-Best tree-search structures. 98

6.7 BER performance degradation introduced by approximating the ℓ^2 norm by the simplified $\ell^{\tilde{1}}$ and $\ell^{\tilde{\infty}}$ norms in the FSE ($n_T = 10$) (a) and K-Best ($K = 7$) (b) tree-search approaches. 99

7.1 Example of a SBE tree search in a 3 user system with $|\mathcal{L}| = 3$. 104
7.2 BER performance of the SBE with different ordering strategies (a) and the CDF of the nodes expanded at the root level during a SBE tree traversal (b). 106
7.3 Accumulated distances of the nodes expanded during a SBE tree traversal with the ℓ^2 norm (a) and with the approximated $\ell^{\tilde{1}}$ (b) and $\ell^{\tilde{\infty}}$ (c) norms. . . 107
7.4 CDF of the nodes evaluated at the root level of a SBE and LC-SBE tree traversal in an 8 antenna system. 108
7.5 BER performance of the LC-SBE and the SBE algorithm with the ℓ^2,$\ell^{\tilde{1}}$ and $\ell^{\tilde{\infty}}$ approaches for distance computation. 109
7.6 Index of the branch selected as the solution vector for the $\ell^{\tilde{1}}$ and ℓ^2-SBE (a), and for the ℓ^2-SBE and the proposed LC-SBE (b). 110
7.7 Run-time constraint values (φ_{TH}) for 4×4 and 8×8 antenna setups with $p = 0.9$. 112
7.8 90-percentile of the number of node expansions per level for the SE, FSE, K-Best, SBE and LC-SBE tree-search algorithms in an 8×8 antenna system. 114
7.9 90-percentile of node expansions for the SE, FSE, K-Best, SBE and LC-SBE in an 8×8 antenna system. 115
7.10 Total number of operations for the SE, K-Best, FSE, SBE and LC-SBE tree-search approaches on an $N = 8$ user system. A zoomed image of the lower section of the main figure is also included for a better appreciation of the complexity results of the SBE-based schemes. 116
7.11 BER performance curves for the SE, SBE and LC-SBE tree-search approaches for 4×4 and 8×8 antenna setups. 117
7.12 BER performance of various tree-search approaches with an overall run-time constraint of $\varpi = 32$ evaluated nodes for an 8×8 antenna setup. 118

List of Tables

3.1 Values for $(1-\eta)10^3$ used in the G_{TB} bound. 30

4.1 Computation of the upper-triangular matrix $\boldsymbol{U}$ with *best-last* ($\mathfrak{O}_{\mathrm{BL}}$) ordering. 54

5.1 Boundary lines to be evaluated for the first 13 nodes in the Schnorr-Euchner sequence. 76
5.2 Unordered sequence of preferred children for different values of n_i. 78
5.3 Amount of clock cycles required to obtain the ordered sequence of child nodes. 85
5.4 Device occupation and maximum achievable frequency for different complex enumerators with MCUs in a fully-pipelined scheme. 86
5.5 Device occupation and maximum achievable frequency for the proposed complex enumerators in a fully-pipelined scheme. 86

6.1 Hardware resource occupation and throughput of the tree-search architectures under study. 100

7.1 Computation of the upper-triangular matrix $\boldsymbol{U}$ with $\mathfrak{O}_{\mathrm{SBE}}$ ordering. 105

Acronyms

AM-GM	arithmetic mean - geometric mean
ASIC	application-specific integrated circuit
AWGN	additive white Gaussian noise
BC	broadcast channel
BER	bit error rate
CDF	cumulative distribution function
COSIC	conditioned ordered successive interference cancelation
DFE	decision-feedback equalization
DPC	dirty paper coding
DPU	distance processing unit
FPGA	field-programmable gate array
FSD	fixed-complexity sphere decoder
FSE	fixed-complexity sphere encoder
GE	gate equivalent
i.i.d.	independent and identically distributed
LC-SBE	low-complexity SBE
LLL	Lenstra-Lenstra-Lovász
LUT	look-up table

MAC	multiple access channel
MCU	metric computation unit
MIMO	multiple-input multiple-output
MMSE	minimum mean squared error
MNA	modified-norm algorithm
MSE	mean squared error
MSU	minimum-search unit
MU-MIMO	multi-user MIMO
NP-hard	non-deterministic polynomial-time hard
QAM	quadrature amplitude modulation
QoS	quality of service
QPSK	quadrature phase shift keying
RVD	real-valued decomposition
SBE	sequential best-node expansion
SD	sphere decoder
SE	sphere encoder
SINR	signal to interference-plus-noise ratio
SNR	signal-to-noise ratio
THP	Tomlinson-Harashima precoding
VLSI	very-large-scale integration
VP	vector precoding
WF	Wiener filter
WPE	winner path extension
ZF	zero forcing

List of Symbols

$\boldsymbol{a}$	Perturbation vector
$\mathbb{A}$	Constellation of the modulation in use
β	Transmit power scaling factor
$\boldsymbol{b}_n$	n^{th} basis vector
$\boldsymbol{C}_{\boldsymbol{x}}$	Covariance matrix of vector $\boldsymbol{x}$
$\mathbb{CZ}$	Set of Gaussian integers
$\mid \boldsymbol{A} \mid$	Determinant of matrix $\boldsymbol{A}$
$\boldsymbol{e}_i$	i^{th} column of the identity matrix
E_{SE}	Power of the precoded symbols or unnormalized transmit power
E_{Tr}	Transmit power constraint
$\text{E}[\cdot]$	Expectation
$\boldsymbol{F}$	Feedback precoding matrix
$\lfloor \cdot \rfloor$	Flooring operation
$F_n(\cdot)$	Schläfli's function
$G(\mathcal{R})$	Dimensionless second moment of polytope $\mathcal{R}$
$\Gamma(\cdot)$	Gamma function
γ	Euler-Mascheroni constant
$\boldsymbol{H}$	Channel matrix

$\boldsymbol{A}^{\dagger}$	Moore-Penrose pseudoinverse of matrix $\boldsymbol{A}$
$H(x)$	Entropy of random variable x
$\boldsymbol{A}^{H}$	Hermitian matrix of $\boldsymbol{A}$
h_{nm}	Channel gain between transmit antenna m and receive antenna n
$I(\mathcal{R})$	Normalized second moment of polytope $\mathcal{R}$
$\boldsymbol{I}_M$	$M \times M$ identity matrix
$\Im(s)$	Imaginary part of the complex variable s
$\Lambda_{\boldsymbol{M}}$	Lattice with generator matrix $\boldsymbol{M}$
M	Number of transmit antennas
$\mathrm{Mod}(\cdot)$	Modulo operator
N	Number of single-antenna users
$\boldsymbol{n}$	Tree configuration vector for the FSE
n_T	Number of candidate branches in the FSE
$\boldsymbol{P}$	Feedforward precoding matrix
$\boldsymbol{\Pi}$	Permutation matrix
$\psi(\cdot)$	Digamma function
$Q(\cdot)$	Nearest-neighbor quantizer
R	Hypersphere radius or sphere constraint
R_n	Rate of user n
$\Re(s)$	Real part of the complex variable s
$\boldsymbol{s}$	Data symbol column vector
$\hat{\boldsymbol{s}}$	Detected data symbol column vector
$\mathcal{S}_n$	n-dimensional sphere
τ	modulo constant
Θ	Covering density or *thickness* of the lattice

$\boldsymbol{A}^T$	Transpose of matrix $\boldsymbol{A}$
$\mathrm{Tr}(\boldsymbol{A})$	Trace of matrix $\boldsymbol{A}$
ϱ	Covering radius
ϱ_{eff}	Effective covering radius
V_n	Volume of an n-dimensional sphere of unitary radius
$\mathcal{V}(\Lambda_{\boldsymbol{M}}, 0)$	Fundamental Voronoi region of lattice $\Lambda_{\boldsymbol{M}}$
$\mathrm{Vol}(\mathcal{R})$	Volume of polytope $\mathcal{R}$
$\boldsymbol{x}$	Transmitted signal column vector
χ	Covering efficiency
$\boldsymbol{y}$	Received signal column vector
$\boldsymbol{w}$	Additive white Gaussian noise column vector

Capítulo 1
Introducción

Los canales inalámbricos *multiple-output multiple-input* (MIMO) han obtenido mucha atención en la última década debido a la notable mejora que ofrecen con respecto a los canales de antena única en términos de eficiencia espectral y fiabilidad. La gran mayoría de trabajos realizados hasta la fecha en torno a los sistemas MIMO se ha centrado en la problemática de las comunicaciones punto a punto. Sin embargo, en los últimos años se han desarrollado técnicas de comunicación MIMO multiusuario que consideran escenarios más complejos, pero más realistas, con múltiples terminales que deben compartir los recursos de tiempo, espacio, ancho de banda y potencia disponibles en una red inalámbrica.

Gran parte del esfuerzo dedicado hoy en día al estudio e innovación de técnicas multiantena se centra en los entornos MIMO multiusuario. Esto se debe a las ventajas inherentes de este tipo de sistema inalámbrico frente a su precedente de usuario único. Entre los principales beneficios de los entornos MIMO multiusuario destacan su mayor inmunidad a varias de las limitaciones en la propagación de los entornos MIMO de usuario único, tales como la pérdida de rango del canal ocasionada por la correlación de las antenas. Estas propiedades se obtienen gracias a la inherente separación entre los usuarios, y por lo tanto entre las antenas, lo que deriva en un reducido coeficiente de correlación. Asimismo, la propagación de visión directa, la cual degrada en gran medida la comunicación en entornos MIMO con multiplexación espacial, no supone un gran problema en los esquemas MIMO multiusuario. Por último, MIMO multiusuario permite obtener una ganancia por multiplexación espacial en la estación base sin la necesidad de que los equipos terminales dispongan de varias antenas. Esto permite la implementación de dispositivos terminales pequeños y de bajo coste, ya que la carga computacional y la complejidad de implementación se transladan a la estación base.

Sin embargo, la utilización de esquemas MIMO multiusuario conlleva la resolución de otros problemas inexistentes en el caso de las comunicaciones MIMO punto a punto. Gran parte de dicha problemática reside en el canal de difusión que comprende la transmisión originada en la estación base y dirigida a los usuarios. En este esquema de comunicaciones la adquisición no cooperativa de los datos de los usuarios es vital para conservar la escalabilidad del modelo. Esta tarea se lleva a cabo mediante la incorporación de técnicas de precodifi-

cación en la estación base, para las cuales es imprescindible el conocimiento instantáneo del estado del canal.

1.1 Motivación y Objetivos

La imposibilidad de cooperación entre terminales en la etapa de detección es la principal causa de que la mayoría de las técnicas de detección de los entornos MIMO de usuario único no se puedan implementar en los esquemas multiusuario. Sin embargo, al realizar un preprocesado de la señal en la estación base, es posible que los terminales receptores adquieran la información transmitida sin ningún tipo de cooperación o interacción entre ellos [Joham04a]. Éste es el concepto principal de las técnicas de precodificación.

La etapa de precodificación se puede realizar de forma lineal [Peel05][Joham05b], aplicando un simple filtro a la señal a transmitir, o de manera no lineal [Fischer02c][Hochwald05], añadiendo elementos de procesamiento no lineales. Estos últimos ofrecen un mejor rendimiento en términos de tasa de error, aunque su complejidad puede ser prohibitiva para ciertas aplicaciones reales.

La precodificación vectorial es una técnica no lineal que mejora considerablemente el rendimiento del sistema mediante la incorporación de un vector de perturbación en la etapa previa al filtrado lineal. A pesar de su excelente rendimiento, la aplicación de esta técnica de transmisión plantea ciertas dificultades. Por un lado, la adquisición de la señal de perturbación implica la búsqueda del punto más cercano en un reticulado infinito. Aunque a día de hoy se han propuesto varias técnicas para la resolución de este problema, el cómputo de la señal perturbadora sigue representando el principal cuello de botella para la implementación de sistemas de precodificación vectorial. Por otro lado, el operador de módulo requerido en los terminales para eliminar el efecto de la señal de perturbación dificulta el análisis teórico de este tipo de sistemas no lineales. Esto se debe a que el patrón estadístico del ruido gaussiano blanco es modificado a su paso por el operador de módulo, resultando en una señal cuya distribución no se corresponde con la bien conocida distribución gaussiana.

Estas dificultades en la implementación y estudio teórico de los sistemas de precodificación vectorial han motivado el planteamiento de los siguientes objetivos:

- Análisis teórico del rendimiento de los sistemas de precodificación vectorial mediante el estudio de las propiedades del reticulado.
- Desarrollo de algoritmos de búsqueda en árbol de baja complejidad y simple estructura.
- Implementación en *hardware* de baja complejidad de algoritmos de búsqueda en árbol para precodificadores de alta tasa de transmisión de datos.

1.2 Contribuciones de la Tesis

En esta sección se describen las contribuciones principales del trabajo de investigación desarrollado. Asimismo, se indican las publicaciones asociadas a las diferentes aportaciones:

- Análisis y obtención de cotas superiores e inferiores para el rendimiento de sistemas de precodificación vectorial en términos de tasa de transmisión de bits. Este trabajo ha sido publicado en [Barrenechea10c] y parcialmente en [Barrenechea10b].
- Resolución de los problemas de optimización y delimitación del rendimiento de los sistemas de precodificación vectorial en aplicaciones de calidad de servicio, maximización de la capacidad suma ponderada y equilibrado de la tasa de transmisión de los usuarios [Barrenechea10b].
- Diseño de un algoritmo de complejidad fija para la búsqueda en árbol del vector de perturbación en un sistema de precodificación vectorial. Este trabajo ha sido publicado en [Barrenechea09b] y [Barrenechea09a] para la variante de filtro de Wiener y en [Barrenechea09c] para la versión de inversión regularizada.
- Implementación en *hardware* de bajo coste de un algoritmo de búsqueda en árbol de complejidad fija y alto rendimiento. Los resultados de este trabajo se han publicado en [Barrenechea10a] y [Barrenechea11b].
- Desarrollo e implementación de un enumerador complejo de naturaleza no secuencial y gran simplicidad computacional para su uso en algoritmos de búsqueda en árbol aplicados a sistemas de precodificación [Barrenechea11d].
- Diseño e implementación de un precodificador vectorial basado en el algoritmo de búsqueda en árbol K-Best sin etapas de ordenamiento y con una alta velocidad de procesamiento de datos. Este trabajo ha sido publicado en [Barrenechea11a].
- Desarrollo de algoritmos iterativos de búsqueda en árbol para una implementación de baja complejidad en esquemas de precodificación vectorial. Este trabajo está pendiente de publicación [Barrenechea11c].

1.3 Estructura de la Tesis

La memoria de la tesis está estructurada en ocho capítulos, de los cuales el primero se ha dedicado a introducir la temática de la tesis y a presentar la motivación que ha llevado al autor a la realización de este trabajo. Asimismo, se han descrito los principales objetivos propuestos para el trabajo de investigación desarrollado.

En el Capítulo 2 se resumen las principales características de las tecnologías MIMO y MIMO multiusuario. Concretamente, se proporciona una perspectiva general de los límites de la capacidad en sistemas MIMO punto a punto, así como un análisis similar para la capacidad de los canales de bajada y de subida de los escenarios multiusuario. Adicionalmente, se define el modelo matemático del sistema utilizado a lo largo de la tesis y se describe el diseño de precodificadores lineales y no lineales, haciendo especial hincapié en el diseño de sistemas de precodificación vectorial, los cuales representan el pilar principal de la temática de esta tesis.

El Capítulo 3 se centra en la problemática de determinar el rendimiento de los sistemas de precodificación vectorial. Dada la relación entre la potencia de los símbolos precodificados y el rendimiento del sistema, se analizan y proponen varias cotas superiores e inferiores para este primer parámetro, con lo que se posibilita la delimitación de la capacidad suma en los sistemas de precodificación vectorial. Asimismo, se plantean diversos problemas de optimización cuya resolución es posible gracias a las cotas obtenidas previamente. Concretamente, se estudian los problemas de la capacidad suma ponderada, la optimización de la calidad de servicio y el equilibrado de las tasas de transmisión de los diferentes usuarios. Adicionalmente, se plantean estos problemas de optimización para un sistema de precodificación lineal. La comparación entre los resultados de este análisis y los previamente obtenidos para la precodificación vectorial servirán para demostrar analíticamente la superioridad de las técnicas de precodificación no lineales, así como para cuantificar la ganancia esperada al introducir algoritmos de procesamiento no lineales en un precodificador lineal.

En el Capítulo 4 se analizan diversos métodos para el cómputo del vector de perturbación en un sistema de precodificación vectorial. En un primer apartado, se describen los métodos más populares de búsqueda distribuída en árbol y se analiza la problemática de su implementación en *hardware*. A continuación, se propone un esquema de complejidad fija, cuya principal característica es la simplicidad de su árbol de búsqueda. El esquema propuesto es comparado a continuación en términos de tasa de error de bit, complejidad computacional y número de nodos evaluados por nivel con los algoritmos de búsqueda en árbol más notorios.

El Capítulo 5 retrata la importancia de la enumeración compleja en los sistemas de precodificación, independientemente del tipo de búsqueda en árbol que se realice. Inicialmente, se describe la problemática de la enumeración en el ámbito complejo y se realiza una revisión literaria de los algoritmos de enumeración más importantes. Seguidamente, se presenta un enumerador de naturaleza no secuencial cuya carga computacional es considerablemente menor que la de los algoritmos presentados en la bibliografía. Para demostrar las ventajas del enumerador propuesto, se ha llevado a cabo la implementación en *hardware* de todos los enumeradores estudiados y se han analizado los resultados de ocupación y latencia de los diferentes algoritmos.

En el siguiente capítulo, se da un paso más en la dirección de la implementación en

hardware y se realiza el diseño de la arquitectura de dos precodificadores vectoriales de complejidad fija: uno de ellos inspirado en el algoritmo de complejidad fija propuesto en un capítulo anterior, y el otro basado en una variante sin etapas de ordenamiento del algoritmo K-Best. Ambos esquemas son de naturaleza no iterativa y poseen una complejidad fija. Al igual que en el capítulo anterior, se realiza un estudio de los resultados de ocupación de ambos algoritmos y, adicionalmente, se muestra el rendimiento en términos de tasa de error de bit de los dos sistemas implementados.

El Capítulo 7 se centra en el diseño de algoritmos de búsqueda en árbol secuenciales para el cómputo del vector de perturbación en sistemas de precodificación vectorial. En este marco, se propone un esquema de estructura simplificada cuyo rendimiento en términos de tasa de error de bit es similar al óptimo establecido por el codificador esférico. Con el objetivo de obtener una reducción de la complejidad computacional adicional, se presenta un segundo algoritmo basado en el cálculo de distancias alternativo y en un novedoso esquema de poda variable. En la última sección de este capítulo, se compara el rendimiento en términos de tasa de error de bit y de nodos computados de los esquemas propuestos con otros algoritmos de búsqueda en árbol, tales como el codificador esférico y el FSE propuesto en el Capítulo 4.

Finalmente, el Capítulo 8 resume los resultados del trabajo de investigación realizado y reúne las principales conclusiones obtenidas. Adicionalmente, se ha incluído en el Apéndice A un listado completo de las publicaciones realizadas a lo largo del periodo de investigación.

Chapter 2
Background and Related Work

2.1 Introduction

A lot of attention has been drawn to multi-antenna technologies in recent years. The capacity increase promised by multiple-input multiple-output (MIMO) technologies [Telatar95, Foschini98] has enabled the implementation of cutting-edge communication applications with an elevated rate demand without the cost of extra bandwidth consumption. After a decade of research on MIMO technologies, the focus has shifted to multi-user environments where the antennas at one of the ends of the communication link are no longer co-located [Gesbert07]. The non-cooperative signal acquisition at the user terminals is enabled by a pre-equalization stage performed at the base station. This procedure, which is referred to as precoding, ensures that the signal at the receivers can be interpreted without the knowledge of the data streams directed to other users or any information on the overall channel state.

The aim of this chapter is to provide a theoretical background on multi-user MIMO precoding techniques. As a starting point, the first two sections will deal with the capacity features of MIMO and multi-user MIMO technologies. Next, the main characteristics of linear and non-linear precoding techniques will be reviewed. This latter section will set the ground for the main research topic of this dissertation: vector precoding (VP) in multi-user environments.

2.2 MIMO

The limited availability of radio frequency spectrum and the complex nature of the wireless scenario represent the limiting factors in the evolution of communication technologies. MIMO techniques represent one of the major breakthroughs in wireless communications, as the utilization of antenna arrays at both transmit and receive devices offers an increased data transmission rate and robustness of the communication link. The core idea of MIMO is the combination of the temporal and spatial domains, which allows to transform the usually undesired multipath propagation into a benefit for the user.

Due to the multiple available subchannels, MIMO enables a flexible trade-off between multiplexing and diversity gains. On one hand, transmit multiplexing allows for a linear increase in the transmitted data rate with respect to the rank of the channel for the same allocated bandwidth and power. Spatial diversity ensures the independence of each subchannel, generating different spatial signatures and allowing to multiplex the transmit data symbols over the set of transmit streams. On the other hand, diversity gain can be achieved by implementing space-time coding, therefore providing the link with an enhanced robustness. Hence, depending on the transmit strategy to be used, the utilization of multiple antennas can provide an increased data rate, higher link reliability or a combination of both advantages.

Nevertheless, the superior performance of MIMO technologies comes at the cost of an increased complexity of the channel estimation and detection algorithms, which need to process several data streams simultaneously. Furthermore, the amount of required hardware resources is also multiplied, which derives in an increased economic cost of the communication link.

The block diagram of a MIMO system with M transmit and N receive antennas is depicted in Figure 2.1. The input-output relationship of the narrowband system for a symbol period can be described by the following equation

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{w},$$

where $\boldsymbol{y} \in \mathbb{C}^{N\times 1}$ and $\boldsymbol{x} \in \mathbb{C}^{M\times 1}$ represent the received and transmitted signal vectors, respectively. The channel response is contained in the matrix $\boldsymbol{H} \in \mathbb{C}^{N\times M}$, where the element $h_{n,m}$ represents the channel gain between transmit antenna m and receive antenna n. The entries of the channel matrix are assumed independent and identically distributed (i.i.d.) with zero-mean circularly symmetric complex Gaussian distribution and $\mathrm{E}[|h_{n,m}|^2] = 1$. Finally, the complex additive white Gaussian noise (AWGN) vector added at the receiving antennas is represented as $\boldsymbol{w} \in \mathbb{C}^{N\times 1}$.

2.2.1 Channel Capacity

The capacity of a channel determines the theoretical limit of the amount of data bits that can be transmitted over the communication link with no error. The capacity of single antenna narrowband memoryless channels was first derived in [Shannon48] and is represented by the following equation:

$$C = \log_2\left[1 + \mu|h|^2\right] \text{ b/s/Hz},$$

where μ stands for the signal-to-noise ratio (SNR) and h represents the normalized complex gain of the wireless channel. The extension to the multi-antenna case was provided in

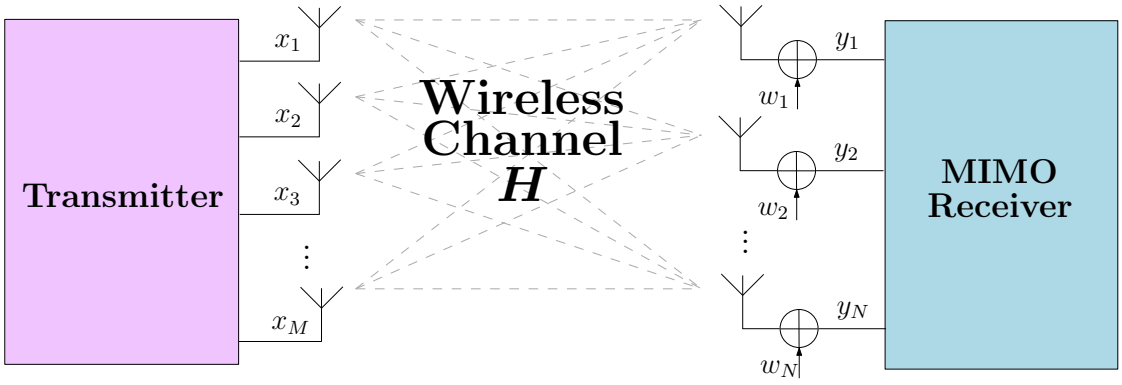

Figure 2.1: Block diagram of an $N \times M$ MIMO system where M stands for the number of transmit antennas and N represents the number of receive antennas.

[Telatar95], where the capacity expression of a flat deterministic MIMO channel for a given transmit power constraint E_{Tr} was derived:

$$C = \max_{\mathrm{Tr}(\boldsymbol{C_x}) \leq E_{\mathrm{Tr}}} \log_2 \left[\left| \boldsymbol{I}_N + \boldsymbol{C_w}^{-1} \boldsymbol{H} \boldsymbol{C_x} \boldsymbol{H}^H \right| \right] \text{ b/s/Hz},$$

with $\mathrm{Tr}(\boldsymbol{A})$, $|\boldsymbol{A}|$ and $\boldsymbol{A}^H$ denoting the trace, determinant and conjugate transpose of matrix $\boldsymbol{A}$, respectively. The covariance matrix of vector $\boldsymbol{z}$ is represented by $\boldsymbol{C_z} = \mathrm{E}[\boldsymbol{z}\boldsymbol{z}^H]$, being $\mathrm{E}[\cdot]$ the expectation operator. Finally, note that $\boldsymbol{I}_N$ stands for the $N \times N$ identity matrix. Based on this expression, the capacity growth with respect to the number of transmit and receive antennas in MIMO systems was proven to scale linearly with $\min(M, N)$ in [Telatar95][Foschini98].

It is particularly noteworthy the influence of the transmit covariance matrix $\boldsymbol{C_x}$ on the achievable capacity. If the channel is unknown to the transmitter, the best transmit strategy is to distribute all the available power equally among the transmission streams [Telatar99]. However, if some feedback from the receiver side of the communication is available, and therefore, the transmitter is aware of the channel conditions, a waterfilling solution can be adopted [Tse05]. By following such an approach the strongest subchannels are assigned the most power, which leads to an overall increase in the performance of the MIMO system.

2.3 Multi-user MIMO

With the advent of new communication technologies, the interest in MIMO has recently evolved towards the development of multi-user schemes which consider more complex albeit realistic scenarios with multiple terminals sharing the time, space, bandwidth and power resources available in a wireless network. Consequently, a great part of the latest research on innovative wireless multi-antenna technologies has been focused on multi-user MIMO (MU-MIMO) environments.

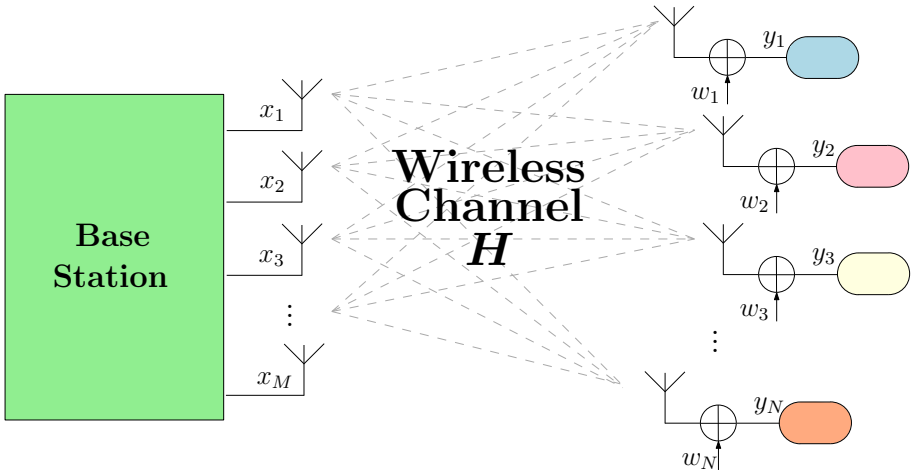

Figure 2.2: Multi-user MIMO broadcast channel with M transmit antennas and N single-antenna user terminals.

A multiple antenna and multi-user system provides a set of advantages over point-to-point MIMO transmissions. One of the main features of MU-MIMO is its greater immunity to propagation shortcomings derived from antenna correlation. Being the antennas hosted at scattered users, the correlation coefficients are inherently low, which allows to overcome the usual problems related to channel rank loss. Another interesting property of MU-MIMO is that direct line of sight propagation, which greatly degrades the quality of the communication link in single-user MIMO systems with spatial multiplexing, does not pose a problem in a multi-user setup. Furthermore, MU-MIMO enables the obtaining of a spatial multiplexing gain at the base station without the requirement of multi-antenna receivers. This allows for the implementation of small, low-cost and low-power terminal devices as the computational load is transferred to the base station.

Nevertheless, the multi-user setup also poses a set of problems that do not exist in the single-user model. For example, the lack of interaction between the users forces the base station to acquire instantaneous knowledge of the channel in order to allow for independent detection of each user's information stream. Additionally, the independence between the receive antennas may also incur in an outage situation if the subchannel directed to a single-antenna user undergoes severe fading. Such a situation in MIMO system can be overcome with simple diversity techniques.

Generally speaking, the multi-user MIMO environment is composed of two channels that communicate the base station with the user terminals: the multiple access channel (MAC), also known as the uplink channel, covers the communication from the terminals to the base station, whereas the broadcast channel (BC), or downlink channel, carries the transmissions that stem from the base station and end at the user terminals, as is shown in Figure 2.2. Hitherto, MU-MIMO techniques for the uplink channel have been widely studied as the detection problem in such systems is equivalent to that of a MIMO channel with multi-user

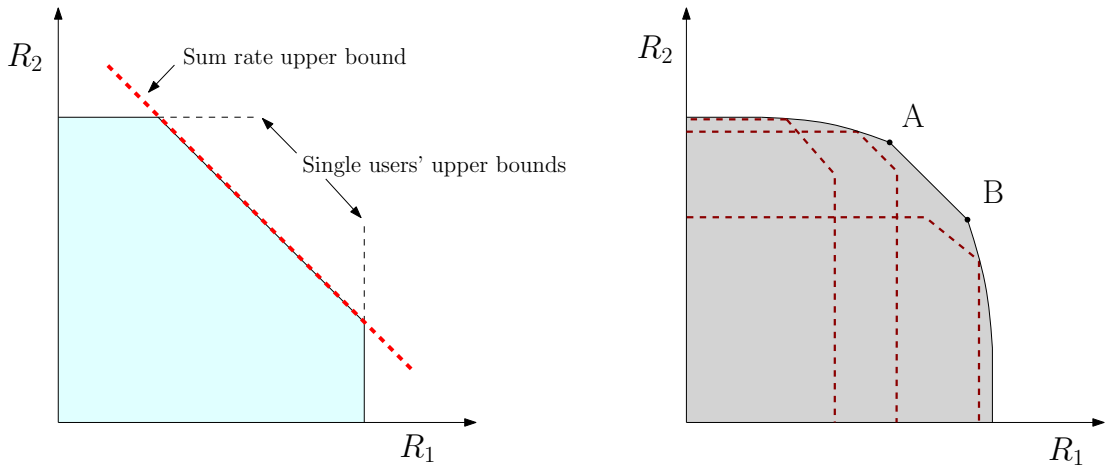

(a) Single-antenna users

(b) Multi-antenna users

Figure 2.3: Capacity region of the two-user MAC.

detection [Verdú98]. However, the study of the downlink channel entails a greater complexity. In the following sections the capacity regions of the uplink and downlink multi-user MIMO channels will be shortly reviewed.

2.3.1 Capacity Region of the Multiple-Access Channel

Unlike in single-user MIMO systems, where the research focuses on the optimization of the total sum rate, a special attention must be paid to the individual rates of the users in multi-user environments. Therefore, the analysis of the theoretical transmission capabilities of the channel is performed via the capacity regions.

The shape of the capacity region of the MAC is determined by the user decoding approach. If a joint-decoding strategy is followed, the capacity region of the MAC with two single-antenna users is shaped as an irregular pentagon [Goldsmith03]. An illustrative example of the pentagon-shaped capacity region is provided in Figure 2.3(a), being R_1 and R_2 the rates of users 1 and 2, respectively. The achievable capacity region is therefore bounded by the single users' upper bounds and the sum rate upper bound. Note that since each user hosts a single antenna, the transmit covariance matrix is a scalar that equals the transmitted power E_{Tr}. If the users are equipped with various antennas, there are several possible approaches to conform the transmit covariance matrix. This way, for a certain power constraint, the capacity region is composed of the union of all the pentagons derived from the set of all possible covariance matrices. Figure 2.3(b) depicts the resulting capacity region of the MAC with two multi-antenna users. For illustrative purposes, the pentagon-shaped capacity regions of three particular transmit covariance matrices are additionally shown in dashed lines. As one can notice, the boundary of the capacity region is curved in most areas except for the sum rate line between points A and B.

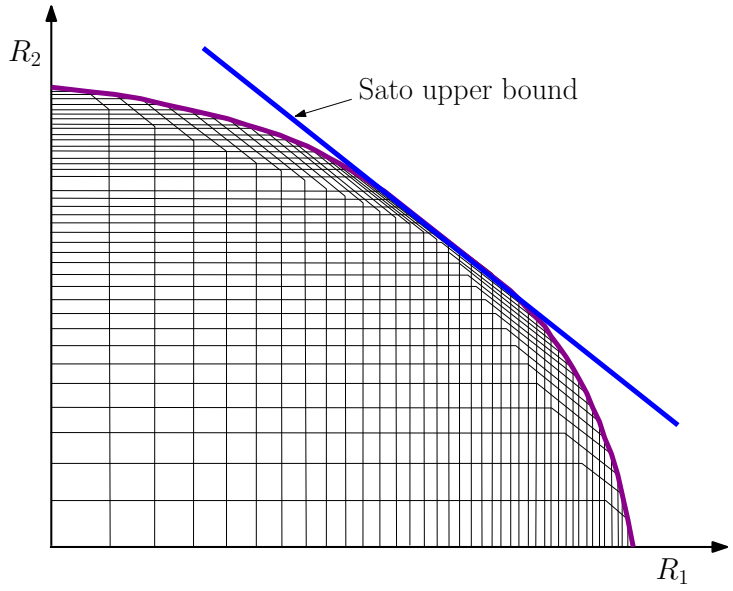

Figure 2.4: Dirty paper broadcast region for multi-user setups with single-antenna users.

2.3.2 Capacity Region of the Broadcast Channel

The achievable rate region for the BC with single-antenna users was presented in [Caire00] and was later extended to the scenario of multi-antenna users in [Yu01]. The achievable rate region in this latest research work was obtained by means of coding by non-causally known interference, a technique also known as dirty paper coding (DPC) [Costa83]. DPC establishes that if the transmitter has non-causal knowledge of the interference in the channel, the resulting capacity region can be equal to that of the channel without the interference if proper coding is employed. Nevertheless, the expressions for the rate derived by the DPC approach are neither convex nor concave functions of the transmit covariance matrices [Goldsmith03], which greatly hinders the computation of the DPC capacity region. A sample DPC region for two single-antenna users is depicted in Figure 2.4. The capacity region in this figure is upper bounded by the Sato bound presented in [Sato78] for the capacity region of general broadcast channels.

The duality between the MAC and BC for constant and scalar channels was first presented in [Jindal01] and later extended to multiple-antenna setups in [Vishwanath02]. The duality theorem states that all achievable rates in the dual MIMO MAC are also achievable in the MIMO BC as long as the sum of the power constraints in the former equals that of the latter. Hence, the capacity region of the BC is equal to the union of the capacity regions of the dual MAC, where the union is taken over all the individual power constraints that sum up to the BC power constraint. This way, the computation of the achievable capacity region in broadcast scenarios can be simplified by solving an equivalent problem in the dual MAC.

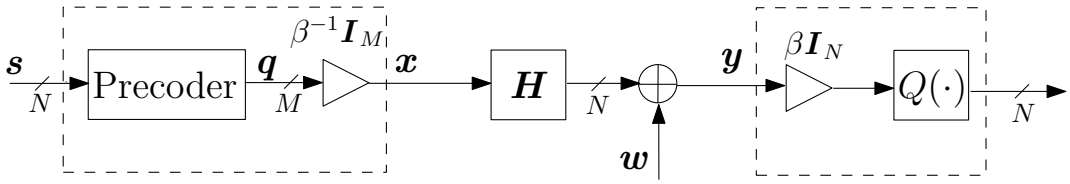

Figure 2.5: Block diagram of a generic precoding system with M transmit antennas and N single-antenna users.

2.4 Precoding Techniques

The lack of cooperation between the terminals at the signal detection stage is the main cause that prevents those well-known detection techniques designed for single-user schemes from being applied into multi-user environments. However, if a preprocessing stage is performed at the base station, it is possible to pre-substract the interference among the transmitted data streams and enable non-cooperative signal detection at the receivers. This technique, which is referred to as precoding, can follow a linear or non-linear approach.

The block diagram of a generic precoding system is shown in Figure 2.5. The user data symbols $\boldsymbol{s} \in \mathbb{C}^{N\times 1}$ are processed by the precoding module and transmitted after a power scaling factor β^{-1} is applied. This scaling stage aims at ensuring that the power of the precoded symbols $E_{\mathrm{SE}} = \mathrm{E}[\boldsymbol{q}^H\boldsymbol{q}]$ does not exceed the preset power quota E_{Tr}. Hence, by applying a scaling factor of $\beta^{-1} = \sqrt{E_{\mathrm{Tr}}/E_{\mathrm{SE}}}$ the average power of the transmitted vector complies with the established transmit power constraint, namely $\mathrm{E}[\boldsymbol{x}^H\boldsymbol{x}] = E_{\mathrm{Tr}}$. The computation of this parameter depends on the power constraint strategy to be used, being usually calculated only once per transmit information block.

At the user terminals the received signal is scaled by β again before the detection process begins. The receivers can gain knowledge of the power control factor by means of training symbols sent from the base station or through a low-rate control channel. The detection stage is usually carried out by means of a nearest-neighbor quantizer $Q(x) = \mathrm{argmin}_{c\in\mathbb{A}}\,|x-c|^2$ which provides hard decisions on the transmitted data symbols:

$$\hat{\boldsymbol{s}} = Q(\boldsymbol{H}\boldsymbol{q} + \beta\boldsymbol{w}). \tag{2.1}$$

From this equation, one can notice that in the event of $E_{\mathrm{SE}} > E_{\mathrm{Tr}}$, or equivalently $\beta > 1$, an increase in the power of the noise vector is experienced at the receivers, which will greatly deteriorate the error-rate performance of the system. In this respect, note that the *detection* SNR experienced at the receivers is subjected to the value of the power scaling factor β, or equivalently to the channel realization and the implemented precoding strategy. Consequently, the SNR will be defined as $E_{\mathrm{Tr}}/\mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})$ for the simulation results presented

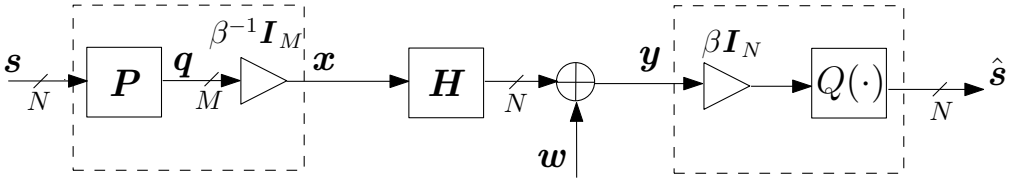

Figure 2.6: Block diagram of a linear precoding system with M transmit antennas and N single-antenna users.

in this dissertation.

Additionally, a unit transmit power constraint per transmit antenna, i.e. $E_{\mathrm{Tr}} = M$, and an uncorrelated noise vector, namely $\boldsymbol{C}_{\boldsymbol{w}} = \sigma^2 \boldsymbol{I}_N$, are considered for all simulations. As an additional remark, the information data streams are extracted from a quadrature amplitude modulation (QAM) constellation of P symbols $\mathbb{A} \triangleq \{a+jb : a, b \in \alpha\{\pm 1, \pm 3, \ldots, \pm\sqrt{P}-1\}\}$ where α is selected so as to ensure unitary average symbol power.

2.4.1 Linear Precoding

Linear precoding techniques aim at pre-equalizing the signal to be transmitted by means of a simple linear transformation, as is shown in Figure 2.6. The precoding filter applied at the base station $\boldsymbol{P} \in \mathbb{C}^{M \times N}$ can be designed following different criteria, such as a zero forcing (ZF) formulation (ZF precoder) or a minimum mean squared error (MMSE) optimization (WF precoder).

2.4.1.1 Zero-Forcing Precoding

If a ZF approach is followed, the removal of all the interference between the transmitted data streams is required. The constraint imposed by the design criteria states that the cascade of precoding filter and channel matrix fulfils $\boldsymbol{H}\boldsymbol{P} = \boldsymbol{I}_N$, which clearly leads to the precoding matrix being equal to the inverse of the channel matrix [Peel05]. More specifically, the right pseudoinverse of the channel is utilized to allow for systems with $M > N$, as shown by the following equation

$$\boldsymbol{P}_{\mathrm{ZF}} = \boldsymbol{H}^H \left(\boldsymbol{H}\boldsymbol{H}^H\right)^{-1}. \tag{2.2}$$

The main drawback of this straightforward approach is the considerable increment in the power of the precoded signal $\boldsymbol{q}$, specially in the event of ill-conditioned channels. As a consequence of this, a substantial noise enhancement and a poor detection SNR are experienced at the receivers, as already shown in (2.1). This problem is aggravated in the fully-loaded, i.e. $M = N$, antenna setup, in which case the precoding matrix in (2.2) is reduced to $\boldsymbol{P} = \boldsymbol{H}^{-1}$. The average power of the linearly precoded signal in such scenario is infinite, as already demonstrated in [Peel05]. Consequently, the characteristic linear growth in the capacity of

multi-antenna systems is unfeasible in fully-loaded ZF linear precoding schemes, being the diversity order equal to that of a single Rayleigh channel in this case. Nevertheless, a certain diversity order increase can be obtained by deploying more antennas at the transmitter side, as the slope of the bit error rate (BER) vs SNR curve in linear ZF precoding systems equals $M - N + 1$.

2.4.1.2 Wiener-Filter Precoding

One of the main disadvantages of the linear ZF precoding technique stems from the requisite to completely null out the interference or *crosstalk* between the different user streams, which potentially leads to a considerable noise boost under certain antenna setups. Hence, by allowing for some interference among the subchannels it is possible to improve the performance of the linear precoder. The optimum *crosstalk* parameter was derived in [Joham05b] following an MMSE problem formulation. The resulting Wiener-filter precoding matrix has a regularized inverse structure with a regularization parameter that is dependant on the noise covariance matrix and the transmit power constraint E_{Tr}:

$$\boldsymbol{P}_{\mathrm{WF}} = \boldsymbol{H}^H \left(\boldsymbol{H}\boldsymbol{H}^H + \frac{\mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})}{E_{\mathrm{Tr}}} \boldsymbol{I}_N \right)^{-1} . \tag{2.3}$$

Having an optimized mean squared error (MSE) the Wiener filter (WF) achieves a better BER performance than other linear precoders [Joham05b]. However, given the SNR-dependency of the *crosstalk* parameter, the diversity order of the linear WF precoder is equal to that of the ZF precoder. This fact is shown in Figure 2.7, where the performance of the linear approaches revised in this section is depicted. Specifically, the setups of a completely loaded system, namely $M = N$, and a system with $M > N$ are analyzed. The performance loss of the linear precoders in the $M = N$ setup is noticeable from the provided BER results.

2.4.2 Non-linear Precoding

A considerable performance improvement over linear precoding techniques can be achieved by including non-linear signal processing algorithms at both ends of the communication link. Nevertheless, the enhanced performance comes at the cost of an increased complexity, which may be prohibitive for real-time practical systems. This is the case of the well-known and capacity achieving DPC technique. Despite its relevance in the theoretic assessment of the capacity of multi-user broadcast channels [cf. Section 2.3.2], Costa's scheme does not provide a practical approach to achieve such benefit. This has led to the development of other non-linear precoding algorithms, such as Tomlinson-Harashima precoding (THP) or VP, that aim at achieving a similar performance with a more reasonable complexity. These

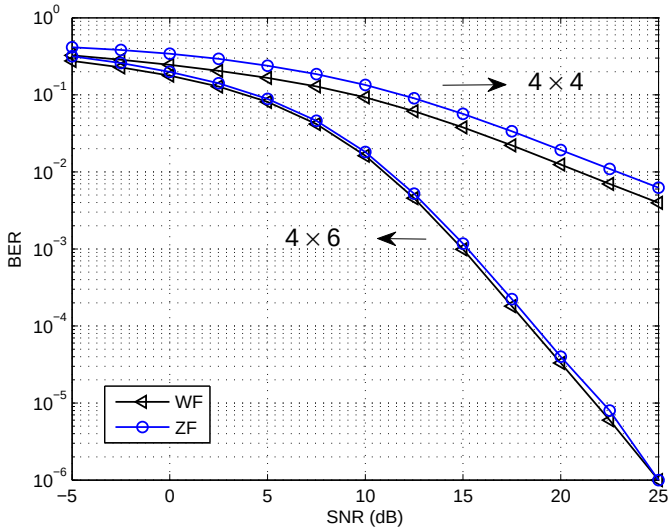

Figure 2.7: BER performance of the ZF and WF linear precoding techniques in the 4×4 and 4×6 antenna setups with 16-QAM modulation.

algorithms are considered to be suboptimal as they do not code the data sequence in time as opposed to DPC [Yu01].

2.4.2.1 Tomlinson-Harashima Precoding

One of the most noteworthy methods of non-linear precoding is THP, which was originally developed in [Tomlinson71] and [Harashima72] to mitigate the effect of intersymbol interference and was later adapted for transmission over MIMO channels in [Fischer02c]. Unlike the aforementioned linear precoding techniques, THP processes the information data sequentially with the aim of successively removing the interference among the transmitted data streams. The interference cancelation is performed in a similar way to the decision-feedback equalization (DFE) method used in the point-to-point MIMO signal detection process. Nevertheless, the symbol constellation is distorted when performing the interference cancelation yielding a higher power of the precoded symbols. To overcome this problem, a modulo operator is inserted in the data processing loop, offering the precoder extra degrees of freedom to select the signal that incurs in the lowest E_{SE}.

The block diagram in Figure 2.8 shows the main features of a THP system, namely, the modulo operators, the interference cancelation loop, which includes the feedback filter $\boldsymbol{F} \in \mathbb{C}^{N \times N}$, and the signal shaping matrix or feedforward filter $\boldsymbol{T} \in \mathbb{C}^{M \times N}$.

The modulo operator can be equivalently described as the addition of integer multiples of the modulo constant τ, in such a way that the input signal is mapped into the fundamental Voronoi region of the lattice $\tau\mathbb{CZ}$, where $\mathbb{CZ} = \mathbb{Z} + j\mathbb{Z}$ represents the set of Gaussian integers. The fundamental Voronoi region of this particular lattice is a square of side τ centered at

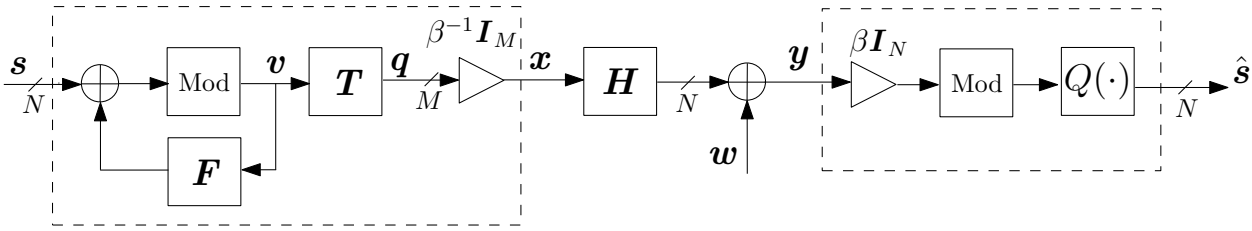

Figure 2.8: Block diagram of a THP system.

the origin, namely $\mathcal{V}_{\mathrm{M}} \triangleq \{x + jy | x, y \in [-\tau/2, \tau/2)\}$. The symbols in this equivalent region require a smaller transmit power, and hence their transmission generates a better detection SNR. The modulo operator works independently on the real and imaginary components of the data signal as shown by the following equation:

$$\mathrm{Mod}(d) = d - \left\lfloor \frac{\Re(d)}{\tau} + \frac{1}{2} \right\rfloor \tau - j \left\lfloor \frac{\Im(d)}{\tau} + \frac{1}{2} \right\rfloor \tau,$$

being $\Re(d)$ and $\Im(d)$ the real and imaginary components of the complex-valued signal d, with $\lfloor \cdot \rfloor$ denoting the flooring operation. Note that $\mathrm{Mod}(d+f) = \mathrm{Mod}(d)$ holds if $f \in \tau\mathbb{CZ}$ and $\mathrm{Mod}(d) = d$ is satisfied if $d \in \mathcal{V}_{\mathrm{M}}$. The value of the τ parameter depends on the type of modulation and is set to $\tau = 2c_{\max} + \Delta$, where $c_{\max}$ is the absolute value of the constellation symbol with the largest magnitude in the real or imaginary axis, and Δ is the minimum spacing between constellation points. The provided expression for the modulo constant yields values of $2\sqrt{2}$, $8/\sqrt{10}$, and $16/\sqrt{42}$ for quadrature phase shift keying (QPSK), 16-QAM and 64-QAM modulation alphabets, respectively. This choice of τ ensures that $\mathbb{A} \in \mathcal{V}_{\mathrm{M}}$, as is illustrated in Figure 2.10(a) for the particular case of QPSK modulation. The mapping procedure needs to be reversed for signal detection, and can be simply performed by adding a modulo operator at the receivers.

The successive cancelation process is enabled by the feedback filter $\boldsymbol{F}$ whose triangular structure with zero main diagonal ensures that only those data symbols that have already been processed are fed back to the interference-cancelation loop. This property of the feedback filter is often referred to as *spatial causality* [Joham05a]. Considering a lower triangular matrix $\boldsymbol{F}$, the user data streams are processed sequentially as $i = 1, \ldots, N$. This way, the signal prior to applying the shaping matrix $\boldsymbol{T}$ is computed as

$$v_i = \mathrm{Mod}\left(s_i + \boldsymbol{e}_i^T \sum_{j=1}^{i-1} \boldsymbol{f}_j v_j \right),$$

where $\boldsymbol{e}_i$ represents the i^{th} column of the $N \times N$ identity matrix and $\boldsymbol{f}_j$ stands for the j^{th} column vector of the feedback matrix $\boldsymbol{F}$. Hence, the first data symbol is transmitted

unaltered, namely $v_1 = s_1$ since $\mathbb{A} \in \mathcal{V}_{\mathrm{M}}$, while the second data symbol is modified according to the interference produced by the first symbol, and so on. Therefore, the data stream of the last coded user is transmitted taking into account the interference of all the other users, and therefore experiences no interference at all, while the rest of the users are subjected to the interference of the users precoded in the subsequent stages.

The performance of the successive interference cancelation procedure is strongly related to the order in which the user streams are precoded. The user permutation scheme is defined by the tuple $\mathfrak{O} = \{\mathfrak{o}_1, \ldots, \mathfrak{o}_N\}$ with $\mathfrak{o}_i \in \{1, \ldots, N\}$, which establishes that the $\mathfrak{o}_i^{\mathrm{th}}$ user is precoded at the i^{th} step. This process is compactly represented by the permutation matrix $\boldsymbol{\Pi}$ which is computed as

$$\boldsymbol{\Pi} = \sum_{i=1}^{N} \boldsymbol{e}_i \boldsymbol{e}_{\mathfrak{o}_i}^T,$$

where $(\cdot)^T$ represents the matrix/vector transpose operator. Note that $\boldsymbol{\Pi}$ is a unitary matrix by construction, i.e. $\boldsymbol{\Pi}^T = \boldsymbol{\Pi}^{-1}$.

Once the interference cancelation procedure has been carried out, the data is shaped by the feedforward filter $\boldsymbol{P}$ prior to transmission. The design of both feedback and feedforward filters is performed jointly and can be subject to a variety of constraints. The two most noteworthy variants of THP filters are built by following a ZF or MMSE criterion, yielding the so-called THP-ZF [Fischer02a][Joham04b][Joham05a] and THP-WF [Joham04b][Kusume05] [Joham05a] solutions, respectively. The expressions for the feedback and feedforward filters of the aforementioned variants of THP are summarized here for completeness. Note that $\boldsymbol{A}^{\dagger}$ stands for the Moore-Penrose pseudoinverse of matrix $\boldsymbol{A}$ and that the projection matrix $\boldsymbol{\Upsilon}_i = \boldsymbol{I}_N - \sum_{j=i+1}^{N} \boldsymbol{e}_{\mathfrak{o}_j} \boldsymbol{e}_{\mathfrak{o}_j}^T$ is used to null out the rows of the channel matrix belonging to receivers which will be precoded at a later step.

$$\boldsymbol{P}_{\mathrm{ZF}} = \sum_{i=1}^{N} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i^T \left(\boldsymbol{\Upsilon}_i \boldsymbol{H} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i\right)^{\dagger} \boldsymbol{e}_{\mathfrak{o}_i} \boldsymbol{e}_i^T$$

$$\boldsymbol{F}_{\mathrm{ZF}} = \boldsymbol{I}_N - \boldsymbol{\Pi} \boldsymbol{H} \boldsymbol{P}_{\mathrm{ZF}}$$

$$\boldsymbol{P}_{\mathrm{WF}} = \sum_{i=1}^{N} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i \left(\boldsymbol{\Upsilon}_i \boldsymbol{H} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i + \frac{\mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})}{E_{\mathrm{Tr}}} \boldsymbol{I}_N\right)^{-1} \boldsymbol{e}_{\mathfrak{o}_i} \boldsymbol{e}_i^T$$

$$\boldsymbol{F}_{\mathrm{WF}} = \sum_{i=1}^{N} \left(\boldsymbol{\Pi} \boldsymbol{\Upsilon}_i \boldsymbol{\Pi}^T - \boldsymbol{I}_N\right) \boldsymbol{\Pi} \boldsymbol{H} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i \left(\boldsymbol{\Upsilon}_i \boldsymbol{H} \boldsymbol{H}^H \boldsymbol{\Upsilon}_i + \frac{\mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})}{E_{\mathrm{Tr}}} \boldsymbol{I}_N\right)^{-1} \boldsymbol{e}_{\mathfrak{o}_i} \boldsymbol{e}_i^T$$

THP precoding schemes offer a significant performance improvement with respect to simple linear precoding techniques due to the interference cancelation procedure performed on the user information streams. Nevertheless, since the first data stream undergoes linear pre-equalization only, the average BER curve will show the same diversity order as if no feedback loop were present, that is $M - N + 1$ [Windpassinger04].

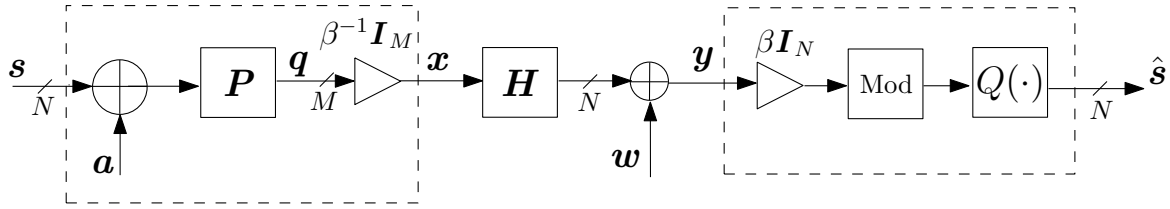

Figure 2.9: Block diagram of a VP system.

2.4.2.2 Vector Precoding

As has been shown in the previous sections, the performance gap between the optimum albeit intractable DPC solution and the linear precoding approaches can be reduced by means of non-linear precoding techniques such as THP. Nevertheless, it is still possible to draw closer to the optimum performance with a reasonable complexity. VP, which resembles the method known as shaping without scrambling for dispersive channels [Fischer02b][Fischer03], is considered to be one of the most promising non-linear precoding approaches due to its close-to-optimum performance.

If the modulo operator at the transmitter side of a THP system is equivalently represented as the addition of an auxiliary vector $\boldsymbol{a}'$ and moved outside the feedback loop, it is possible to combine the feedback and feed-forward filters into a single matrix $\boldsymbol{P}' = \boldsymbol{T}(\boldsymbol{I}_N - \boldsymbol{F})^{-1}$ and compute the signal prior to the scaling stage as $\boldsymbol{q} = \boldsymbol{P}'(\boldsymbol{s} + \boldsymbol{a}')$. This derives in the system model shown in Figure 2.9, which represents the block diagram of a VP system.

The concept of VP was originally developed in [Hochwald05] and has been a topic of extended research ever since, e.g. [Schmidt05][Callard06][Boccardi06]. By perturbing the signal prior to transmission as shown in Figure 2.9, it is possible to reduce the average power of the signal before the scaling stage, namely $E_{\mathrm{SE}} = \mathrm{E}[||\boldsymbol{q}||^2]$, and therefore improve the overall system performance. The modulo operator at the receivers provides the transmitter with additional degrees of freedom to choose the perturbation vector that is most suitable. This process is illustrated in Figure 2.10, where the infinite repetition of the modulation constellation before and after precoding in shown. Note that the proper selection of the modulo constant allows for the extensions of the modulo alphabet in the complex plane to be centered around $\tau\mathbb{CZ}$ [Dietrich08]. The congruent constellations in Figure 2.10(a) are distorted after the precoding procedure, as is shown in Figure 2.10(b). For some cases, such as the one depicted in this figure, there exists a symbol in one of the congruent constellations whose power is smaller than that of the precoded symbol $\boldsymbol{Ps}$. Hence, the transmission of the equivalent $\boldsymbol{P}(\boldsymbol{s}+\boldsymbol{a})$ symbol incurs in a smaller value of the power scaling factor β, which ultimately results in a better SNR at the receivers. Note that the perturbation signal must be composed of integer multiples of the modulo constant τ, namely $\boldsymbol{a} \in \tau\mathbb{CZ}^N$, so that it can be removed at the receivers by means of a simple modulo operation. This feature enables the receivers to eliminate the perturbation signal without the knowledge of the particular

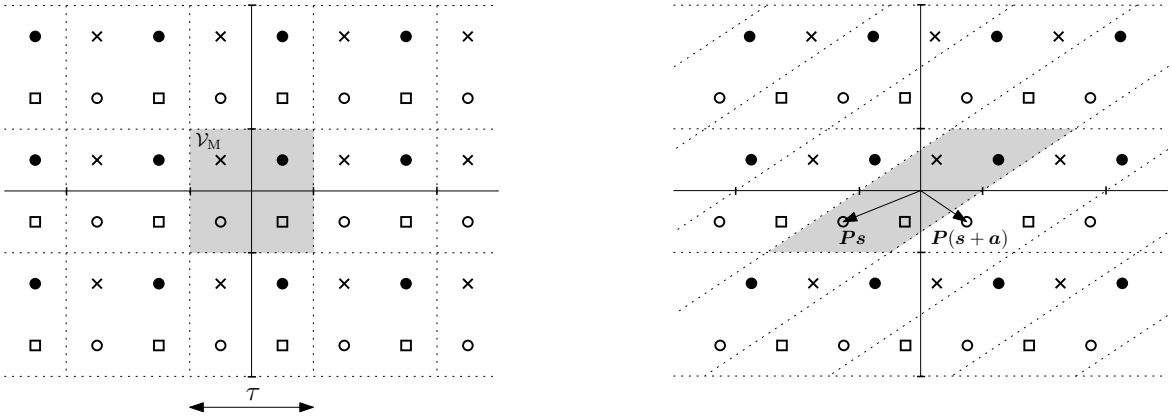

(a) (b)

Figure 2.10: Illustration of the infinite repetition of the modulation constellation before precoding (a) and after (b). Figure (b) shows that it is possible to transmit the same symbol with a smaller power if a perturbation vector is added to the user data stream.

perturbation vector of choice.

Unlike THP, where the interference cancelation is performed successively, VP optimizes the perturbation signal $\boldsymbol{a}$ directly. This way, the successive precoding procedure in THP only takes into account previously precoded symbols but fails to include those precoded at later stages, whereas VP performs the optimization considering all transmit symbols. Consequently, THP is considered a constrained type of VP, and therefore, it will always be outperformed by the latter as long as the same criteria is used for the filter design.

The design of the precoding filter $\boldsymbol{P}$ can be based on different criteria. If the goal is to completely cancel out the interference between the users, a ZF approach should be followed [Hochwald05][Schmidt05]. The requisite of null distortion established by the ZF formulation requires that $\boldsymbol{H}\boldsymbol{P} = \boldsymbol{I}_N$, which clearly results in the precoding filter given by Equation (2.2). The perturbation signal should then be selected so as to minimize the transmit power prior to applying the power control factor β^{-1}:

$$\boldsymbol{a}_{\mathrm{ZF}} = \underset{\hat{\boldsymbol{a}} \in \tau\mathbb{C}\mathbb{Z}^N}{\operatorname{argmin}} \left\| \boldsymbol{P}_{\mathrm{ZF}}(\boldsymbol{s} + \hat{\boldsymbol{a}}) \right\|_2^2. \tag{2.4}$$

As is the case with the linear precoding approaches, the performance of the VP system can be further enhanced by regularizing the inverse involved in the definition of the precoding filter in (2.2) so that a smaller value of E_{SE} can be achieved [Hochwald05]. This new system formulation represents the so-called regularized VP solution (VP-Reg) shown in Equations (2.5) and (2.6).

$$\boldsymbol{P}_{\mathrm{Reg}} = \boldsymbol{H}^H \left(\boldsymbol{H}\boldsymbol{H}^H + \xi_{\mathrm{Reg}} \boldsymbol{I}_N \right)^{-1} \tag{2.5}$$

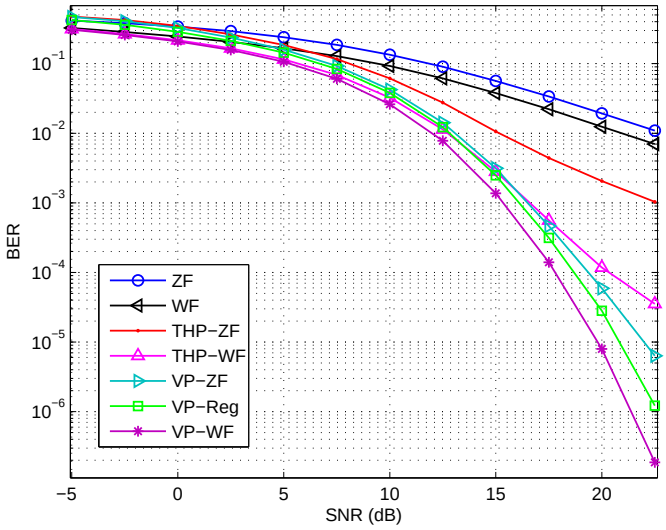

Figure 2.11: BER performance of non-linear and linear precoding schemes.

$$\boldsymbol{a}_{\mathrm{Reg}} = \underset{\hat{\boldsymbol{a}} \in \tau \mathbb{C}\mathbb{Z}^N}{\operatorname{argmin}} \left\| \boldsymbol{P}_{\mathrm{Reg}} (\boldsymbol{s} + \hat{\boldsymbol{a}}) \right\|_2^2 \tag{2.6}$$

No closed-form expression for the regularization factor or crosstalk parameter ξ_{Reg} was provided in the aforementioned research work. Nevertheless, experimental values of ξ_{Reg} were given for some antenna setups, e.g., $\xi_{\mathrm{Reg}} \approx \mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})/(5NE_{\mathrm{Tr}})$ for $N = 4$. Similarly to the linear precoding case, the regularized solution of the precoding filters for VP always outperforms the ZF solution in terms of BER.

In [Schmidt05] and [Schmidt08] a novel approach for VP filter design was proposed. Instead of aiming for a minimum power of the precoded symbols, the aforementioned approach achieves the optimum compromise between noise enhancement and residual interference, yielding the MMSE solution for VP (VP-WF). Given that this approach minimizes the MSE, it is clear that VP-WF outperforms all the other variants of VP presented hitherto. The precoding matrix resulting from the aforementioned optimization problem is given by (2.3). As opposed to the other VP techniques, the matrix involved in the cost function evaluated to select the optimum perturbation vector is not equal to the precoding matrix [see (2.4) and (2.6)]. In this case, any matrix $\boldsymbol{U}$ that fulfils

$$\boldsymbol{U}^H \boldsymbol{U} = \left(\boldsymbol{H}\boldsymbol{H}^H + \frac{\mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})}{E_{\mathrm{Tr}}} \boldsymbol{I}_N \right)^{-1} \tag{2.7}$$

is used to compute the perturbation vector, which gives

$$\boldsymbol{a}_{\mathrm{WF}} = \underset{\hat{\boldsymbol{a}} \in \tau \mathbb{C}\mathbb{Z}^N}{\operatorname{argmin}} \left\| \boldsymbol{U} (\boldsymbol{s} + \hat{\boldsymbol{a}}) \right\|_2^2. \tag{2.8}$$

The matrix decomposition in (2.7) can be performed by means of the Cholesky decomposition of $\left(\boldsymbol{H}\boldsymbol{H}^H + \xi \boldsymbol{I}_N\right)^{-1}$, with $\xi = \mathrm{Tr}(\boldsymbol{C}_{\boldsymbol{w}})/E_{\mathrm{Tr}}$ or by the QR factorization of the augmented matrix $\tilde{\boldsymbol{H}} = [\boldsymbol{H}\ \sqrt{\xi}\boldsymbol{I}_N]^H$, with $\boldsymbol{U} = \boldsymbol{R}^{-H}$, for example. These two matrix decomposition methods yield triangular $\boldsymbol{U}$ matrices, whose special structure can be exploited to ease the closest-point lattice search in (2.8).

It is worth pointing out that, unlike the rest of the precoding techniques studied so far, with the exception of the computationally-intractable DPC, VP is the only precoding approach that can achieve full-diversity order. This fact was proven for VP systems using lattice-reduction techniques in [Taherzadeh05, Taherzadeh07], and for different transmit power constraints in [Jaldén08].

The full diversity order of VP techniques can be appreciated in the BER vs SNR curve provided in Figure 2.11, where the error rate performance of the addressed non-linear approaches has been assessed for a 4×4 antenna setup with 16-QAM modulation. Specifically, the THP-ZF and THP-WF algorithms as described in Section 2.4.2.1 are considered, along with the recently reviewed VP-ZF, VP-Reg and VP-WF techniques. The performance results of the linear algorithms reviewed in Section 2.4.1 have been also included for completeness. The simulation results depicted in Figure 2.11 show that the incorporation of non-linear signal processing techniques results in a substantial error-rate performance improvement when compared to the linear precoding schemes. However, note that both THP and the addressed linear precoders yield the same diversity order, as already discussed in Section 2.4.2.1. Moreover, the WF variants of all the precoders yield better BER performance results than their ZF counterparts. Due to the use of modulo arithmetics and an MMSE-optimized model, VP-WF is the precoding scheme that achieves the best error-rate performance among the addressed precoding approaches. Consequently, VP-WF will be the precoding method of choice in the remainder of this thesis (we shall refer to it as simply VP), with the exception of Chapter 3 where the VP-ZF variant will be used instead for the sake of analytical simplicity.

Note that the computation of the perturbation signal in all the VP variants, namely (2.4), (2.6) and (2.8), entails the search for the closest-point in an infinite lattice, which is known to be in the class of non-deterministic polynomial-time hard (NP-hard) problems [Grötschel93]. Hence, the addition of the perturbation vector poses two main challenges: on one hand, the efficient and low-complexity computation of the perturbing signal, and on the other hand, the analytical assessment of a precoding system influenced by a signal that represents the result of a computationally complex process and whose statistical properties are unknown. The work presented in this thesis will try to address this issues by proposing low-complexity algorithms for the computation of the perturbation signal and by presenting useful bounds to overcome the initially complex performance assessment of VP techniques.

2.5 Chapter Summary

This chapter has reviewed the main features of MIMO and multi-user MIMO techniques. More specifically, a brief overview of the capacity limit of single-user MIMO systems has been provided, along with a similar analysis for the uplink and downlink channels of multi-user environments.

On a more practical scope, the key features of linear and non-linear precoding algorithms have been described in the last section. Starting from the more straightforward linear precoding techniques, it has been shown that the performance of ZF filters can be enhanced by loosening the constraint of null interference among user streams. Next, several non-linear precoding approaches that make use of modulo arithmetics, such as THP and VP, have been reviewed. Due to the power-limiting capacity of the modulo function, non-linear precoders have shown a considerable performance improvement with respect to their linear counterparts. However, since at least one of the users' data streams is only linearly transformed in the THP model, the diversity order yielded by this non-linear technique remains the same as that of the linear precoders. VP offers an additional error-rate performance enhancement and diversity order increase over THP by replacing the successive interference cancelation by a direct optimization procedure. Provided simulation results show the superior performance of VP techniques and their potential to reach the full diversity order of multi-antenna systems.

Chapter 3

Asymptotic Sum Rate Analysis and Rate Optimizations for VP

3.1 Introduction

Despite its enhanced performance, the analysis of VP systems from an information-theoretic point of view is hindered by the effect of the perturbation vector on the unnormalized transmit power E_{SE}. This parameter, which represents the power of the precoded symbols prior to the power scaling stage, is pivotal when determining the performance of any precoding system as it determines the effective noise power present at the detection stage. Closed-form expressions for E_{SE} have not been derived for VP so far, which has led to an increased interest in developing useful bounds.

The first lower bound on this parameter was proposed in [Hochwald05] for an arbitrary number of user terminals and constellation sizes. The presented bound on E_{SE} was given as a function of the perturbed data and the eigenvectors of the precoding matrix, and hence, required numerical evaluation. The approaches in [Müller08, Zaidel08] introduce lower bounds on the unnormalized transmit power in the limit of the system, e.g. $N, M \rightarrow \infty$, by analyzing the process of data perturbation from a statistical physics perspective. Nevertheless, the bounds provided in the aforementioned publications require a number of assumptions and are given as a function of fixed-point integrals which need to be numerically evaluated. These facts limit the insight given by these lower bounds into the E_{SE} value.

Recently, a lower bound on E_{SE} has been derived using a lattice-theoretic approach which has enabled the establishment of an upper bound on the achievable sum rate capacity of VP systems with ZF precoding [Razi09][Ryan09]. This work has been extended to precoding systems with WF filtering in [Razi11] by estimating the entropy of the data symbols at the output of the modulo receiver. However, an upper bound on E_{SE} is still missing to be able to define the performance range of VP systems.

In this chapter, several upper and lower bounds on the unnormalized transmit power will be discussed and introduced. By means of these bounds, it will be possible to completely

delimit the performance of VP-ZF. Additionally, based on the derived bounds, three different optimizations with respect to the rates will be investigated: weighted sum rate maximization, quality of service (QoS) optimization and rate balancing. Finally, ergodic expressions and simulation results on the performance gap between VP and linear precoding will be provided for the aforementioned applications.

Given the lattice-theoretic nature of the derived bounds, an introductory section that summarizes the basic features of lattice theory will be presented first with the aim of setting a theoretic foundation for the subsequent analysis.

3.2 Concepts of Lattice Theory

Every lattice of dimension n is generated by the linear combination of a set of linearly independent vectors $\{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n\}$ that represent the basis of the lattice. This basis is not unique, so different sets of basis vectors can represent the same lattice. Given a set of basis vectors, the generator matrix of the lattice can be represented as $\boldsymbol{M} = [\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n]$. The rest of the vectors in the lattice are generated from the matrix $\boldsymbol{M}$ as linear combinations of the form $\boldsymbol{pM}$, where $\boldsymbol{p} \in \mathbb{Z}^n$. The lattice generated from a certain generator matrix $\boldsymbol{M}$ will be denoted as $\Lambda_{\boldsymbol{M}}$.

The volume is one of the most representative features of a lattice as it remains unchanged even if an equivalent set of basis vectors is utilized. The volume of a lattice is directly related to its generator matrix and is computed as [Conway98][Ryan09]:

$$\mathrm{Vol}\left(\Lambda_{\boldsymbol{M}}\right) = \begin{cases} |\boldsymbol{M}| & \text{for } \boldsymbol{M} \in \mathbb{R}^n \\ |\boldsymbol{M}^H \boldsymbol{M}| & \text{for } \boldsymbol{M} \in \mathbb{C}^n. \end{cases} \tag{3.1}$$

The volume of the lattice is actually equal to the volume of any of its congruent Voronoi regions. These regions are composed of those points in the Euclidean space that are at least as close to Λ_i as they are to any other Λ_j, where Λ_i and Λ_j represent any two points of the lattice. The Voronoi cells are therefore convex polytopes whose union is the whole Euclidean space. These regions were originally described in [Dirichlet50][Voronoi07] and have been extensively studied thereafter [Conway82][Conway84]. The shape of the Voronoi regions depends on the lattice structure. This way, the Voronoi cells of the integer lattice $\mathbb{Z}^n$ are n-dimensional cubes, while in the case of the 2-dimensional hexagonal lattice (also known as the A_2 lattice) they are shaped as regular hexagons. A certain Voronoi cell can therefore be uniquely described in terms of the generator matrix of the lattice and the lattice point at the center of the cell, namely $\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)$. The Voronoi cell with $\Lambda_i = 0$ is also known as the fundamental Voronoi region and is often used as a representative of all the other congruent cells. In Figure 3.1 some congruent Voronoi regions of an arbitrary 2-dimensional lattice are depicted.

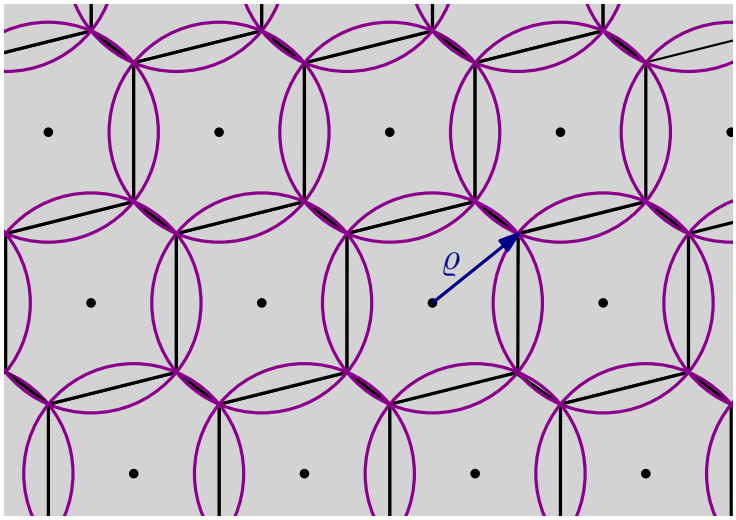

Figure 3.1: Congruent Voronoi regions and covering radius ϱ of an arbitrary lattice.

The longest distance from the vertices of the Voronoi region with respect to Λ_i is known as the covering radius ϱ, which also represents the radius of the smallest sphere that circumscribes the Voronoi region, as is shown in Figure 3.1. Hence, n-dimensional overlapping spheres of radius ϱ centered at each point of the lattice will cover the whole Euclidean space, and no smaller radius will do. The problem of finding the covering radius is in the class of NP-hard problems [vanEmde Boas81].

This leads us to the well-known covering problem, which deals with finding the most economical way of covering the whole n-dimensional Euclidean space with equal and overlapping spheres [Conway98]. Clearly, the best covering strategy will be the one that allows for the smallest amount of overlapping between the spheres. Consequently, the assessment of the quality of a covering strategy is usually performed by means of the *thickness* or covering density of the lattice. This measure represents the average number of spheres that contain a point of the space, and can be computed as

$$\Theta = \frac{V_n \varrho^n}{|\boldsymbol{M}|} \geq 1, \tag{3.2}$$

where $V_n = \pi^{n/2}/(n/2)!$ represents the volume of an n-dimensional sphere of unitary radius. The quality of a certain covering strategy can also be assessed by means of the covering efficiency $\chi = \varrho/\varrho_{\text{eff}}$, where the effective radius of the covering ϱ_{eff} represents the radius of a virtual n-dimensional sphere with the same volume as the Voronoi region of the lattice, namely

$$\text{Vol}\left(\Lambda_{\boldsymbol{M}}\right) = V_n \ \varrho_{\text{eff}}^n. \tag{3.3}$$

The second moment of inertia of the Voronoi regions is also a distinctive feature of a

lattice. It represents the average square distance of the points in $\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)$ with respect to Λ_i. Two are the most remarkable variants of this measure: the dimensionless second moment $G[\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)]$ and the normalized second moment $I[\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)]$, which are given by the following equations:

$$G\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right] = \frac{1}{n} \operatorname{Vol}[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)]^{-\frac{n+2}{n}} \int\limits_{\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)} \|x - \Lambda_i\|^2 dx \tag{3.4}$$

$$I\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right] = \operatorname{Vol}\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right]^{-1} \int\limits_{\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)} \|x - \Lambda_i\|^2 dx. \tag{3.5}$$

As one can notice from Equation (3.4), the normalization with respect to the number of dimensions derives in $G[\mathcal{V}(\Lambda_{\boldsymbol{M}}, \Lambda_i)]$ being a dimensionless parameter, as stated by its own name. The two forms of the second moment of inertia represent, in essence, the same problem. Therefore, Equations (3.4) and (3.5) can be easily combined, yielding the following expression:

$$I\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right] = n \operatorname{Vol}\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right]^{\frac{2}{n}} G\left[\mathcal{V}\left(\Lambda_{\boldsymbol{M}}, \Lambda_i\right)\right]. \tag{3.6}$$

3.3 Bounds on the Unnormalized Transmit Power

The performance of any precoding system is closely related to the power of the signal prior to the scaling stage. Nevertheless, the computation of E_{SE} involves the expectation over the result of a closest-point search problem, which hinders the assessment of the expected performance of VP systems. However, it is possible to circumvent this problem by setting upper and lower bounds on the unnormalized transmit power, which will eventually enable the determination of the performance range of VP systems. To this end, the analysis of several lower and upper bounds on E_{SE} is performed in this section.

The VP system model described in Section 2.4.2.2 will be slightly modified for the current analysis to allow for the definition of optimization problems and to simplify the analytical study of E_{SE}. Hence, a VP-ZF precoding model will be used instead, since the SNR-dependency of the precoding matrix in the VP-WF variant would greatly complicate the analysis of the lattice features. Following the arguments in [Razi09], the user data symbols will be assumed to be uniformly distributed over $\mathcal{Q} \triangleq \{d : |\Re(d)| < 1/2, |\Im(d)| < 1/2\}$. This latter assumption eases the tractability of the performance analysis employing modulo arithmetics. The error introduced by assuming uniformly distributed symbols can be neglected specially if higher order modulations are used, as is shown in [Ryan08]. Following from this, the modulo constant will be set to $\tau = 1$. At the receiver side, the re-scaling stage performed prior to applying the modulo operator will be redefined to enable the definition of optimization problems on the individual user data rates. This way, the k^{th} user will weight

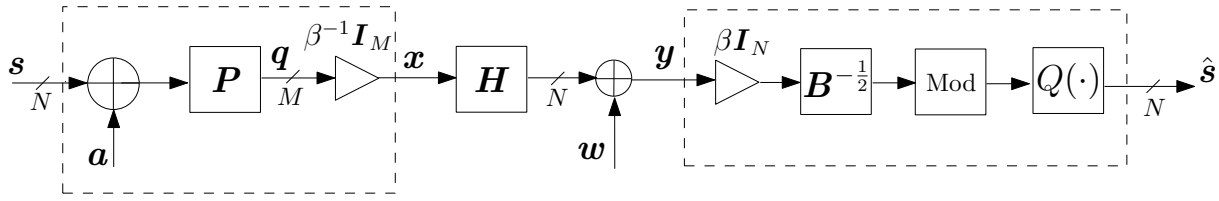

Figure 3.2: Block diagram of the VP system used for the study of the unnormalized transmit power.

the received signal y_k with $\beta\ b_k^{1/2}$, where the non-negative weights $\{b_1, \ldots, b_N\}$ represent the entries of the diagonal matrix $\boldsymbol{B} \in \mathbb{R}^{N \times N}$. Also, the elements of the AWGN vector are assumed to be $w_k \sim \mathcal{CN}(0,1)$. The block diagram of the VP model used for the current analysis is shown in Figure 3.2.

The signal detected at the user terminals after the incorporation of the power loading matrix into the VP system model is given by the following expression:

$$\hat{\boldsymbol{s}} = Q\left\{\text{Mod}\left[\boldsymbol{B}^{-\frac{1}{2}}\boldsymbol{H}\boldsymbol{P}\left(\boldsymbol{s}+\boldsymbol{a}\right)+\beta\boldsymbol{B}^{-\frac{1}{2}}\boldsymbol{w}\right]\right\}. \tag{3.7}$$

When a ZF formulation is followed, no interference is allowed among user streams, and therefore, the cascade of precoding matrix, channel matrix and the inverse square root of the power loading matrix should equal the identity matrix, i.e. $\boldsymbol{B}^{-\frac{1}{2}}\boldsymbol{H}\boldsymbol{P} = \boldsymbol{I}_N$. This leads to the precoding matrix being designed as $\boldsymbol{P} = \boldsymbol{H}^{\dagger}\boldsymbol{B}^{\frac{1}{2}}$ for the VP system under study. Hence, the power of the precoded symbols is given by:

$$E_{\text{SE}} = \text{E}\left[\min_{\boldsymbol{a}\in\mathbb{C}\mathbb{Z}^N}\left\|\boldsymbol{H}^{\dagger}\boldsymbol{B}^{\frac{1}{2}}\left(\boldsymbol{s}+\boldsymbol{a}\right)\right\|_2^2\right]. \tag{3.8}$$

As it is known from lattice theory [Conway85], the encoded signal $\boldsymbol{q} = \boldsymbol{H}^{\dagger}\boldsymbol{B}^{\frac{1}{2}}(\boldsymbol{s}+\boldsymbol{a})$ lies within the fundamental Voronoi region of the lattice rendered by the generator matrix $\boldsymbol{P} = \boldsymbol{H}^{\dagger}\boldsymbol{B}^{\frac{1}{2}}$, i.e. $\boldsymbol{q} \in \mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)$, if the lattice search to find the perturbation vector is performed optimally. Therefore, the magnitude of E_{SE} equals the normalized second moment of the Voronoi region $\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)$ [Conway85], which enables its representation by means of the lattice parameters of $\Lambda_{\boldsymbol{P}}$.

Given that the real and imaginary components of the complex-valued signals in Figure 3.2 are considered as separate dimensions, we get $n = 2N$. Also, following from Equation (3.1), the volume of the Voronoi region with complex generator matrix $\boldsymbol{P}$ equals $\text{Vol}[\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)] = |\boldsymbol{W}^{-1}\boldsymbol{B}|$, where $\boldsymbol{W} = \boldsymbol{H}\boldsymbol{H}^H$ is an $N \times N$ Wishart distributed matrix with M degrees of freedom, assuming that the entries of $\boldsymbol{H}$ are i.i.d. zero-mean unit-variance complex Gaussian distributed. Thus, the average power of the precoded symbols can be described by means of

the lattice parameters of $\Lambda_{\boldsymbol{P}}$ as shown in (3.6), which gives

$$E_{\mathrm{SE}} = 2N\ G\left[\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)\right]\ |\boldsymbol{W}^{-1}\boldsymbol{B}|^{\frac{1}{N}}. \tag{3.9}$$

As one can notice from Equation (3.9), it is possible to delimit the operation range of E_{SE} by setting appropriate lower and upper bounds on $G[\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)]$. The study and discussion of these bounds will be performed in the following subsections.

3.3.1 Lower Bounds

In the search for the optimum quantization strategy, Zador derived the upper and lower bounds on the second moments of inertia of the ideal quantizer (G_{opt}) [Zador63], which are given by

$$\frac{1}{(n+2)\pi}\Gamma\left(\frac{n}{2}+1\right)^{2/n} \leq G_{\mathrm{opt}} \leq \frac{1}{n\pi}\Gamma\left(\frac{n}{2}+1\right)^{2/n}\Gamma\left(\frac{2}{n}+1\right), \tag{3.10}$$

where $\Gamma(x)$ stands for the Gamma function, with $\Gamma(x) = (x-1)!$ if $x \in \mathbb{Z}$. Interestingly, both bounds converge to $G_{\mathrm{opt}} \to 1/(2\pi e)$ in the limit, i.e. when $n \to \infty$.

In [Zador63] it was argued that no polytope or convex body of other kind attained a smaller dimensionless second moment than the n-dimensional sphere $\mathcal{S}_n$. The lower bound given for the ideal quantizer in (3.10) therefore stands for $G(\mathcal{S}_n)$ and is known as the *sphere bound*. This bound was used to set a lower limit to the unnormalized transmit power E_{SE} of a VP system in [Ryan08] yielding very promising results.

Despite the relevance of the *sphere bound* in the theoretical field, the design of a quantizer with these features is not possible in practical systems. This is due to the fact that this lower bound is only reached if the Voronoi regions are shaped as spheres, but no sphere arrangement is able to reach a full coverage of the space. Motivated by this result, a significantly stronger lower bound was derived in [Conway85][Conway98] based on a geometrical argumentation on the features of the Voronoi regions. This tighter lower bound is defined as

$$G_{\mathrm{opt}} \geq G_{\mathrm{CS}} = \frac{n+3-2H_{n+2}}{4n(n+1)}\left\{(n+1)(n!)^4 F_n^2\left[\frac{1}{2}\arccos\left(\frac{1}{n}\right)\right]\right\}^{\frac{1}{n}},$$

where $H_m = \sum_{i=1}^{m} 1/i$ represents the harmonic sum and $F_n(\cdot)$ denotes the Schläfli's function [Conway98] which can be recursively computed as

$$F_n(\alpha) = \frac{2}{\pi}\int_{\rho}^{\alpha} F_{n-2}(\beta)\ d\theta,$$

with $\sec(2\beta) = \sec(2\theta) - 2$, $\rho = \frac{1}{2}\arccos(\frac{1}{n-1})$ and initial values $F_0(\alpha) = F_1(\alpha) = 1$. As shown in [Conway85], this bound also asymptotically converges to $G_{\mathrm{CS}} \to 1/(2\pi e)$.

Figure 3.3 depicts the dimensionless second moments of several well-known lattices, such

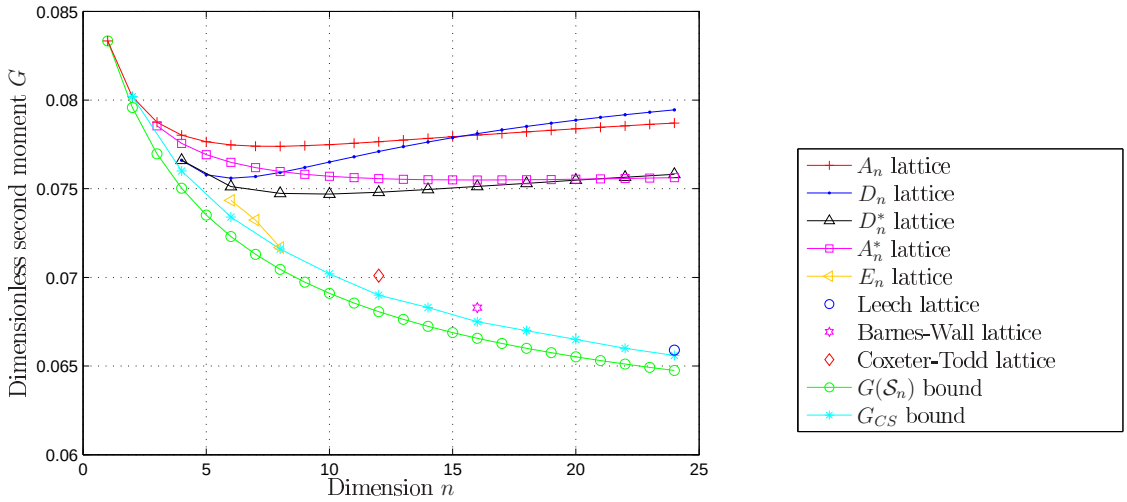

Figure 3.3: Dimensionless second moment of several well-known lattices along with the *sphere bound* and the G_{CS} lower bound.

as the Leech lattice, the Barnes-Wall lattice or the A_n and D_n lattices along with their dual lattices A_n^* and D_n^*, respectively. Additionally, the *sphere bound* and the conjectured new lower bound (G_{CS}) are also shown. The results depicted in this figure show the validity of the two lower bounds reviewed in this section. Also, note that the G_{CS} bound represents a significantly stronger lower bound than $G(\mathcal{S}_n)$, and therefore, it will yield more accurate results when employed to bound the unnormalized transmit power in (3.9).

3.3.2 Upper Bounds

Even though the aforementioned lower bounds for $G\left[\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)\right]$, and by extension lower bounds for E_{SE}, are very useful in order to determine the behavior of VP, it is even more important to set an upper bound to know what minimum performance of VP systems can be expected under ideal conditions. Clearly, by combining the aforementioned upper and lower bounds, the performance range of VP can be determined.

As stated in [Zamir96][Krithivasan07], it is possible to upper bound the dimensionless second moment of a polytope $\mathcal{R}$ in terms of the efficiency of its covering as

$$G(\mathcal{R}) \leq G(\mathcal{S}_n)\, \frac{n+2}{n}\, \chi^2.$$

By combining the previous expression with Equations (3.2) and (3.3), we get an upper bound in terms of the covering density:

$$G(\mathcal{R}) \leq \frac{\Theta^{2/n}}{n V_n^{2/n}}. \tag{3.11}$$

Furthermore, according to [Rogers58b][Rogers58a], the covering density of a lattice can

N	2	3	4	5	6	7	8
M=N	2.068	0.126	0.225	63.80	85.38	118.2	203.5
M=N+1	1.892	0.009	0.011	4.179	7.277	51.48	11.94
M=N+2	16.44	13.42	1.481	0.246	0.076	0.056	10.06

Table 3.1: Values for $(1-\eta)10^3$ used in the G_{TB} bound.

be upper bounded by $\Theta \leq n\log n + n\log\log n + 5n$, which leads to the following upper bound

$$G(\mathcal{R}) \leq G_{\mathrm{Rog}} = \frac{(n\log n + n\log\log n + 5n)^{2/n}}{nV_n^{2/n}}.$$

Despite its validity as an upper bound, G_{Rog} does not approximate the actual value of $G[\mathcal{V}(\Lambda_{\boldsymbol{M}}, 0)]$ tightly and derives in too pessimistic performance results for VP. This is the reason why a tighter upper bound is proposed following the approach of [Messenger09], where the covering density is studied for random template banks. In this case, instead of the strict requirement of complete coverage, a certain mismatch is introduced, that is, complete coverage is required only with a certain confidence degree η. As stated in [Messenger09], the covering density can be then expressed as

$$\Theta = -\log(1-\eta).$$

The values for η have been obtained by extensive simulations and are shown in Table 3.1. Plugging this latest formula into (3.11), the expression for the proposed upper bound is obtained:

$$G(\mathcal{R}) \leq G_{\mathrm{TB}} = \frac{\left(\log\frac{1}{1-\eta}\right)^{2/n}}{nV_n^{2/n}}.$$

Figure 3.4 provides values for E_{SE} averaged over at least 1000000 channel realizations and considering uniform power distribution among users, i.e. $\boldsymbol{B} = \boldsymbol{I}_N$. The upper and lower bounds discussed and presented in this section are also shown in this figure. As one can notice, all the bounds converge to E_{SE} as the amount of users increases. However, the proximity of both lower bounds to the actual E_{SE} value is considerably higher, while the G_{Rog} upper bound differs significantly, specially when a reduced number of antennas is used. As for the proposed G_{TB} upper bound, its performance improves noticeably as more antennas are deployed, yielding a similar average power to E_{SE} with $N = 8$.

3.4 Bounded Sum Rate in VP Systems

The analysis of the sum rate and other performance-related features of VP systems entails several difficulties. On one hand, the noise is affected by the non-linear modulo operation

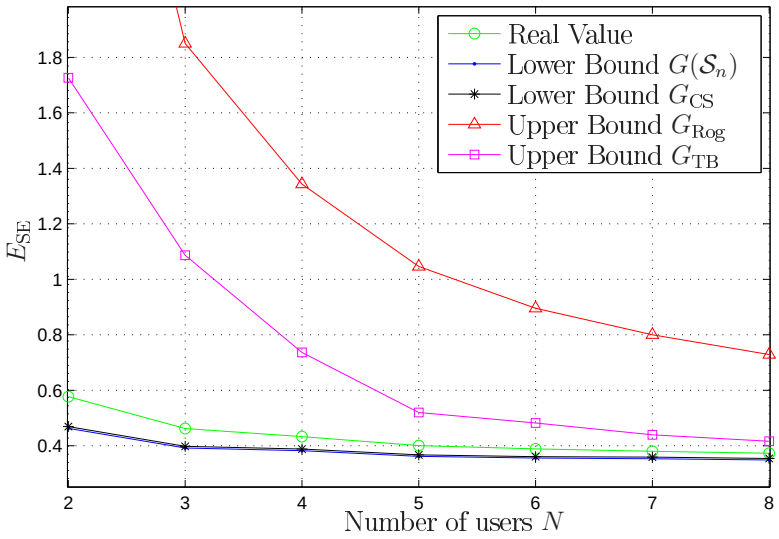

Figure 3.4: Unnormalized transmit power E_{SE} along with the upper bounds (G_{TB} and G_{Rog}) and lower bounds (*sphere bound* and G_{CS}) derived for this parameter. The provided results are given as a function of the number of single-antenna users N.

which maps it to the region $\mathcal{Q}$. On the other hand, the computation of E_{SE} involves the expectation over the result of a closest-point search problem. Due to these facts, the analysis of the sum rate of VP systems has been an open issue until recently, when a closed-form expression in terms of the unnormalized transmit power has been proposed in [Ryan09]. Due to the relevance of this work, its main conclusions and results will be summarized here.

The mutual information between transmitted and detected symbols for each user k can be computed by means of the following expression:

$$I(\hat{s}_k; s_k) = H(\hat{s}_k) - H(\hat{s}_k|s_k), \tag{3.12}$$

where $H(x)$ represents the entropy of the variable x and $H(x|y)$ stands for the entropy of x conditional on y. As discussed in [Ryan09], being $\hat{\boldsymbol{s}}$ uniformly distributed over a $2N$-dimensional cube, it follows that $H(\hat{s}_k) = 0$. Additionally, we get that $H(\hat{s}_k|s_k) = 2H(\hat{\varsigma}_k|\varsigma_k)$, where ς_k represents the real part (or equivalently the imaginary part) of s_k, that is $\varsigma_k = \Re(s_k)$. From (3.7) we infer that the conditional entropy $H(\hat{\varsigma}_k|\varsigma_k)$ equals the entropy of the real (or imaginary) noise component after the modulo operation, namely $H(\hat{\varsigma}_k|\varsigma_k) = H(\zeta_k)$, with $\zeta_k = \Re[\mathrm{Mod}(\beta b_k^{-1/2} w_k)]$. Due to this non-linear operator, ζ_k has a Gaussian modulo distribution with variance $\gamma_k = \beta^2 b_k^{-1}/2$ given by [Razi09]

$$f(\zeta_k) = \begin{cases} \sum_{t=-\infty}^{\infty} \frac{1}{\sqrt{2\pi\gamma_k}} e^{-\frac{|t-\zeta_k|^2}{2\gamma_k}} & \zeta_k \in [-\frac{1}{2}, \frac{1}{2}] \\ 0 & \zeta_k \notin [-\frac{1}{2}, \frac{1}{2}]. \end{cases}$$

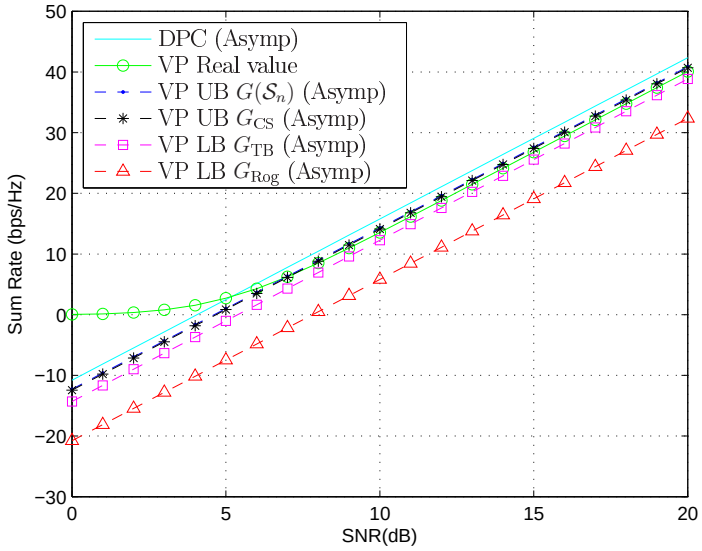

Figure 3.5: Sum rate vs. SNR for asymptotic DPC, VP, and bounded asymptotic VP for an 8×8 antenna system.

Therefore, the entropy of the noise component after the modulo operation is

$$\begin{aligned} H(\zeta_k) &= \int_{-\frac{1}{2}}^{\frac{1}{2}} f(\zeta_k) \log_2 \left[f\left(\zeta_k\right) \right] \, \mathrm{d}\zeta_k \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \sum_{t=-\infty}^{\infty} \frac{1}{\sqrt{2\pi\gamma_k}} e^{-\frac{|t-\zeta_k|^2}{2\gamma_k}} \log_2 \left(\sum_{q=-\infty}^{\infty} \frac{1}{\sqrt{2\pi\gamma_k}} e^{-\frac{|q-\zeta_k|^2}{2\gamma_k}} \right) \mathrm{d}\zeta_k. \end{aligned}$$

By plugging this latest result into (3.12) we ultimately get the individual user rate in a VP system:

$$R_{\mathrm{VP},k} = I(\hat{s}_k; s_k) = -\log_2 \left(2\pi e \gamma_k\right) + 2\Omega(\gamma_k),$$

where the term $\Omega(\gamma)$ captures the effect of the modulo operation on the Gaussian noise. This parameter is defined as

$$\Omega(\gamma) = \frac{\log_2(e)}{2} + \int_{-\frac{1}{2}}^{\frac{1}{2}} \sum_{t=-\infty}^{\infty} \frac{1}{\sqrt{2\pi\gamma}} e^{-\frac{|t-\zeta|^2}{2\gamma}} \left(\log \sum_{q=-\infty}^{\infty} e^{-\frac{|q-\zeta|^2}{2\gamma}} \right) \mathrm{d}\zeta. \tag{3.13}$$

As argued in [Razi09], $\Omega(\gamma)$ is an increasing function of γ, which results in $\Omega(\gamma) \to 0$ as $E_{\mathrm{Tr}} \to \infty$. As a consequence of this, the impact of the modulo-filtered noise can be neglected when a high-SNR analysis is carried out. Thus, the asymptotic rate for user k in the VP system under study is simplified to

$$R_{\mathrm{VP},k} \cong \log_2 \left(\frac{E_{\mathrm{Tr}} \, b_k}{\pi e E_{\mathrm{SE}}} \right), \tag{3.14}$$

where $a \cong b$ means that $\lim_{E_{\mathrm{Tr}}\to\infty} a - b = 0$. Upper and lower bounds on this expression can be obtained by combining this latest result with (3.9) and the bounds derived in Section 3.3. Note that, given the position of E_{SE} within the previous expression, the upper and lower bounds on this parameter yield lower and upper bounds on the sum rate, respectively.

Figure 3.5 depicts the sum rate capacity of an 8×8 antenna system with $\boldsymbol{B} = \boldsymbol{I}_N$, along with the asymptotic capacity of DPC, VP, and the upper and lower bounded asymptotic VP. As is shown in this plot, the asymptotic VP expression yielded by the *sphere* and G_{CS} bounds is remarkably tight to the real VP sum rate. The asymptotic VP sum rate computed by means of the G_{TB} upper bound also closely approximates the real value of the VP sum rate. Additionally, the provided simulation results show the small performance gap between the theoretical upper bound set by DPC and VP when a moderate amount of antennas is used, e.g. $N = 8$.

3.5 Optimization Problems for Asymptotic VP

In this section, the problem of assessing the performance of VP systems is extended to the analysis of optimization procedures, where the transmit power assigned to each user is determined following a certain design criterion. To that end, the upper and lower bounds on the performance of VP systems described in Section 3.3 will be used to facilitate the task of solving the optimization problems. More specifically, three different optimization procedures with respect to the user rates are investigated: weighted sum rate maximization, QoS optimization and rate balancing. For the sake of simplicity, the high-SNR regime will be considered for the remainder of the section, where the effect of the modulo-filtered noise (3.13) can be neglected.

In order to highlight the performance gain of non-linear precoding systems with respect to their linear counterparts, the aforementioned optimization problems are additionally solved for a linear precoder. In this case, an equivalent system to the one depicted in Figure 3.2 is considered, where the signal perturbation stage and the modulo operator are removed. The resulting linear precoding system is shown in Figure 3.6. Additionally, given that there is no perturbation process or non-linear operators, the more conventional Gaussian signaling will be adopted in this case.

Finally, ergodic expressions and simulation results on the performance gap between VP and linear precoding are provided for the aforementioned applications.

3.5.1 Weighted Sum Rate

This section analyzes the weighted sum rate (WSR) maximization for VP and linear precoding. The optimal matrix $\boldsymbol{B}$ and the maximum achievable weighted sum rate with non-

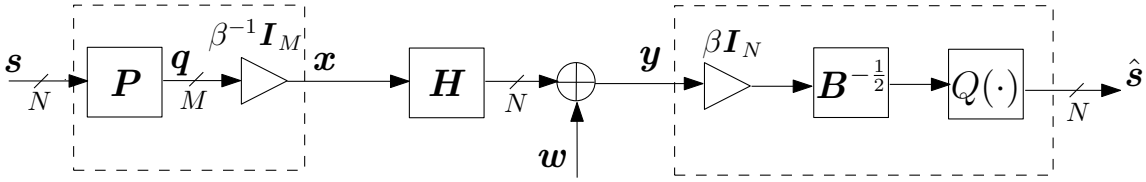

Figure 3.6: Linear precoding system with power loading matrix $\boldsymbol{B}^{-\frac{1}{2}}$.

negative weights $\boldsymbol{v} = [v_1, \ldots, v_K]^T > \boldsymbol{0}$ are derived for the high-SNR region based on the following problem formulation:

$$\text{WSR} = \max_{\{b_k\}} \sum_{k=1}^{N} v_k R_k \quad \text{s.t.: } b_k \geq 0 \ \forall k. \tag{3.15}$$

Additionally, analytical expressions for the performance gap between VP and linear precoding will be provided in the framework of both the weighted and the regular sum rate capacity.

3.5.1.1 Linear Precoding

For the linear precoding system with Gaussian signaling depicted in Figure 3.6, the achievable rate for user k can be expressed as a function of the signal to interference-plus-noise ratio (SINR) associated with the received signal at each user terminal, namely

$$y_k = \frac{1}{\beta} b_k^{1/2} s_k + w_k.$$

Given that the ZF premise enforces null-interference, and considering that $w_k \sim \mathcal{CN}(0,1)$, the asymptotic rate of a certain user k in a linear precoding system can be expressed as

$$R_{\text{lin},k} \cong \log_2 \left(\frac{E_{\text{Tr}}\ b_k}{E_{\text{SE,lin}}} \right), \tag{3.16}$$

where the power of the precoded symbols in computed as

$$E_{\text{SE,lin}} = \text{E}\left[\left\| \boldsymbol{H}^{\dagger} \boldsymbol{B}^{\frac{1}{2}} \boldsymbol{s} \right\|_2^2 \right] = \text{Tr}\left[\boldsymbol{B} \left(\boldsymbol{H}\boldsymbol{H}^H \right)^{-1} \right]. \tag{3.17}$$

Hence, combining the expressions (3.16) and (3.17) the asymptotic user rate for the linear system described in Figure 3.6 is obtained:

$$R_{\text{lin},k} \cong \log_2 \left(\frac{E_{\text{Tr}}\ b_k}{\sum_{i=1}^{N} b_i \alpha_i} \right), \tag{3.18}$$

where α_i represents the i^{th} diagonal element of the matrix $\boldsymbol{W}^{-1}$. Note that weighting all b_k-s with the same scalar, i.e., $b'_k = cb_k \; \forall k$ for some $c \neq 0$, does not change the rate in (3.18). Consequently, and for the sake of analytical simplicity, we set $\sum_i b_i \alpha_i = 1$. Solving the optimization problem in (3.15) for linear precoding in compliance with the Karush-Kuhn-Tucker conditions leads to the optimal coefficients

$$b_k = \frac{v_k}{\alpha_k V}, \tag{3.19}$$

with $V = \sum_{k=1}^{N} v_k$. Therefore, the resulting asymptotically optimum weighted sum rate for linear precoding is derived:

$$\text{WSR}_{\text{lin,opt}} = V \log_2 \left(\frac{E_{\text{Tr}}}{V} \right) + V \log_2 \left[\prod_{i=1}^{N} \left(\frac{v_i}{\alpha_i} \right)^{v_i/V} \right]. \tag{3.20}$$

For the case of unitary and equal weights, the solution to the weighted sum rate optimization in (3.20) reduces to the asymptotically optimum conventional sum rate

$$\text{SR}_{\text{lin,opt}} = N \log_2 \left(E_{\text{Tr}} \right) - N \log_2 \left(N \right) - \log_2 \left(\prod_{i=1}^{N} \alpha_i \right).$$

3.5.1.2 Vector Precoding

The optimization problems presented in this section will make use of the asymptotic expression given in (3.14) for the information transmission rate of VP systems.

Similarly to the linear case, weighting all b_k-s with the same scalar, i.e., $b'_k = cb_k \; \forall k$ for some $c \neq 0$, does not change the VP rate in (3.14) due to the dependence of E_{SE} on $\boldsymbol{B}$ [see (3.8)]. Since there is no closed-form expression for E_{SE}, it is necessary to approximate this value by an upper bound when performing the weighted sum rate maximization. As all the bounds discussed in Section 3.3 lead to the unsatisfactory result that only one b_k is different from zero, the ultimate upper bound for E_{SE} will be used instead, i.e., the unnormalized transmit power of a linear precoder $E_{\text{SE}} \leq \sum_{i=1}^{K} b_i \alpha_i$. Solving the weighted sum rate optimization problem in (3.15) leads to the same optimal power-loading coefficients as in (3.19). Combining the latter with (3.9) and (3.14) results in the following expression for the weighted sum rate:

$$\begin{aligned} \text{WSR}_{\text{VP,opt}} \geq \; & \text{WSR}_{\text{lin,opt}} - V \log_2 \left(\zeta \right) \\ & + \underbrace{V \log_2 \left(\prod_{j=1}^{N} \alpha_j^{1/N} \right) - V \log_2 \left(\left| \boldsymbol{W}^{-\frac{1}{N}} \right| \right)}_{\Upsilon(\boldsymbol{H})} + \underbrace{V \log_2 \left(\frac{V}{N} \prod_{i=1}^{N} v_i^{-\frac{1}{N}} \right)}_{\Psi(\boldsymbol{v})}, \end{aligned} \tag{3.21}$$

with $\zeta = 2\pi e G_{\mathrm{UB}}$, where G_{UB} stands for an upper bound on $G[\mathcal{V}(\Lambda_{\boldsymbol{P}}, 0)]$. This upper bound can be replaced by any of the bounds analyzed in the previous section, although the G_{TB} bound is recommended for accurate results. Note that the loss term involving ζ results from the employment of uniformly distributed symbols in the VP model.

In the particular case of $v_k = 1\ \forall k$, the expression for the conventional sum rate is obtained, which reads as

$$\mathrm{SR}_{\mathrm{VP,opt}} \geq N \log_2\left(\frac{E_{\mathrm{Tr}}}{N}\right) - \log_2\left(\left|\boldsymbol{W}^{-1}\right|\right) - N \log_2\left(\zeta\right).$$

3.5.1.3 Ergodic Performance Gap

When examining the performance gap between linear precoding and VP in terms of the weighted sum rate (3.21), one can notice that the superiority of the non-linear scheme is subjected to the channel conditions, the transmit data signaling model and the distribution of the weights. Nevertheless, note that the terms $\Psi(\boldsymbol{v})$ and $\Upsilon(\boldsymbol{H})$ are always positive or equal to zero, and hence, they will always contribute to the performance advantage of VP over linear precoding. More specifically, the gain following from different weights, denoted as $\Psi(\boldsymbol{v})$, is non-negative for any weight distribution, since the arithmetic mean V/N is always larger than or equal to the geometric mean of the weights as stated by the arithmetic mean - geometric mean (AM-GM) inequality. Additionally, the Hadamard's inequality [Hadamard93] stipulates that for a given hermitian and semi-definite positive matrix $\boldsymbol{G}$ with diagonal values g_{ii} $i \in \{1, \ldots, Z\}$ the following inequality holds:

$$|\boldsymbol{G}| \leq \prod_{i=1}^{Z} g_{ii},$$

following from which we get that $\Upsilon(\boldsymbol{H}) \geq 0$.

From Equation (3.21), the performance gap between the analyzed precoding schemes can be shown for specific channel conditions and design features. For example, for diagonal channels and equal weights $v_j = v_1\ \forall j$, the performance gap is reduced to

$$\mathrm{WSR}_{\mathrm{VP,opt}} = \mathrm{WSR}_{\mathrm{lin,opt}} - N v_1 \log_2(\zeta).$$

The loss term $N v_1 \log_2(\zeta)$ stems from the use of a less advantageous signaling model by the vector precoder. Therefore, when considering the same distribution of the user data symbols, the performance of VP and linear precoding is equal in terms of the weighted sum rate. This is an expected result since the perturbation vector chosen by the VP algorithm in the event of a diagonal channel will always be the null vector, which reduces the non-linear precoding scheme to a simple linear precoder with a modulo operator at the receivers.

Nevertheless, this statement cannot be generalized to any other channel realization, and therefore, the computation of the ergodic performance gap is required in order to gain some insight into the expected performance advantage of VP over the more straightforward linear precoding.

If the expectation of the weighted sum rate gap is taken with respect to $\boldsymbol{H}$, the asymptotic ergodic gain of VP can be expressed as

$$\mathrm{E}[\mathrm{WSR}_{\mathrm{VP}} - \mathrm{WSR}_{\mathrm{lin}}] \geq \Psi(\boldsymbol{v}) - V\log_2(\zeta) + \frac{V}{N\log(2)}\sum_{i=1}^{N-1}\frac{i}{M-i}. \tag{3.22}$$

This result follows from $\boldsymbol{W}$ being a Wishart distributed matrix of dimension N and M degrees of freedom, namely $\boldsymbol{W} \sim \mathcal{W}_N(M, \boldsymbol{I}_N)$, whose ergodic characteristics have been widely studied. For instance, in [Tulino04] it was shown that for a Wishart distributed matrix the following property holds

$$\mathrm{E}[\log|\boldsymbol{W}|] = \sum_{\ell=0}^{N-1}\psi(M-\ell), \tag{3.23}$$

where $\psi(\cdot)$ stands for the Digamma function [Tulino04], which is defined as

$$\psi(m) = -\gamma + \sum_{\ell=1}^{m-1}\ell^{-1},$$

being γ the Euler-Mascheroni constant. Additionally, note that the α_i-s are inverse Wishart distributed [Gupta99], namely $\alpha_i \sim \mathcal{W}_1^{-1}(M-N+1, 1)$, which leads to

$$\mathrm{E}[\log(\alpha_i)] = -\psi(M-N+1). \tag{3.24}$$

By plugging the ergodic expressions in (3.23) and (3.24) into the performance gap definition, the equation for the expected performance gap in (3.22) is obtained.

As an additional result, specializing to $v_k = 1\ \forall k$ gives the ergodic sum rate gain of VP over linear precoding:

$$\mathrm{E}[\mathrm{SR}_{\mathrm{VP}} - \mathrm{SR}_{\mathrm{lin}}] \geq \frac{1}{\log(2)}\sum_{l=1}^{N-1}\frac{l}{M-l} - N\log_2(\zeta).$$

3.5.1.4 Simulation Results

The results shown in Figure 3.7 represent the achievable gain range of VP with respect to linear precoding. The bounding of the VP results has been performed by means of the G_{TB} and G_{CS} bounds. Additionally, the actual value of the analyzed performance gap is plotted

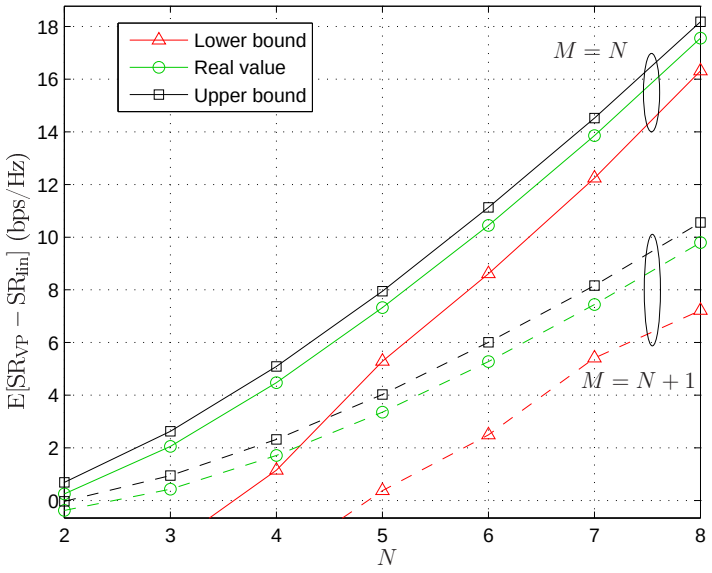

Figure 3.7: Ergodic performance gap on the sum rate for systems with $M = N$ and $M = N + 1$.

as a reference. Note that the values given for the sum rate can also be employed to compute the gain for the weighted sum rate (for $V = N$), in which case the addition of the positive factor $\Psi(\boldsymbol{v})$ would be required.

As is shown in Figure 3.7, a better performance of VP with respect to linear precoding can be expected for all the analyzed antenna setups. What is more, a more significant gain of VP can be achieved in those systems with the same amount of transmit and receive antennas. This is due to the well-known performance loss experienced by linear ZF precoders in fully-loaded systems. Additionally, the simulation results depicted in this figure show that the performance gain increases noticeably as more antennas are deployed, achieving a performance gain of roughly 18 bps/Hz with just 8 transmit and receive antennas.

3.5.2 Quality of Service

This section deals with the QoS optimization, which establishes the premise that a certain rate for each user has to be guaranteed with a minimum effort. In other words, the transmit power is minimized under constraints ensuring the minimum rates ρ_k for all users. The QoS optimization problem can be described by the following problem formulation:

$$\min_{\{b_k\}} E_{\mathrm{Tr}} \qquad \text{s.t.: } R_k \geq \rho_k \ \forall k. \tag{3.25}$$

In the following sections, the optimal matrix $\boldsymbol{B}$ and the minimum achievable transmit power will be derived for the linear and VP precoding approaches.

3.5.2.1 Linear Precoding

For a linear precoder with the transmit rate defined in (3.18), it can be shown that the constraint in (3.25) is fulfilled for

$$b_k = 2^{\rho_k} \frac{\sum_{i=1}^{N} \alpha_i b_i}{E_{\mathrm{Tr,lin}}}.$$

Taking the sum $\sum_{k=1}^{N} \alpha_k b_k = E_{\mathrm{Tr,lin}}$ shows that the minimum power for linear precoding is given by

$$E_{\mathrm{Tr,lin}} = \sum_{k=1}^{N} \alpha_k 2^{\rho_k}. \tag{3.26}$$

3.5.2.2 Vector Precoding

For vector precoding, the combination of the expression in (3.14) and the active QoS constraints leads to

$$b_k = \frac{\zeta N |\boldsymbol{W}|^{-1/N} |\boldsymbol{B}|^{1/N}}{E_{\mathrm{Tr}}} 2^{\rho_k}.$$

The minimum power for the VP system under study is then obtained by analyzing $\prod_{k=1}^{N} b_k$, which gives

$$E_{\mathrm{Tr,VP}} \leq \zeta N |\boldsymbol{W}|^{-1/N} \prod_{k=1}^{N} 2^{\frac{\rho_k}{N}}. \tag{3.27}$$

3.5.2.3 Ratio of Ergodic Powers

The analysis of the ratio of ergodic powers for any user rate ρ_k is hindered by the different structures of the solutions in (3.26) and (3.27). Therefore, in order to ease the evaluation of the expected power of the precoding techniques under a QoS formulation, the particular case of equal rates will be considered. The expression of the ratio of powers therefore reads as

$$\frac{E_{\mathrm{Tr,VP}}}{E_{\mathrm{Tr,lin}}} \leq \zeta \underbrace{\frac{|\boldsymbol{W}|^{-1/N}}{\mathrm{Tr}(\boldsymbol{W}^{-1})/N}}_{\Xi},$$

which follows from $\sum_{j=1}^{N} \alpha_j = \mathrm{Tr}(\boldsymbol{W}^{-1})$. Once more, due to the AM-GM inequality, the second term in this expression is known to be $\Xi \leq 1$. Consequently, the superiority of the VP approach is subjected to the ζ term that stems from the presence of uniformly distributed symbols in the VP model. As is the case with the weighted sum rate optimization, the effect of this term can be neglected if the same signaling model is considered for both precoding schemes.

In any case, in the limit of the system, e.g. $N \to \infty$, we get that $\zeta \to 1$ due to the convergence of the upper and lower bounds discussed in Section 3.3, which indicates that VP is able to offer a certain QoS with a smaller transmit power than linear precoding under such conditions.

In order to evaluate the power requirements of the linear and VP precoding models in a more general scenario, the ratio of ergodic powers will be computed. To that end, the following property of Wishart distributed matrices [Tulino04] will be used along with the ergodic expression in (3.24):

$$\mathrm{E}[\mathrm{Tr}(\boldsymbol{W})^{-1}] = \frac{N}{M-N}.$$

Thus, the desired ergodic power ratio reads as:

$$\frac{\mathrm{E}[E_{\mathrm{Tr,VP}}]}{\mathrm{E}[E_{\mathrm{Tr,lin}}]} \leq \zeta N \left(\frac{M}{N} - 1\right) \prod_{\ell=0}^{N-1} \frac{\Gamma(M - \frac{1}{N} - \ell)}{\Gamma(M-\ell)}.$$

Note that the ratio of ergodic powers given by the above expression equals 0 for $M = N$. This follows from the well-known effect of a fully-loaded antenna setup in linear ZF precoders which results in $\mathrm{E}[E_{\mathrm{SE,lin}}] = \infty$ [Peel05].

3.5.2.4 Simulation Results

The ratio of ergodic powers between the vector and linear precoding approaches is depicted in Figure 3.8 for $M = N + 1$ and $M = N + 2$ antenna setups. The results for the upper and lower bounds shown in this figure have been acquired by means of the G_{TB} and G_{CS} bounds, respectively.

For a small number of single-antenna users, e.g. $N = 2$, the proposed upper bound does not yield very accurate results for any of the antenna setups. Nevertheless, as the amount of users increases, both bounds converge and approach the real value of the ratio.

As for the performance of the different precoding schemes, the advantage of using VP over linear precoding is more notorious as more users are added to the system. Specifically, for a $M = 5$ and $N = 4$ system configuration, the provided simulation results show that VP requires half of the power used by the linear approach for the same QoS constraint. Nonetheless, the data depicted in Figure 3.8 shows that the advantage of using VP over the more straightforward linear approach fades as the amounts of antennas at both ends of the communication link are less balanced. Put in other words, higher power ratios are obtained for the $M = N + 2$ antenna setup when compared to the $M = N + 1$ configuration or the fully-loaded system model.

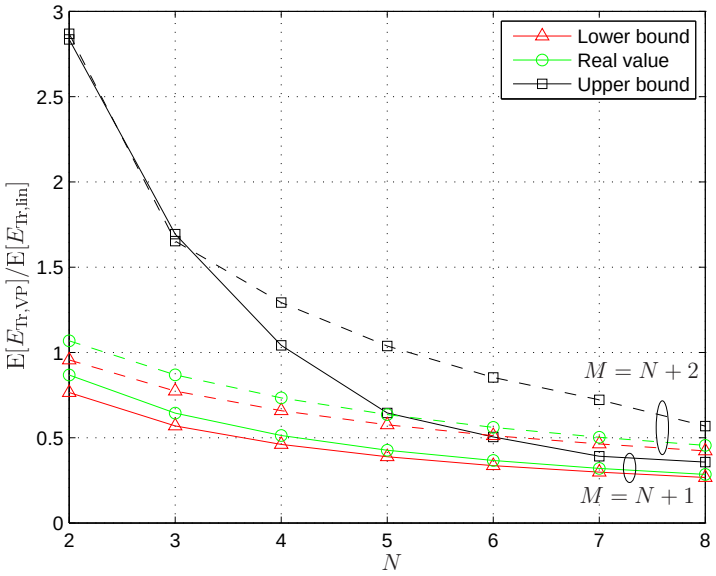

Figure 3.8: Ratio of ergodic powers of the linear and VP precoding approaches.

3.5.3 User Balancing

When trying to achieve a fair resource allocation, a rate balancing optimization can be employed to maximize the rates of the users, where the ratio with respect to some requirements $\rho_k \ \forall k$ is the same for all users. The problem formulation for the user balancing optimization is as follows

$$\max_{\delta,\{b_k\}} \delta \qquad \text{s.t.: } R_k = \delta \rho_k \ \forall k.$$

3.5.3.1 Linear Precoding

The solution to the aforementioned optimization problem for linear precoding systems (constraining the total transmit power to E_{Tr}) can be found by solving the following equation:

$$\frac{1}{E_{\mathrm{Tr}}} \sum_{i=1}^{N} \alpha_i \ 2^{\delta \rho_i} = 1,$$

where we set $\sum_{j=1}^{N} b_j \alpha_j = 1$ without loss of generality. In order to get the value of 2^{δ}, it is required to solve a root finding problem with real exponents ρ_i, which greatly hinders the computation of the optimal δ. However, it is possible to obtain a closed-form solution for the particular case of equal rate requirements, i.e. $\rho_k = \rho \ \forall k$. In such a scenario, the asymptotic value of the balancing parameter is

$$\delta_{\mathrm{lin}} = \frac{1}{\rho} \log_2 \left(\frac{E_{\mathrm{Tr}}}{\sum_{i=1}^{N} \alpha_i} \right).$$

3.5.3.2 Vector Precoding

The balancing optimization problem for VP can be easily solved by combining (3.9) and (3.14). Given that the entries of the matrix $\boldsymbol{B}$ can be scaled arbitrarily without affecting the performance of the system, the product $\prod_{i=1}^{K} b_i^{1/N} = 1$ is taken, which leads to the following optimum balancing parameter for VP

$$\delta_{\mathrm{VP}} \geq \frac{N}{\sum_{i=1}^{N} \rho_i} \log_2 \left(\frac{E_{\mathrm{Tr}}}{\zeta N |\boldsymbol{W}|^{-\frac{1}{N}}} \right).$$

3.5.3.3 Ergodic Performance Gap

Due to the limitation set by the linear precoding approach, only the case for $\rho_k = \rho \ \forall k$ will be analyzed. The performance gap for the balancing problem under such conditions is represented as follows

$$\delta_{\mathrm{VP}} - \delta_{\mathrm{lin}} \geq \frac{1}{\rho} \left[-\log_2 (\Xi) - \log_2(\zeta) \right].$$

As is the case with the previous optimization problems, the knowledge of $\Xi \leq 1$ is insufficient to determine the superiority of the VP approach for an arbitrary channel realization. Consequently, the expected performance gap will be computed, which reads as

$$\mathrm{E}[\delta_{\mathrm{VP}} - \delta_{\mathrm{lin}}] \geq \frac{1}{\rho} \left(\mathrm{E} \left[\log_2 \frac{\mathrm{Tr}(\boldsymbol{W}^{-1})}{|\boldsymbol{W}^{-\frac{1}{N}}|} \right] - \log_2(\zeta N) \right).$$

Unfortunately, evaluating $\mathrm{E}\left\{\log_2\left[\mathrm{Tr}(\boldsymbol{W}^{-1})\right]\right\}$ is difficult. Thus, a closed-form expression has been derived only for $N = 2$ and $M \geq N$. For such case, the problematic ergodic expression can be rewritten as

$$\mathrm{E}\left\{\log_2\left[\mathrm{Tr}(\boldsymbol{W}^{-1})\right]\right\} = \log_2(e) \left\{ \mathrm{E}\left[\log(\varpi)\right] - \mathrm{E}\left[\log|\boldsymbol{W}|\right] \right\},$$

where $\varpi = W_{11} + W_{22}$ is a Wishart distributed variable with $2M$ degrees of freedom, namely $\varpi \sim \mathcal{W}_1(2M, 1)$. We now apply the properties of Wishart matrices described in (3.23) and (3.24) to the equation above. The expression so obtained finally simplifies to

$$\mathrm{E}[\delta_{\mathrm{VP}} - \delta_{\mathrm{lin}}] \geq \frac{\log_2(e)}{\rho} \left[\sum_{\ell=M-1}^{2M-1} \ell^{-1} - \frac{1}{2(M-1)} - \log_2 \left(\frac{\zeta N}{e} \right) \right]. \tag{3.28}$$

3.5.3.4 Simulation Results

Figure 3.9 depicts the ergodic performance gap between linear precoding and VP. The data shown in this figure have been obtained through extensive simulation, except for the case of

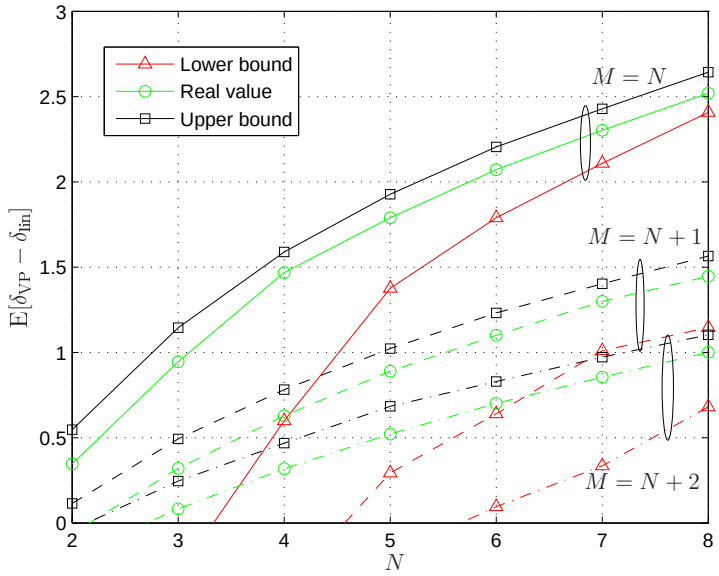

Figure 3.9: Ergodic performance gap for the user balancing optimization.

$N = 2$, where the expression in (3.28) has been used instead. The upper and lower bounds have been computed by means of the G_{CS} and G_{TB} bounds, respectively.

The data in Figure 3.9 shows the accuracy of the upper bound for any of the antenna setups under study. Nevertheless, the precision of the lower bound is strongly dependant on the transmit and receive antenna configuration. This way, more accurate results of the lower bound can be expected for those antenna configurations where the amounts of transmit and receive antennas are more balanced.

As for the ergodic difference of the balancing parameter, the data in Figure 3.9 shows the conventional performance advantage of VP with respect to linear precoding in the fully-loaded antenna setup. What is more, the performance enhancement obtained when deploying an additional antenna at both the transmitter and the receiver side is more prominent in the $M = N$ antenna configuration. Generally speaking, it can be observed that the slopes of the $\mathrm{E}[\delta_{\mathrm{VP}} - \delta_{\mathrm{lin}}]$ vs. N curves depend on the excess of transmit antennas, in such a way that the smaller the M/N ratio is, the steeper the curves are.

3.6 Chapter Summary

In this chapter, the problem of assessing the performance of VP systems has been addressed. The main difficulties in this analysis are twofold: the addition of the perturbation signal, which results from the evaluation of an NP-hard problem, and the modulo operators located at the receivers, which map the received signal and the noise into a certain Voronoi region.

The performance of any precoding system is strongly related to the power of the precoded symbols. Nevertheless, the expected power of the symbols precoded following a VP approach

cannot be evaluated due to the data perturbation process. To overcome this issue, several upper and lower bounds have been reviewed and proposed in order to delimit the range of the unnormalized transmit power, and ultimately, bound the performance of VP.

Three different applications have then been analyzed for the proposed bounds: the weighted sum rate maximization, the power resulting from a QoS formulation and the performance when balancing the rates of the users. With the aim of proving the superiority of VP over the more straightforward linear precoding approaches, the aforementioned optimization problems have also been solved for a ZF linear precoder. Next, the ergodic performance gap with respect to VP has been presented and analyzed.

Simulation results on the sum rate and the ergodic performance gap with respect to linear precoding for the analyzed optimization problems have also been provided. The results presented in this chapter show a better performance of VP for all the applications and antenna setups under study. The performance gain of the non-linear technique is more prominent in those systems with a higher amount of transmit/receive antennas as well as in fully-loaded systems, e.g. $M = N$.

The performance results of VP systems given in this chapter include a penalty term that stems from a uniform symbol distribution, whereas the analysis of the linear precoder has been carried out considering the more advantageous Gaussian signaling. In a realistic scenario, where the symbols to be transmitted are taken from a finite constellation, the linear precoding model studied in this chapter would suffer from an added performance degradation. On the other hand, if the order of the modulation is high enough, the distribution of transmit symbols will approach a uniform distribution, and therefore, only a slight performance degradation with respect to the provided performance results of the non-linear technique is to be expected in a practical implementation.

All in all, the gain of VP over linear precoding has been shown and documented over several cases of study. Furthermore, due to the assumptions made when analytically assessing the performance of the precoding techniques, the gain of VP over linear precoding is expected to be even higher in a realistic scenario.

Chapter 4

The Fixed-complexity Sphere Encoder

4.1 Introduction

As it has been shown in previous chapters, VP provides a considerable performance advantage with respect to linear precoding techniques or other non-linear precoding schemes such as THP. Nonetheless, the computation of the perturbing signal represents the main challenge for its practical implementation. Since the publication of the first VP model in [Hochwald05], many algorithms have been proposed in the literature to replace the computationally intractable exhaustive search problem posed by the cost function in (2.8).

Lattice reduction approaches have been widely used as a means to compute a suboptimum perturbation vector with a moderate complexity. The key idea of lattice-reduction techniques relies in the usage of an equivalent and more advantageous set of basis vectors to allow for the suboptimal resolution of the problem in (2.8) by means of a simple rounding operation. This method is used in [Windpassinger04], where the Lenstra-Lenstra-Lovász (LLL) reduction algorithm [Lenstra82] is used to yield the Babai's approximate closest point solution [Babai86]. Similar approaches can be found in [Seethaler06], [Hur07] and [Liu07]. Despite achieving full diversity order in VP systems, the performance degradation caused by the quantization error due to the rounding operations still remains. Moreover, many lattice reduction algorithms have a considerable computational complexity, which poses many challenges to a prospective hardware implementation.

The solution to the cost function in (2.8) can also be found by *searching* for the optimum solution within a subset of candidate vectors. To ease the candidate evaluation process, the computation of the vector norms in the aforementioned expression is replaced by a distributed distance computation model, whose structure resembles that of a tree. These approaches, also known as tree-search techniques, perform a traversal through a tree of hypotheses with the aim of finding an appropriate perturbation vector.

One of the most popular tree-search algorithms is the sphere encoder (SE), which represents the adaptation of the sphere decoder (SD) approach in [Viterbo99] [Damen03] to precoding scenarios [Hochwald05] [Schmidt08]. Despite its optimal performance, the vari-

able complexity of the algorithm, as well as its strong dependance on the channel conditions and the noise variance, has motivated the design of other tree-search approaches that achieve a suboptimal performance with a bounded complexity. This is the case of the K-Best algorithm [Anderson84] which was applied to VP scenarios in [Zhang05][Habendorf07]. In spite of its fixed complexity, the required sorting stages represent the main bottleneck of the K-Best tree traversal.

Tree-search techniques can additionally be used to alleviate the performance degradation suffered by lattice reduction approaches. For example, the schemes in [Airy06] and [Park08] incorporate a SE search around the solution provided by the lattice-reduction algorithm to improve the quality of the resulting perturbation vector.

The main focus of the present chapter is on the presentation of a tree-search algorithm for VP that having a close-to-optimum performance is still suitable for a high-throughput hardware implementation.

4.2 Tree-search Algorithms for VP

The distributed distance computation model for tree-search-based precoding systems will be analyzed in this section. Additionally, two of the most widely-used tree-search schemes, namely the SE and the K-Best approaches, will be briefly reviewed.

Note that in the remainder of this thesis the term (Euclidean) *distance* will be used to denote calculations of the form $|a+s|$. The strict definition of the distance between two points would require the evaluation of $|-a'+s|$ instead, where $a'=-a$. Therefore, when referring to the distance between the variables a and s it is implied that the calculation of $|a+s|$ is being alluded to, and not the more conventional $|s-a|$. This notation has been adopted for a better understanding of the tree-search algorithms in precoding scenarios, as their original description contemplates the term *distance* in its common form.

If the matrix decomposition procedure in (2.7) is performed in such a way that a triangular $\boldsymbol{U}$ matrix is obtained, it is possible to gather all the solution vector hypotheses into an organized tree-like structure so that the search for the optimum perturbation vector is facilitated. The nodes at the top level of this tree structure (level $i=N$) are referred to as root nodes and represent each one of the possible values within the set $\tau\mathbb{C}\mathbb{Z}$. The next levels of the tree are composed by expanding all the eligible nodes, namely every element in $\tau\mathbb{C}\mathbb{Z}$, for each one of the nodes expanded at the previous level, which we shall refer to as parent nodes. Additionally, a leaf node will denote any node at the bottom of the tree (level $i=1$). The path leading from a root node to a any node within the tree structure will be referred to as a branch.

This way, the computation of the Euclidean distances in (2.8) can be distributed across multiple stages as

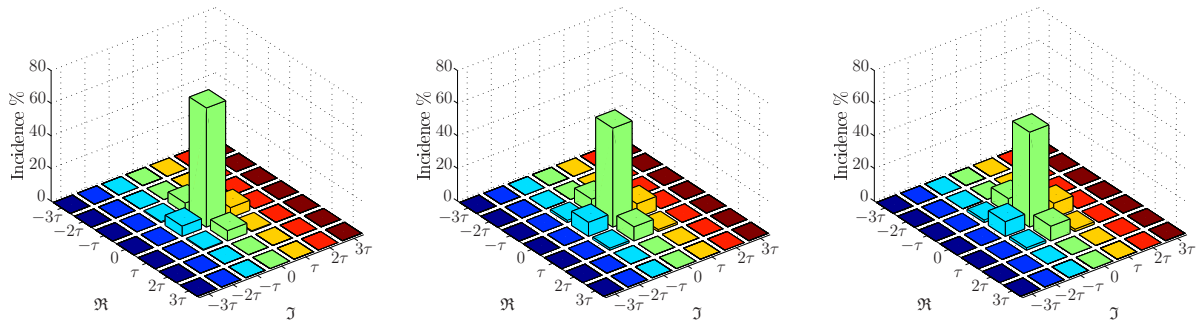

(a) 0 dB (b) 15 dB (c) 30 dB

Figure 4.1: Incidence of the lattice values in the optimum solution vector for different SNR regimes.

$$D_i = u_{ii}^2 |a_i + z_i|^2 + \sum_{j=i+1}^{N} u_{jj}^2 |a_j + z_j|^2 = d_i + D_{i+1}, \tag{4.1}$$

where the intermediate point that reflects the previous perturbation element choices is computed as

$$z_i = s_i + \sum_{j=i+1}^{N} \frac{u_{ij}}{u_{ii}} (a_j + s_j). \tag{4.2}$$

Hence, the search tree is organized in such a way that the Euclidean distances in (2.8) can be computed in a distributed fashion across the different levels of the tree. Specifically, the partial Euclidean distance (PED) of level i is denoted as d_i, while the accumulated Euclidean distance (AED) down to level i is represented by D_i. The search for the perturbation vector is then performed by traversing the tree of N levels (each one representing a user) starting from the root level $i = N$, and working backwards until $i = 1$. Note that, as opposed to the point-to-point MIMO detection scenario, the amount of child nodes that stem from the same parent node does not depend on the modulation constellation in use. Since the elements of the solution vector belong to the expanded search space $\tau\mathbb{CZ}$, the amount of child nodes that originate from each parent node in the tree equals $|\tau\mathbb{CZ}| = \infty$ in theory.

However, depending on the tree-traversal strategy to be followed, the cardinality of this set can be reduced either artificially, by limiting the search space to the group $\mathcal{L}$ of closest points to the origin, or by identifying the set of eligible nodes following a distance control policy (also known as the *sphere constraint*). Regarding the former search set limitation strategy, note that the lattice values that comprise the optimum solution vector in the VP framework are more concentrated around the origin of the lattice in the low-SNR region, as the null vector is a more probable solution in this scenario [Christensen07]. This will derive in a low error-rate performance penalty when artificially constraining the search space to the set of closest elements to the origin. On the other hand, those lattice values that are further

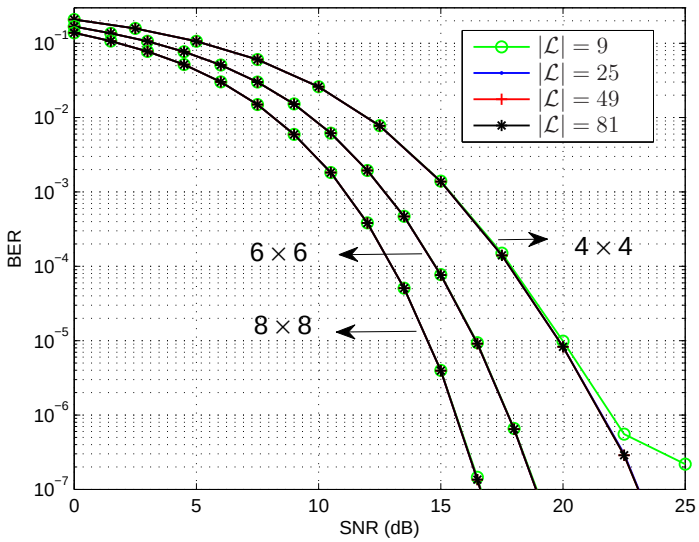

Figure 4.2: Effect of employing a reduced search set on the BER performance of VP systems.

away from the origin are more often selected as part of the optimum perturbation vector at high SNRs. The effect of the SNR on the distribution of the selected perturbation elements is shown in Figure 4.1 for an 8×8 antenna system. As one can notice, the 0 element accounts for roughly 80 % of the vector perturbation elements at low SNR (e.g. 0 dB), whereas it is only present in 60 % of the selected perturbation vectors in the high-SNR regime (e.g. 30 dB).

Consequently, the limitation of the search set provides a computational-complexity advantage to the structure of the tree-search algorithms. We shall define the restricted search space $\mathcal{L}$ as the square region centered around the origin of the lattice such that $\mathcal{L} \triangleq \{x+yj : |x|, |y| \leq (\sqrt{|\mathcal{L}|}-1)/2\}$. The effect of artificially constraining the search space on the error-rate performance of VP systems is shown in Figure 4.2. The data provided in this figure show that the limitation of the cardinality of the search set has a smaller impact on the error-rate performance of the precoding system when a higher number of antennas are deployed, e.g. 6×6 and 8×8 configurations. For the 4 antenna case on the other hand, a noticeable performance degradation is only perceived when a very restricted lattice is used. In the light of these results, a search set of cardinality $|\mathcal{L}| = 25$ will be considered in the remainder of this document, unless otherwise stated.

4.2.1 Sphere Encoder for VP Systems

Depth-first tree-search algorithms traverse the tree recursively in both forward and backward directions. The main feature that identifies this type of tree traversal is that the expansion of the child nodes of a parent node is prioritized over the expansion of its siblings. This way, a leaf node is visited after the first N node expansions allowing for efficient pruning techniques to be implemented based on, for example, the Euclidean distance of the first

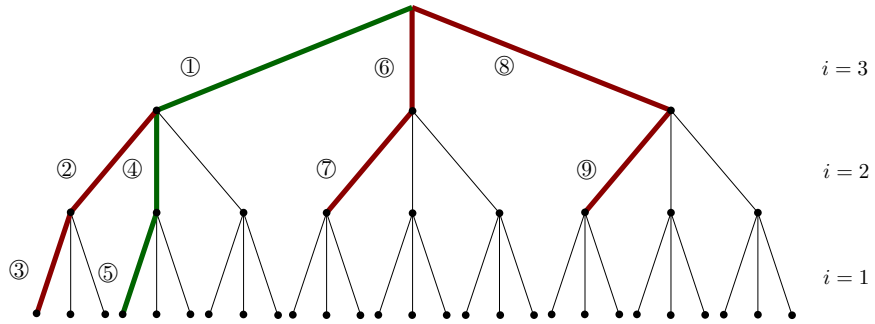

Figure 4.3: SE tree traversal example in a system with $N = 3$ and $|\mathcal{L}| = 3$.

computed branch or statistics of the precoded symbols as proposed in [Habendorf06]. The pruning criterion is updated every time a leaf node with a smaller AED is reached. However, the efficient node expansion scheme comes at the cost of a variable complexity.

One of the most noteworthy depth-first techniques is the SE algorithm, which restricts the search for the perturbation vector to a set of points that lie within a hypersphere of radius R centered around a reference signal. Therefore, the set of eligible points is restricted to those nodes that satisfy the sphere constraint, namely $D_i \leq R$. The radius of the hypersphere can be initialized to $R = \infty$ if the depth-first tree traversal is performed in such a way that the most promising nodes are visited first. The good performance of the SE algorithm is a consequence of the identification and management of the admissible set of nodes at each stage of the tree search.

An arbitrary example of the SE tree traversal through an $N = 3$ and $|\mathcal{L}| = 3$ tree is depicted in Figure 4.3, where the ordered sequence of node expansions is also shown. Every time a forward iteration is performed ($i \rightarrow i - 1$) the algorithm selects and computes the distance increments of the nodes that fulfil the sphere constraint and continues the tree search with the most favorable node according to the Schnorr-Euchner enumeration [Schnorr91] (the node resulting in the smallest d_i), e.g. the expanded node in ②. At each level of the tree, those nodes that do not fulfil the sphere constraint, along with all their descendants, are discarded or pruned from the search tree. This process is repeated until a leaf node is reached (③, which will result in a radius update), or no nodes that satisfy the sphere constraint are found. In any case, the SE will proceed with a backward iteration ($i \rightarrow i + 1$) where a radius check will be performed among the previously computed set of candidate points. If a node with $D_i < R$ is found (④), the tree traversal is resumed with a forward iteration. The optimum solution has been found when the hypersphere with the updated radius contains no further nodes.

The radius reduction strategy, along with the tracking of potentially valid nodes at each level of the algorithm, prevents unnecessary distance computations, but ultimately results in a rather complex system architecture.

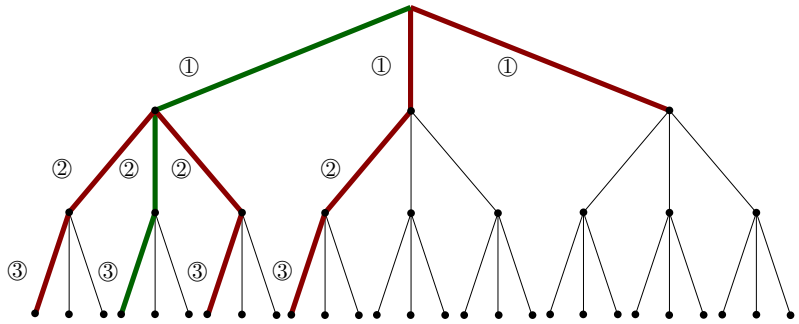

Figure 4.4: K-Best tree traversal example in a system with $N = 3$, $|\mathcal{L}| = 3$ and $K = 4$.

4.2.2 K-Best Tree Traversal for VP

Breadth-first tree-search algorithms traverse the tree without performing any backward iterations, in such a way that all the eligible nodes are expanded at every level. These algorithms benefit from a high data processing throughput that stems from the parallel processing of the branches during the tree traversal. However, the strict application of the breadth-first tree-search concept derives in a high computational load that results from the evaluation of all the perturbation vector hypotheses. Therefore, it is a common practice to bound the complexity of the breadth-first algorithms by restricting the amount of nodes considered at each level.

As one can guess from its name, the K-Best precoder selects the K best branches at each level of the tree regardless of the sphere constraint or any other distance control policy. The K-Best tree traversal strategy is illustrated in Figure 4.4 for a system with $N = 3$, $|\mathcal{L}| = 3$ and $K = 4$. The parallel branch-processing nature of the algorithm is reflected in the equal node expansion sequence of all the nodes at a certain level of the tree. At each stage i of the K-Best tree search, an ordering procedure has to be performed on the eligible $K|\mathcal{L}|$ candidate branches based on their AEDs down to level i. After the sorting procedure, the K paths with the minimum accumulated distances are passed on to the next level of the tree. Once the final stage of the tree has been reached, the branch with the minimum Euclidean distance is selected as the K-Best solution.

Even though only K branches are considered as potential solutions at the final stage of the algorithm, the amount of calculated distance increments that need to be computed is remarkably higher than KN and depends strongly on the sorting approach implemented at each stage. Sorting algorithms for K-Best-based VP systems are out of the scope of this section and will be analyzed in a forthcoming chapter. Additionally to the unnecessary distance computations, the full-sorting procedures performed at every level contribute to a considerable resource occupation and power consumption of a K-Best hardware implementation.

4.3 The Fixed-complexity Sphere Encoder

The proposed fixed-complexity sphere encoder (FSE) algorithm has been designed with the objective of providing an implementation-friendly tree-search algorithm for VP. This novel scheme aims at overcoming the two main drawbacks of the SE, namely the variable complexity and the sequential nature of the tree search, along with the major bottleneck of the fixed-complexity K-Best, that is, the computationally-intensive sorting stages. The proposed scheme is based on the fixed-complexity sphere decoder (FSD) presented in [Barbero06b] for the detection of single-user MIMO systems.

The main difference of the proposed FSE with respect to the optimum albeit sequential SE is that the search is not constrained only to those nodes whose PEDs are within a certain distance from a reference signal. To the contrary, the search is performed in an unconstrained fashion in terms of distance control. The tree-search structure is defined by a tree configuration vector $\boldsymbol{n} = [n_1, \ldots, n_N]$ instead, which specifies the number of child nodes to be evaluated at each level (n_i). Therefore, only n_i PEDs are computed per parent node at each level, yielding a total candidate branch count of $n_T = \prod_{i=1}^{N} n_i$. The selection of the expanded nodes is performed based on the Schnorr-Euchner enumeration [Schnorr91], hence avoiding the intricate sorting stages required by the K-Best tree search.

The tree is therefore traversed from level $i = N$ down to level $i = 1$, expanding only the most promising n_i nodes at each stage. When the bottom of the tree is reached, that is $i = 1$, the path with the smallest AED among the tree branches considered during the FSE tree traversal is selected as the solution. A sample representation of an FSE tree search is depicted in Figure 4.5 for a $N = 3$ user system with a constrained lattice of $|\mathcal{L}| = 3$ elements and $\boldsymbol{n} = [1, 2, 2]$. Note that, as is the case with the K-Best tree traversal, all the node expansions with the same sequence number can be computed simultaneously.

The proposed tree-search algorithm therefore features two main concepts: on one hand, the ordering strategy, which determines the order in which the users' data streams are visited during the tree traversal, and on the other hand, the tree configuration vector $\boldsymbol{n}$ which shapes the structure of the search tree. These two characteristics of the FSE algorithm will be analyzed in the following sections.

4.3.1 Matrix Preprocessing for the FSE Tree Search

The close-to-optimum performance of the original FSD algorithm proposed for the single-user MIMO detection scenario is due to a unique combination of an unconventional ordering strategy and an optimized tree-search structure. The key factor of the FSD ordering approach is that, in order to minimize the probability of a wrong decision, the weakest stream is placed at the top of the tree, where all possible nodes are considered for expansion. However, this

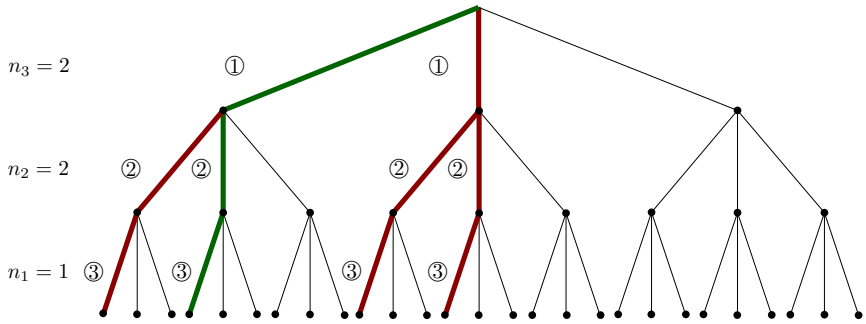

Figure 4.5: FSE tree traversal example in a system with $N = 3$, $|\mathcal{L}| = 3$ and tree configuration vector $\boldsymbol{n} = [1, 2, 2]$.

approach is not applicable in precoding scenarios as the set of candidate nodes at the top level of the tree does no longer have a finite number of elements. This fact, along with the inherent differences between decoding and precoding scenarios, motivates the implementation of a different ordering strategy for the FSE algorithm.

The introduction of a certain user ordering strategy $\mathfrak{O}$ into the VP system model derives in the following reformulated expressions for the matrix decomposition procedure in (2.7) and the optimum perturbation vector computation in (2.8):

$$\boldsymbol{U}^H\boldsymbol{U} = \boldsymbol{\Pi}_{\mathfrak{O}} \left(\boldsymbol{H}\boldsymbol{H}^H + \xi \boldsymbol{I}_N\right)^{-1} \boldsymbol{\Pi}_{\mathfrak{O}}^T$$

$$\boldsymbol{a} = \boldsymbol{\Pi}_{\mathfrak{O}}^T \underset{\hat{\boldsymbol{a}} \in \tau\mathbb{C}\mathbb{Z}^N}{\operatorname{argmin}} \left\| \boldsymbol{U}(\boldsymbol{\Pi}_{\mathfrak{O}}\boldsymbol{s} + \hat{\boldsymbol{a}}) \right\|_2^2. \tag{4.3}$$

The quality of the pruning process in a K-Best tree search can be enhanced by placing the most *reliable* streams at the top of the tree, so that a better solution can be found in the subsequent node selection stages. This ordering approach has been reported to achieve good error-rate performance results when applied to precoding scenarios. More specifically, the sorted QR matrix ordering ($\mathfrak{O}_{\mathrm{SQR}}$) presented in [Wübben01] has been used along with the K-Best fixed-complexity tree-search scheme in [Habendorf07].

A more suitable ordering strategy can be used for the FSE by exploiting the similarities between the VP and THP schemes. After all, the VP system model is transformed into the THP precoding scheme if an FSE tree traversal with $\boldsymbol{n} = [1, \ldots, 1]$ is used to obtain the precoding vector. Hence, the THP solution is always contained within the set of candidate branches considered by the FSE regardless of the tree configuration vector in use. In a K-Best system however, such an statement cannot be made as the THP solution vector may be dropped out during any of the sorting stages of the algorithm. The following section summarizes the main features of the preferred ordering for THP systems.

4.3.1.1 Ordering Strategy for THP (*Best-last* Rule)

THP can be regarded as the precoding counterpart of DFE, where a successive interference cancelation procedure is performed on the received signal vector. With the aim of minimizing the error propagation through each detection step, the spatial streams in a DFE approach are ordered following a *best-first* rule [Kusume07]. This ordering strategy states that the strongest data stream is detected first while the worst stream is left for the final detection stage.

Nevertheless, the picture is different in precoding scenarios. In a THP system for example, the first data stream is transmitted unaltered (only the channel shaping matrix $\boldsymbol{T}$ is applied to this signal), whereas the last precoded signal has to avoid interfering all the previous data streams. Hence, the task of precoding the last user is much more arduous. Consequently, following a *best-last* rule instead and precoding the best stream last is a sensitive decision, as this data stream has less degrees of freedom. This reversed ordering approach represents the optimum ordering strategy for THP systems ($\mathfrak{O}_{\mathrm{BL}}$) [Joham04b].

Once the strategy of the ordering procedure has been defined, we shall focus on the criteria used to select the best/worst stream at each iteration. Unlike in point-to-point MIMO detection, where the ordering was determined based on the SNR associated with each transmitted stream, the MSE of the perturbed symbols in a block transmission will be the parameter used to establish the optimum user permutation strategy for VP. Note that, the perturbation vector in (2.8) is computed so as to minimize the MSE for a given block of data symbols. Consequently, the best user stream will be the one that contributes the least to the MSE, or equivalently the one with a smaller diagonal value of the $\boldsymbol{\Psi} = (\boldsymbol{H}\boldsymbol{H}^H + \xi \boldsymbol{I}_N)^{-1}$ matrix.

Table 4.1 shows the pseudocode for the joint user order determination and triangular matrix computation. The reverse direction in which the matrix preprocessing (from $i = 1$ to N) and the tree traversal (from $i = N$ down to $i = 1$ since $\boldsymbol{U}$ is upper triangular) are carried out is noticeable from the algorithm displayed in this table. The calculation of the triangular matrix in (2.7) is performed following the computationally-efficient *Cholesky factorization with symmetric permutation* method presented in [Kusume05] [Kusume07], whose complexity is similar to that of a sorted QR decomposition.

4.3.1.2 Performance of the FSE with Different Ordering Strategies

The performance of the proposed FSE tree search is depicted in Figure 4.6 for a 4×4 antenna system and a tree configuration vector $\boldsymbol{n} = [1, 1, 2, 5]$. The triangular matrix used for the fixed-complexity tree traversal has been rearranged following the aforementioned $\mathfrak{O}_{\mathrm{SQR}}$ and $\mathfrak{O}_{\mathrm{BL}}$ ordering strategies. The BER performance curve of the unordered FSE tree search has additionally been included for completion. The benefits of performing an appropriate

$$
\begin{aligned}
&\boldsymbol{\Psi} = \left(\boldsymbol{H}\boldsymbol{H}^H + \xi \boldsymbol{I}_N\right)^{-1} \\
&\boldsymbol{\Pi} = \boldsymbol{I}_N \\
&\text{for } i = 1, \ldots, N \\
&\quad q = \underset{\hat{q} \in \{i, \ldots, N\}}{\operatorname{argmin}} \boldsymbol{\Psi}(\hat{q}, \hat{q}) \\
&\quad \boldsymbol{\Xi} = \tilde{\boldsymbol{I}}_N^{(i,q)} \\
&\quad \boldsymbol{\Pi} = \boldsymbol{\Xi}\boldsymbol{\Pi} \\
&\quad \boldsymbol{\Psi} = \boldsymbol{\Xi}\boldsymbol{\Psi}\boldsymbol{\Xi}^T \\
&\quad \boldsymbol{D}(i,i) = \boldsymbol{\Psi}(i,i) \\
&\quad \boldsymbol{\Psi}(i:N,i) = \boldsymbol{\Psi}(i:N,i)/\boldsymbol{\Psi}(i,i) \\
&\quad \boldsymbol{\Psi}(i+1:N,i+1:N) = \boldsymbol{\Psi}(i+1:N,i+1:N) \\
&\qquad -\boldsymbol{\Psi}(i+1:N,i)\boldsymbol{\Psi}(i+1:N,i)^H \boldsymbol{D}(i,i) \\
&\text{end} \\
&\boldsymbol{L} = \text{lower triangular part of } \boldsymbol{\Psi} \\
&\boldsymbol{U} = \boldsymbol{D}^{1/2}\boldsymbol{L}^H
\end{aligned}
$$

Table 4.1: Computation of the upper-triangular matrix $\boldsymbol{U}$ with *best-last* ($\mathfrak{O}_{\mathrm{BL}}$) ordering.

preprocessing stage on the triangular matrix are clearly visible from the data displayed in this figure. As one can notice, a performance gain of 5 dB at a BER of 10^{-5} can be attained by simply rearranging the order in which the users are processed within the search tree. The data displayed in this figure also show that a similar BER performance is obtained by the FSE at the low-SNR range with any of the considered ordering strategies. Nevertheless, the FSE with the $\mathfrak{O}_{\mathrm{SQR}}$ ordering suffers a considerable error-rate performance degradation in the high-SNR regime.

Considering the similar computational complexity and notable performance difference between the aforementioned ordering strategies, the $\mathfrak{O}_{\mathrm{BL}}$ ordering scheme is proposed for an FSE-based VP system.

4.3.2 FSE Tree Configuration Vector

The tree configuration vector of the FSE provides a flexible error-rate performance and complexity trade-off. Heuristically, one can argue that the higher the number of calculated paths is, the better the performance of the overall system will be, and viceversa. This fact is generally true, as an extended set of candidate vectors increases the probability of finding the optimum solution vector. Nevertheless, it should be noted that, for a given non-prime n_T value, several tree configuration vectors can be defined, where each one will provide a certain error-rate performance.

In the original description of the algorithm, namely the FSD technique for signal detection, the use of the $\boldsymbol{n} = [1, 1, \ldots, n_T]$ tree configuration vector was recommended in consideration of the specific features of the FSD ordering strategy. Nevertheless, due to the

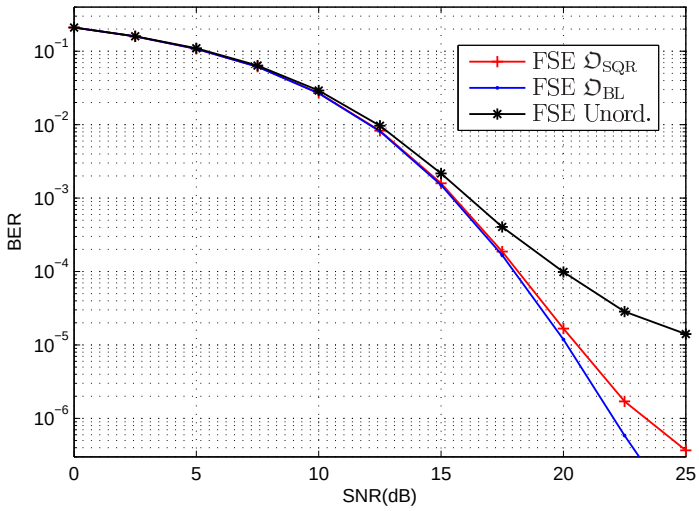

Figure 4.6: BER performance of the proposed FSE algorithm with $\boldsymbol{n} = [1, 1, 2, 5]$ and ordering strategies $\mathfrak{O}_{\mathrm{SQR}}$ and $\mathfrak{O}_{\mathrm{BL}}$. Additionally, the unordered case is included for completion.

key differences between MIMO detection and multi-user precoding scenarios outlined in Section 4.3.1, the guidelines for the design of the optimum tree configuration vector presented in [Barbero06b] are not valid for the current application of the sort-free fixed-complexity tree search.

However, there are several lessons to be learned from the tree-search structure of the FSD. For example, by distributing the values within the tree configuration vector such that $n_1 \leq n_2 \leq \ldots \leq n_N$, a broader range of eligible values can be achieved at the top of the tree, which enables the selection of more suitable perturbation values at the rest of the levels of the tree. Additionally, n_1 should be set to 1 as setting a higher value for this element would not provide any error-rate improvement if the Schnorr-Euchner enumeration is followed.

We shall refer to the set of W factoring prime integers of n_T as $\mathcal{W} \triangleq \{\eta_1, \eta_2, ..., \eta_W\}$ with $\prod_{j=1}^{W} \eta_j = n_T$. For a given value of n_T, the set of eligible tree configuration vectors can be represented as

$$\begin{aligned}\mathcal{N}_{n_T} = \{&[1, \ldots, 1, n_T] \\ &[1, \ldots, \eta_1, n_T/\eta_1] \\ &[1, \ldots, \eta_2, n_T/\eta_2] \\ &\quad\vdots \\ &[1, \ldots, \eta_1, \eta_2, n_T/(\eta_1\eta_2)] \\ &\quad\vdots \\ &[1, \ldots, \eta_1, \eta_2, \ldots, \eta_W]\}\,.\end{aligned}$$

Figure 4.7 shows the BER performance curves of FSE-based VP systems in the 6×6 and 8×8 antenna setups considering the tree configuration vector candidate sets $\mathcal{N}_{12}$ and $\mathcal{N}_{24}$, respectively. The data depicted in these figures show that the tree configuration vectors of the type $\boldsymbol{n}=[1,\ldots,n_T]$ achieve the worst performance among the eligible vectors in the set $\mathcal{N}_{n_T}$ for both antenna setups. Moreover, those tree structures with more dispersedly distributed values of n_i achieve the best performance, being the performance gap with respect to other tree configuration vectors with higher values of n_i more noticeable in the 8 antenna case.

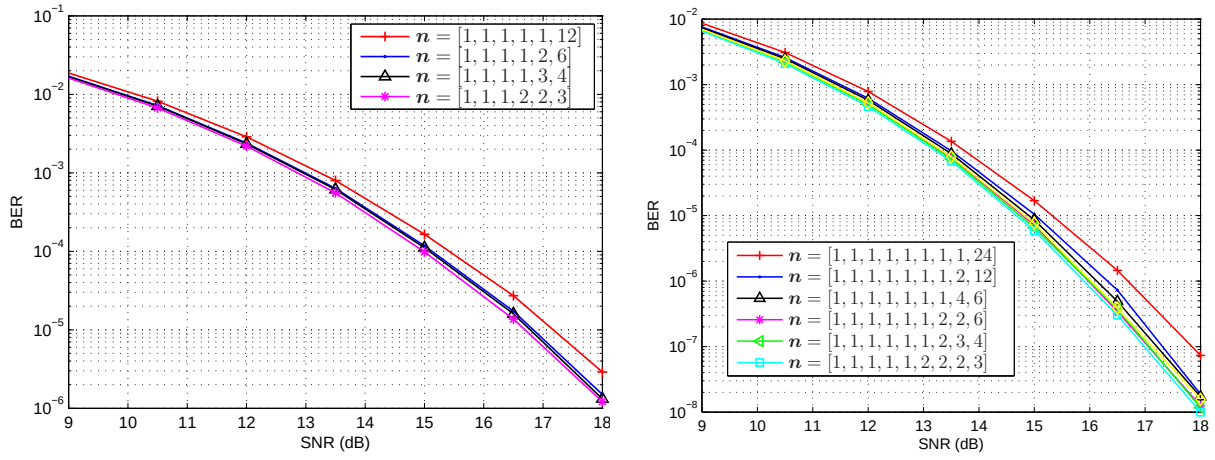

(a) (b)

Figure 4.7: BER performance of the FSE with different tree configuration vectors in a 6×6 (a) and an 8×8 (b) antenna setup.

Apart from the impact on the error-rate performance of the FSE, the implementation of a certain tree configuration vector also affects the computational complexity of the tree traversal. As already stated, all the candidate tree configuration vectors in $\mathcal{N}_{n_T}$ yield the same amount of total computed branches or AEDs. Nevertheless, this does not imply that the amount of computed PEDs is equal for all the considered configuration vectors. Therefore, the selection of the tree configuration vector will determine the computational complexity of the FSE, as most of its computational load is due to PED calculations. The amount of required PED calculations (C_{PED}) for a given FSE tree structure is given by the following formula:

$$C_{\mathrm{PED}}=\sum_{i=1}^{N}\prod_{j=i}^{N}n_i. \tag{4.4}$$

Additionally to the Euclidean distance computations, the number of intermediate point calculations should also be taken into account when assessing the computational complexity of a certain tree-search structure, as the high volume of complex-valued multiplications involved in their computation represents an important part of the computational load of the FSE tree traversal. Therefore, the amount of intermediate points to be computed (C_{IP}) can

be calculated as:

$$C_{\mathrm{IP}} = \sum_{i=1}^{N-1} \prod_{j=i+1}^{N} n_i. \tag{4.5}$$

Clearly from equations (4.4) and (4.5), those tree configuration vectors with reduced values of n_i at each level will render the smallest computational complexity among the candidate vectors in $\mathcal{N}_{n_T}$. To better support this argument, the number of PED computations and z_i calculations has been computed for all the tree configuration vectors with $n_T = 24$ in an $N = 8$ user system. Additionally, the total number of operations has been calculated for each one of this cases, where for the sake of simplicity, all the arithmetic operations (multiplication, addition and substraction) have been considered to have the same weight on the final operation count. The results of this study are depicted in Figure 4.8. As already anticipated, the provided data reflects the dependency between the computational complexity of the FSE tree traversal and its tree-search structure. More specifically, those tree-search architectures with a greater amount of expanded nodes at the top levels require a higher number of operations than those with a more distributed node expansion scheme.

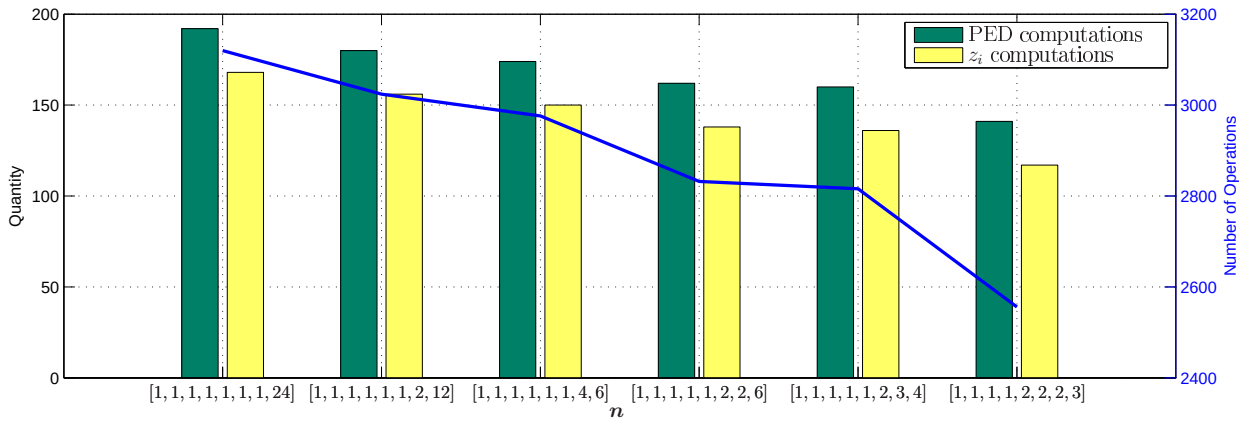

Figure 4.8: Number of PED computations (C_{PED}) and z_i calculations (C_{IP}) required for several tree configuration vectors with $n_T = 24$. Additionally, the total amount of operations is depicted for each one of the FSE tree-search structures.

However, when it comes to hardware implementation, the multiplication count of a certain design is a factor of mayor importance as the embedded multipliers are a scarce resource in field-programmable gate array (FPGA) devices. In this case too, the structure of the tree search will determine the amount of required multiplication units, as is shown in Figure 4.9. The results shown in the aforementioned figure have been obtained by considering that 3 multipliers are required for the product of two complex variables. This has been achieved by rearranging the terms involved in the computation of the complex multiplication from the most straightforward approach structured as $(a+jb)(c+jd) = (ac-bd)+j(ad+cb)$, where 4

multiplication operations are performed, to a more computationally-efficient method, namely $(a+bj)(c+dj) = [a(c-d)+d(a-b)]+j[b(c+d)+d(a-b)]$, that reduces the number of required multipliers due to the repeated factor $d(a-b)$ [Barbero06a]. The data shown in Figure 4.9 reflect the diverse computational complexity of the tree configuration vectors in the set $\mathcal{N}_{24}$. The tree configuration vectors that expand more nodes at the higher levels of the tree search require more intermediate point computations and PED calculations, as already seen in Figure 4.8, which ultimately results in a high volume of allocated multipliers. Hence, by selecting the most disperse tree configuration vector within the candidate set $\mathcal{N}_{24}$, the amount of required embedded multipliers can be reduced in a 42% when compared to the originally proposed $\boldsymbol{n} = [1, \ldots, 1, n_T]$ tree structure.

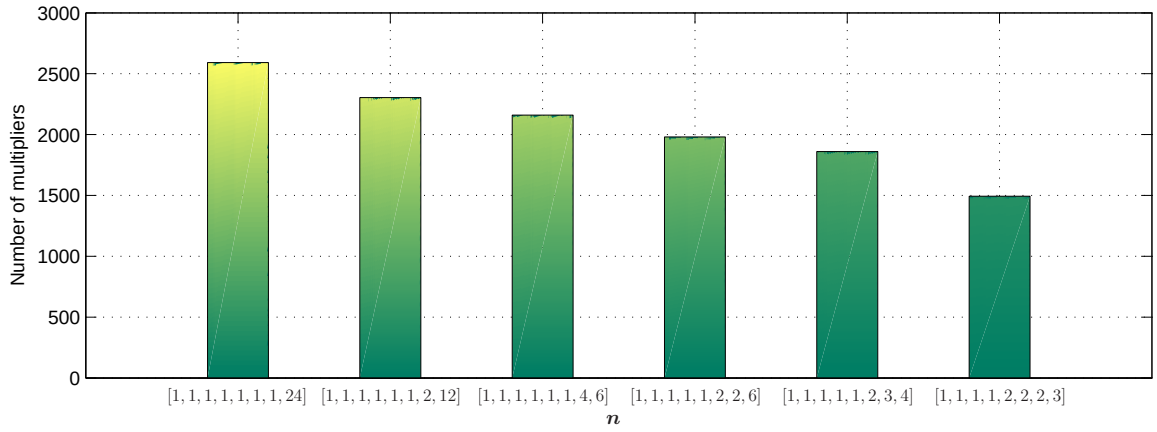

Figure 4.9: Amount of required multipliers for different tree configuration vectors in an FSE tree search.

In the light of the error-rate performance and computational complexity results, it is concluded that the optimum tree configuration vector for the precoding scenario is obtained by factorizing the amount of desired computed branches n_T and arranging the resulting values such that $n_1 \leq n_2 \leq \ldots \leq n_N$. Hence, if the tree traversal is performed starting from $i = N$ and working backwards until $i = 1$, the tree configuration vector is set as

$$n_i = \begin{cases} \max \quad \mathcal{W}^{(i)} & \text{for } N - i < W \\ 1 & \text{otherwise} \end{cases},$$

where $\mathcal{W}^{(i)}$ represents the set of factoring prime integers that have not yet been selected at level i.

During the development of this research work, an alternative precoding version of the FSD has been published in [Mohaisen11]. The presented algorithm, also referred to as FSE, deals with the parallel computation of the branches in a fixed-complexity tree search. Nevertheless, the differences with respect to the algorithm presented in this chapter are

numerous. First, no guidelines for the configuration of the fixed-complexity tree search are provided. Additionally, the approach in [Mohaisen11] features a real-valued decomposition (RVD) model, with the consequent double length of the tree search, and does not consider the use of a matrix preprocessing method to enhance the error-rate performance of the system. Due to these facts, the approach in [Mohaisen11] achieves a considerably worse BER performance than the FSE scheme proposed in this dissertation. More specifically, an error-rate performance degradation of 2.5 dB at a BER of 10^{-4} with respect to the K-Best approach in [Zhang05] is reported, whereas the scheme here proposed achieves a better performance than the aforementioned K-Best precoder when a reasonably similar number of candidate branches are considered, as it will be shown in forthcoming sections.

4.4 Simulation Results

The complexity and error-rate performance of the reviewed and presented tree-search techniques will be assessed in this section. Note that, the provided results depict the performance of fully-loaded VP systems, namely $M = N$, with 16-QAM modulation. Furthermore, a constrained search set of $|\mathcal{L}| = 25$ elements has been considered for both the error-rate simulations and the complexity analysis.

4.4.1 Computational Complexity

Two main factors contribute to the computational complexity of a certain tree search algorithm: on one hand, the amount of node expansions performed during the tree traversal, and on the other hand, the number of operations required to determine which nodes are to be considered at each level.

4.4.1.1 Number of Evaluated Nodes

The amount of nodes to be visited during the tree traversal is given by design parameters in fixed-complexity tree-search techniques, that is, K and $\boldsymbol{n}$ for the K-Best and FSE algorithms, respectively. The choice of these design parameters for the current analysis has been performed so as to yield a close-to-optimum performance of the fixed-complexity schemes [see Figure 4.13]. Nevertheless, for the recursive SE algorithm, this complexity measure is not only subjected to design features such as the user ordering strategy, but also to environmental factors, i.e. the SNR.

The SNR-dependant regularization factor in (2.7) causes the diagonal values of the triangular matrix $\boldsymbol{U}$ to be more sparse at high SNRs. As a consequence to this, the difference between the PEDs of the nodes closer to the root and the initial sphere constraint results in a very lenient pruning and consequently, in a high volume of expanded nodes during

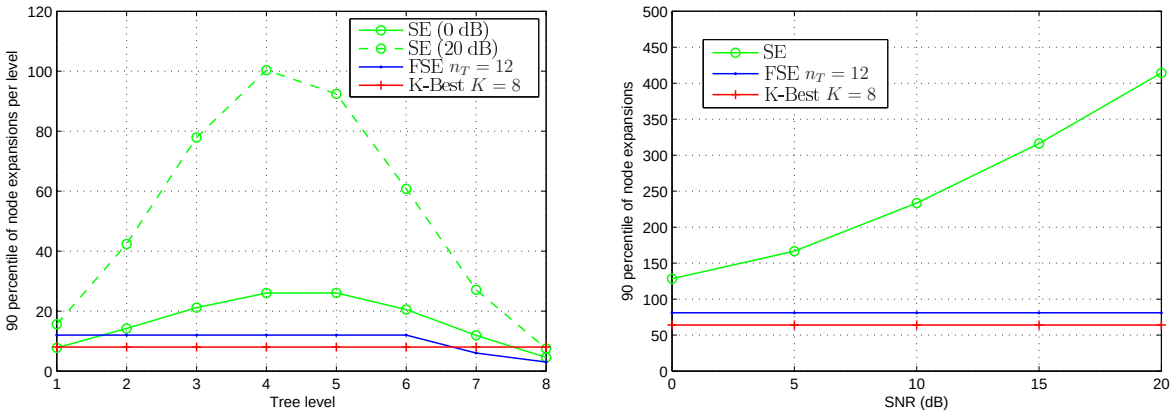

(a) (b)

Figure 4.10: 90-percentile of the number of node expansions for the SE, FSE and K-Best tree-search algorithms in an 8×8 antenna system. The amount of expanded nodes per level is given in (a), whereas the total amount of evaluated nodes vs. SNR is considered in (b).

the SE tree traversal. This fact is represented in Figure 4.10(a), where the 90-percentile of the number of evaluated nodes per level is shown for different tree-search techniques. As is shown in this figure, the run-time of the SE algorithm is considerably longer in the high-SNR range (20 dB), specially in the middle part of the search tree where the advantages of the node-pruning policy are yet unnoticeable.

Additionally, the different level-wise node distribution patterns for the iterative and non-iterative schemes can be observed in the aforementioned figure. The SE undergoes an escalating node expansion phase in the upper half of the tree due to the high amount of nodes that fulfil the sphere constraint at the early stages of the tree traversal. Nevertheless, there is a decrease in the evaluated nodes per level in the subsequent stages due to the implemented pruning strategy. On the contrary, due to of the lack of sphere constraint, the fixed-complexity techniques yield a constant node distribution pattern, with the exception of the ramp-up phase in the initial stages of the FSE tree search where the amount of computed partial branches $n_T^{(i)} = \prod_{j=i}^{N} n_j$ is smaller that n_T.

On a more general perspective, the dependency between the total amount of visited nodes and the SNR is depicted in Figure 4.10(b). As one can notice, the complexity of the SE in terms of volume of expanded nodes is noticeably higher than that of the FSE and K-Best approaches, specially in the high-SNR regime.

4.4.1.2 Number of Operations

Apart from the number of evaluated nodes per level, the amount of operations to be performed in order to select the set of nodes that will be visited during the tree traversal is also an important complexity measure that needs to be taken into account. Hence, the node se-

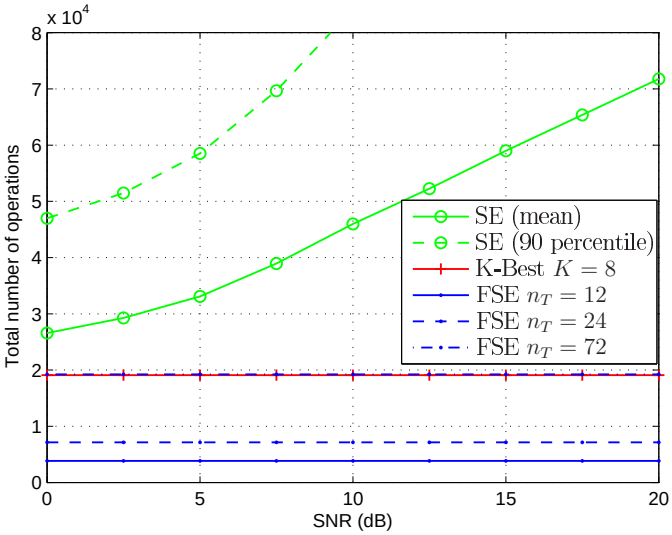

Figure 4.11: Total number of operations for the SE, K-Best and FSE tree-search approaches in an $N = 8$ user system.

lection procedure performed by each tree-search strategy will dictate the amount of excessive PED computations and the computational load derived from possible sorting procedures.

For example, if a SE tree traversal is implemented, every time a forward iteration is performed, the PED of the node to be expanded is computed along with the PEDs of the nodes that fulfil the sphere constraint in that level. These nodes and their associated PEDs compose the candidate list of a certain tree level. Hence, if the sphere search is finished after the first N node expansions, that is, only a single branch is computed, the amount of PED calculations will be significantly higher than N. For the case of the K-Best on the other hand, the selection procedure requires that the PEDs of all possible $K|\mathcal{L}|$ nodes are computed and subsequently sorted in order to select the K nodes to be expanded at a certain level. Therefore, the additional distance computations and the sorting stages should also be taken into account when assessing the computational complexity of this algorithm. When traversing an FSE tree, the nodes to be expanded can be selected based on their distance to the intermediate point when $n_i > 1$, or by a simple rounding operation when $n_i = 1$.

With the aim of assessing the overall computational complexity of the different tree-search approaches, the number of arithmetic (addition, substraction, multiplication, division) and logical (comparison, swapping, branching) operations has been measured. This way, not only the computational complexity of the PED calculations is reflected in the final operation count, but also the complexity due to the sorting and node-selection stages. All the aforementioned operations are considered to have the same weight in the final operation count. The computational complexity analysis that is shown in Figure 4.11 has been performed on an $N = 8$ tree search with a constrained grid of $|\mathcal{L}| = 25$ elements. The data displayed in this figure show the high computational complexity of the SE approach when compared to

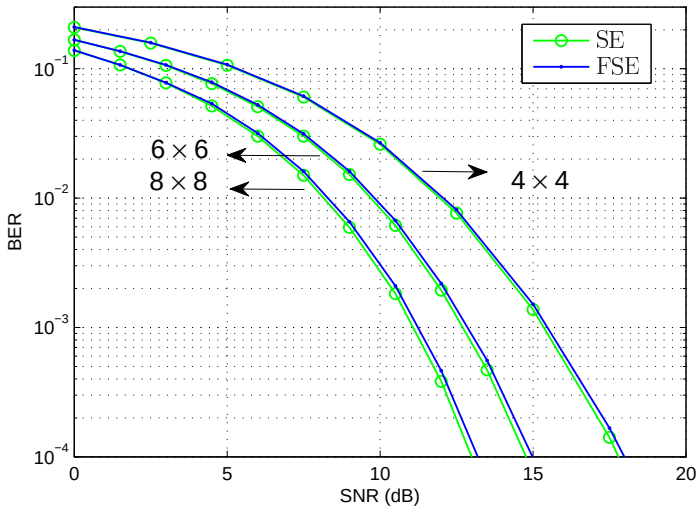

Figure 4.12: BER performance curves of the FSE applied to VP systems in 4×4, 6×6 and 8×8 antenna setups. The total amount of evaluated branches has been set to 10, 12 and 24 for the aforementioned antenna configurations, respectively.

the fixed-complexity schemes. This result is due to two main factors: on one hand, the SNR-dependant node expansion process discussed in Section 4.4.1.1, and on the other hand, the high computational cost derived from maintaining the candidate lists required to implement the sphere constraint.

As for the fixed-complexity approaches, the K-Best performs a considerably higher number of operations when compared to the FSE, even though the amount of expanded nodes for the former is more reduced, as already seen in Figure 4.10(b). The complexity gap between the K-Best scheme and the FSE approach is such that it requires an FSE tree search with $n_T = 72 \gg K = 8$ to yield a similar operation count. This is mainly due to the excessive distance computations required by the selection algorithm and the sorting stages performed at each level in the K-best tree-search model.

The amount of computed distances and computational complexity of the selection stages required for all the tree-search techniques can be reduced by using an appropriate node enumeration technique. This issue will be the main topic of the upcoming Chapter 5.

4.4.2 BER Performance

First of all, we shall focus on the BER performance of the proposed fixed-complexity scheme. Figure 4.12 depicts the BER performance curves for the FSE applied to fully-loaded VP systems of dimension $N = 4$, $N = 6$ and $N = 8$, where the total amount of branches to be evaluated has been set to $n_T = 10$, $n_T = 12$ and $n_T = 24$, respectively. The performance gap between the FSE and the optimum SE is negligible in all the antenna configurations under study. More specifically, a small performance loss of 0.2 dB can be achieved by the FSE at

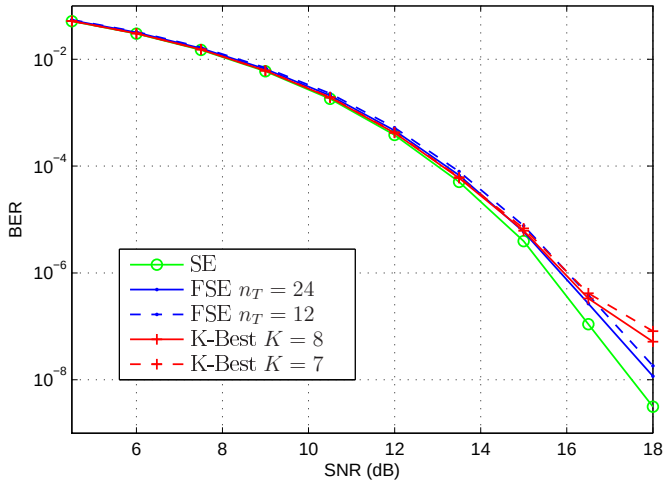

Figure 4.13: BER performance of the SE, the FSE with $n_T = 12$ and $n_T = 24$ and the K-Best with $K = 7$ and $K = 8$ in an 8×8 VP system.

an uncoded BER of 10^{-4}.

It is noteworthy that the FSE achieves a close-to-optimum performance by evaluating a considerably smaller amount of branches than the original algorithm designed for single-user MIMO detection. As documented in [Barbero06a], the computation of 16 ($\boldsymbol{n} = [1, \ldots, 1, P]$) and 256 ($\boldsymbol{n} = [1, \ldots, P, P]$) branches is required for a quasi-ML performance of the FSD algorithm with 16-QAM modulation ($P = 16$) in 4×4 and 8×8 antenna setups, respectively. The data in Figure 4.12 show that the proposed fixed-complexity tree search can also achieve a close-to-optimum performance under the same system parameter considerations by means of the evaluation of just 10 and 24 branches, respectively. What is more, the modulation-dependant tree-search structure of the FSD algorithm contributes to the poor scalability of the model with respect to the size of the modulation in use (4096 branch computations are required for an 8 level tree search with 64-QAM modulation), whereas the tree-search structure of the FSE remains unaltered in the event of a change of modulation order.

In order to showcase the main performance differences between the FSE and K-Best fixed-complexity approaches, a broader error-rate analysis will be performed, where lower BER values will be targeted. As is shown in Figure 4.13, the BER performance of the FSE and K-Best fixed-complexity approaches for a fully-loaded $N = 8$ user system is close to the optimum set by the SE in the low and mid-SNR range. Nevertheless, the slope of the BER vs SNR curve of the K-Best schemes degrades severely for high SNR values. Moreover, note that the performance of the FSE with $n_T = 12$ is considerably better than that of the K-Best with $K = 8$ in the high-SNR regime, even if the complexity of the FSE in terms of total operation count has been shown to be remarkably smaller than that of the K-Best [see Figure 4.11].

4.5 Chapter Summary

This chapter addresses the problem of an efficient computation of the perturbation vector in VP systems. Several tree-search techniques have been proposed in the literature to solve the selection of the perturbation vector, being the SE one of the most widely used tree-search techniques due to its optimum error-rate performance. The good performance of the algorithm is a consequence of the identification and management of the admissible set of nodes at each stage of the tree search, which ultimately leads to a variable complexity of the algorithm and a rather intricate hardware architecture of the tree traversal.

Due to the simplicity of its architecture and the possibility of parallel processing, the fixed-complexity K-Best algorithm has been regarded as a prominent candidate for an efficient and high-speed hardware implementation of vector precoders. Despite its non-iterative nature, which greatly simplifies the tree traversal, the sorting stages required to select the candidate branches contribute to the high computational complexity of the algorithm. With the aim of overcoming the main shortcomings of the K-Best precoder, a sort-free fixed-complexity algorithm has been proposed in this chapter. The presented FSE scheme traverses the tree in a non-iterative fashion, where the selection of the nodes at each level is dictated by a tree configuration vector. This design parameter offers a flexible trade-off between error-rate performance and complexity of the algorithm. Furthermore, the optimum combination of matrix preprocessing strategy and structure of the FSE tree search that results in a minimum computational complexity and best error-rate performance of the FSE structure has been introduced.

The complexity in terms of evaluated nodes and total operation count has been assessed for the aforementioned tree-search schemes. The SNR dependency of the iteration complexity of the optimum SE, along with the continuous radius checks performed during the tree traversal, contribute to the high computational load of this recursive scheme. The provided results on the amount of evaluated nodes and total operation count show a reduced and constant complexity of the non-iterative approaches. Moreover, the lack of sorting stages in the FSE tree-search architecture derives in a considerable complexity reduction when compared to the K-Best or SE approaches.

On the whole, the proposed FSE scheme has been shown to achieve a close-to-optimum error-rate performance with a remarkably small and constant complexity. Hence, the performance degradation suffered by the proposed technique can be neglected considering the great computational savings that are achieved when compared to the optimum SE algorithm.

Chapter 5

Complex-plane Enumeration for Precoding Systems

5.1 Introduction

Many of the algorithms proposed in the literature to solve the closest-point search problem for VP systems resort to a distributed distance computation model. The resulting equivalent system can be subsequently solved by traversing the tree of candidate nodes following a variety of strategies, as already discussed in the previous chapter.

The identification of the set of most favorable nodes plays a vital role in the way the tree is traversed regardless of the tree-search approach to be followed. The aforementioned task can be performed by means of the Schnorr-Euchner enumeration strategy [Schnorr91], which states the order in which a certain set of child nodes is to be visited.

The incorporation of an adequate enumeration approach will affect the features of the VP system differently depending on the nature of the tree traversal. For example, in a VP system based on the SE algorithm, a more efficient radius reduction, and therefore a shorter run-time of the algorithm, can be achieved by expanding the more probable nodes first rather than visiting them based on a pre-defined order as established by the Fincke-Pohst enumeration approach [Fincke85]. On the contrary, bounded breadth-first algorithms benefit from a constant processing time, and therefore, the advantages of implementing a suitable enumeration strategy are reflected on other features of the tree search. For instance, the incorporation of an Schnorr-Euchner enumerator can greatly simplify the architecture of a K-Best tree search as it enables the implementation of distributed sorting strategies, such as the ones described in [Wenk06] and [Mondal08], to overcome the computationally expensive sorting stages performed at every level. These techniques take advantage of the ordered sequence of nodes given by the Schnorr-Euchner enumeration to greatly simplify the task of selecting the K most promising branches. On a different note, the alternative sort-free structure of the FSE algorithm is enabled by the incorporation of a Schnorr-Euchner enumerator which selects the nodes to be expanded at every level.

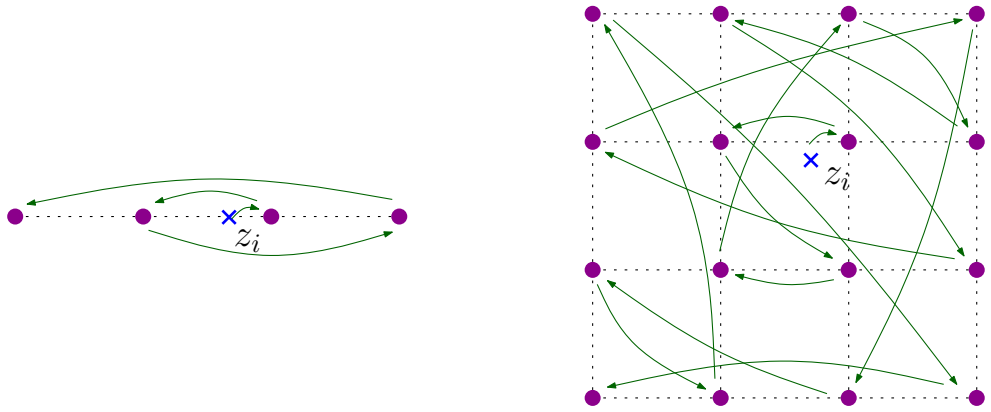

Figure 5.1: Schnorr-Euchner ordered sequence of nodes in a real-valued model (a) and in the complex plane (b).

The intricacy of the complex-plane enumeration has led to the selection of real-valued equivalent models as the preferred approach when designing the hardware architecture of tree-search algorithms in the literature. Nevertheless, the simplicity of the enumerator comes at the cost of an expanded tree, whose depth is twice that of the original. This derives in a higher resource occupation and longer delays, which may also affect the final throughput of the system. Furthermore, the benefits of working directly on the complex signals in VP systems has been recently reported in [Müller08], whereas the advantages of performing the tree search in the complex-plane were outlined and studied in [Burg05] and [Barbero06a].

This chapter deals with the efficient implementation of complex-plane enumeration algorithms applied to precoding systems. As an starting point, a background study of the most prominent Schnorr-Euchner enumeration algorithms will be performed, followed by the presentation of a novel non-sequential complex-plane enumerator. Additionally, and with the aim of showcasing the simple architecture of the presented enumerator, the fully-pipelined high-throughput implementation of the proposed algorithm and other state-of-the-art complex-plane enumerators will be carried out in the following sections.

5.2 Schnorr-Euchner Enumeration

The Schnorr-Euchner enumeration principle states the order in which a certain set of child nodes is to be visited according to their distance to the intermediate point z_i (4.2). This results in the nodes being expanded in ascending order of their associated PED.

When a real equivalent model is used, the ordering of the child nodes is a simple task as the nodes just need to be visited in a zig-zag fashion, as is shown in Figure 5.1(a). However, Figure 5.1(b) depicts that such an straightforward scheme cannot be followed if complex-valued signals are utilized. The provided illustrative example clearly shows the non-triviality

of the enumeration process.

The exhaustive computation of the PEDs of all children nodes and their subsequent sorting is a simple approach to determine the ordered sequence of most favorable nodes. Nonetheless, this procedure derives in a great amount of hardware resources dedicated to computing the distance increments of the $|\mathcal{L}|$ child nodes. What is more, the sorting stage entails a great data movement, which ultimately results in an increased power consumption. With the aim of avoiding the high complexity and resource demand of full-sorting enumeration, several complex-plane enumeration algorithms have been proposed in the literature. In the following subsections, a brief overview of the main complex-plane enumeration strategies will be performed.

5.2.1 Enumeration by Identification of the Admissible Set

A novel approach for complex-plane enumeration based on trigonometric calculations was presented in [Hochwald03]. The proposed scheme was originally developed to determine the admissible interval of constellation points taking into account the sphere constraint set by the sphere decoder (R). The proposed approach can easily be adapted to precoding scenarios by arranging the lattice symbols in concentric circles whose intersection with the search disk centered at $z_i = \hat{r}e^{i\hat{\theta}}$ is identified as the admissible set. More specifically, if the elements of the lattice $\mathcal{L}$ are represented as $r_p e^{i\theta_p}$, where r_p represents the radius of each one of the concentric circles, the admissible interval of θ_p for each circular set is given by the equation

$$\left\lceil \frac{|\mathcal{L}|}{2\pi}\left(\hat{\theta} - \cos^{-1}\left[\frac{r_p^2 + \hat{r}^2 - \bar{R}^2}{2r_p\hat{r}}\right]\right)\right\rceil \leq \frac{|\mathcal{L}|}{2\pi}\theta_p \leq \left\lfloor \frac{|\mathcal{L}|}{2\pi}\left(\hat{\theta} + \cos^{-1}\left[\frac{r_p^2 + \hat{r}^2 - \bar{R}^2}{2r_p\hat{r}}\right]\right)\right\rfloor, \quad (5.1)$$

with $\bar{R} = R/u_{i,i}$. Once the boundaries have been established, the symbols are visited in a zig-zag fashion within each concentric circle. An illustration of the admissible set of points within $\mathcal{L}$ is depicted in Figure 5.2 for arbitrary values of z_i and $\bar{R}$.

One of the main issues with this enumeration approach is its applicability to bounded breadth-first systems where the search is performed in an unconstrained way, i.e. the PEDs of the eligible nodes do not comply with any sphere constraint. Additionally, the complex trigonometrical computations in (5.1) render it unsuitable for an efficient hardware implementation.

5.2.2 Enumeration by Node Arrangement in Circular Subsets

The intricacies of performing the highly complex trigonometrical calculations required by (5.1) motivated the work in [Burg05], where an implementation-friendly version of the ideas in [Hochwald03] was presented. Figure 5.3 illustrates the procedure for complex enumeration

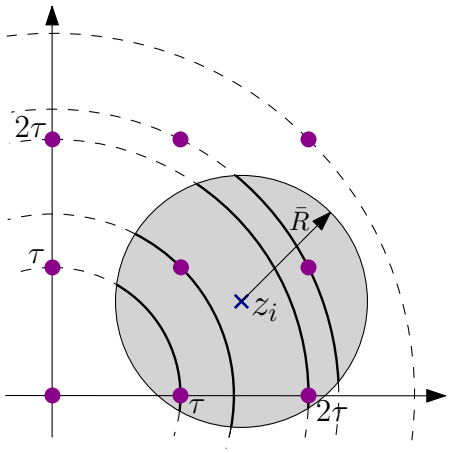

Figure 5.2: Identification of the admissible set of candidates by means of the sphere constraint.

for a restricted lattice of $|\mathcal{L}| = 25$ elements. Note that, since the lattice of precoding symbols also includes null components, the parameters of the algorithm in [Burg05] have been slightly modified to adjust to the VP case. Therefore, in the system under study, the symbols are arranged in $P_{\mathcal{S}} = 6$ subsets were the 0 element is the unique member of one of the subsets.

As an starting point, the value of z_i is mapped into the first quadrant to reduce the amount of decision boundaries to be checked. Consequently, all the enumerated values will have to be mapped back to their original quadrants once the enumeration process is finished. Next, the initial symbols within each concentric circle are chosen. According to the original description of the algorithm in [Hochwald03], the preferred child node within a certain circular set is the one with the minimum phase difference with respect to the intermediate point, namely the node that yields a smaller $|\hat{\theta} - \theta_p|$ value. Nevertheless, given the structure of the lattice, the selection of the most favorable node within each set can be easily performed by checking the position of z_i with respect to the bisector line spanned between the real and imaginary axis ($x = y$ boundary line) when needed, as is shown in Figure 5.3.

Once all the initial symbols have been identified (nodes with a double circle in Figure 5.3), their PEDs are computed and subsequently compared, which results in the selection of the node associated with the smallest distance increment as the starting point for the enumeration ($a_i^{(1)}$). The process is continued by selecting the following symbol that minimizes the phase difference with respect to z_i within the subset of $a_i^{(1)}$. This local enumeration procedure is performed in a zig-zag fashion, where the direction of the enumeration is determined by an additional set of boundary conditions ($3x = y$, $3y = x$ and $x = y$). Next, the PED of the newly enumerated symbol is computed and compared to the previously calculated $P_{\mathcal{S}} - 1$ distance increments, which leads to the selection of the second enumerated value ($a_i^{(2)}$). The algorithm proceeds accordingly until no more symbols need to be evaluated or all the subsets are empty.

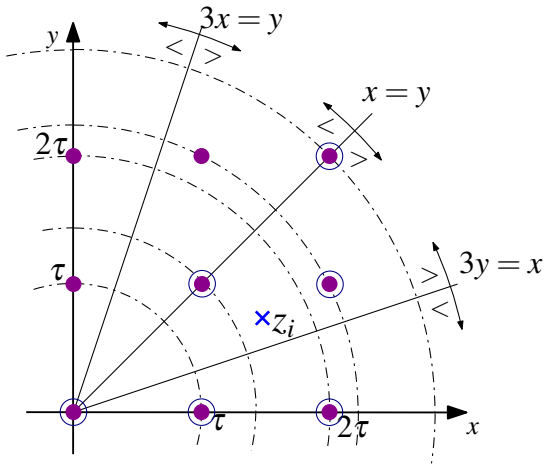

Figure 5.3: Boundaries and node arrangement for the Schnorr-Euchner enumeration by circular subsets.

Note that $P_{\mathcal{S}}$ incremental distance calculations need to be performed at the initial stage of the algorithm and just one PED computation for each one of the following enumerated points. This implies that $P_{\mathcal{S}} - 1$ redundant distance calculations need to be performed in order to identify the set of ρ preferred nodes. Furthermore, the node selection process requires ρ minimum-search units (MSUs), being each one of them carried out by $P_{\mathcal{S}} - 1$ compare-and-select blocks.

If a sphere constraint is implemented, the amount of subsets to be considered can be reduced by adapting the approach presented in [Chen07] to precoding scenarios. The algorithm in [Chen07] is based on identifying the group of *active* circular subsets based on the noise variance and certain signal statistics. After all, for a given search radius $\bar{R}$, only certain circle sets have eligible candidates in them, as is shown in Figure 5.2, and therefore, the number of PED computations can be greatly reduced by considering only these for the enumeration procedure.

5.2.3 Enumeration by Node Arrangement in Unidimensional Subsets

A similar approach to the one in [Burg05] is proposed in [Hess07] , where the arrangement of nodes in different subsets is performed based on their real (or equivalently imaginary) part, as is shown in Figure 5.4. This allows for a smaller amount of subsets ($P_{\mathcal{S}} = 5$), specially for high values of $|\mathcal{L}|$. More specifically, the difference in the number of subsets scales with the number of elements in the lattice as $\frac{1}{8}(|\mathcal{L}| - 4\sqrt{|\mathcal{L}|} + 3) \geq 0$. Additionally, the local enumeration is performed based on the imaginary (or equivalently real) part of the lattice elements, and hence, the same zig-zag pattern can be shared among all subsets.

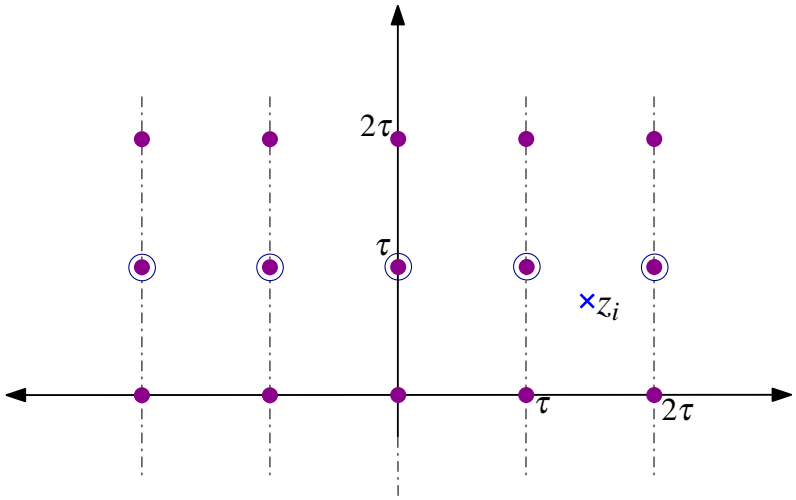

Figure 5.4: Node arrangement for the Schnorr-Euchner enumeration by unidimensional subsets.

5.2.4 Enumeration by Neighbor Expansion

This novel approach [Shabany08b] uses the simple zig-zag enumeration in the real and imaginary components of the lattice symbols as a mean to determine the order for the complex Schnorr-Euchner enumeration. Therefore, the first step of the algorithm is to independently order the elements in the real and imaginary axis according to their proximity to z_i. This process is depicted in Figure 5.5, where the elements $\{a^1_{\Re}, a^2_{\Re}, a^3_{\Re}\}$ and $\{a^1_{\Im}, a^2_{\Im}, a^3_{\Im}\}$ represent the ordered sequence of nodes in the real and imaginary axis, respectively. Clearly, the initial value in the Schnorr-Euchner sequence $(a_i^{(1)})$ will be a combination of the first enumerated values in the real and imaginary axes, namely $a_i^{(1)} = a^1_{\Re} + a^1_{\Im} j$. To determine the rest of the values, a candidate list $\mathcal{C}$ needs to be defined.

The process of node selection and candidate identification is depicted in Figures 5.5(a) and 5.5(b) for the first two nodes in the ordered sequence. After the initial enumeration point has been selected, the two adjacent symbols according to the unidimensional enumeration are added to the list of candidates. Therefore, the symbols $a^2_{\Re}+a^1_{\Im} j$ and $a^1_{\Re}+a^2_{\Im} j$ are included in $\mathcal{C}$ and their corresponding PEDs are computed. The one with the smallest PED is selected as $a_i^{(2)}$ and its two adjacent symbols are added to the candidate list, namely $a^3_{\Re} + a^1_{\Im} j$ and $a^2_{\Re} + a^2_{\Im} j$ for the provided example. Note that a single new symbol will be included in $\mathcal{C}$ if one of the adjacent symbols has previously been selected. This derives in a variable length of the candidate list, whose number of elements cannot be determined beforehand.

5.2.5 Other Suboptimum Enumeration Techniques

The enumeration techniques presented so far deal with obtaining the ordered sequence of child nodes as dictated by the Schnorr-Euchner enumeration. Nevertheless, it is possible to

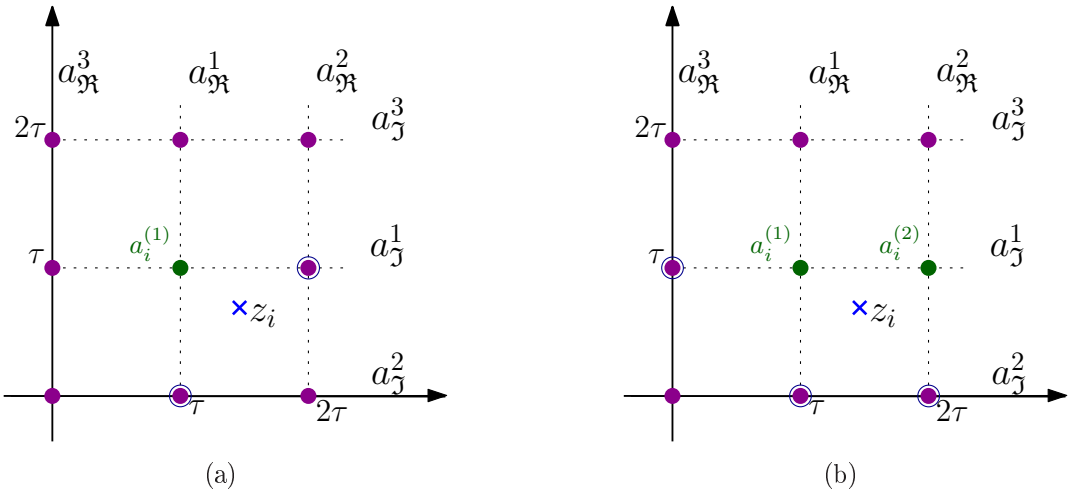

Figure 5.5: Schnorr-Euchner enumeration by neighbor expansion.

obtain a certain complexity reduction by relaxing the requirement of perfect enumeration. The provided hardware resource saving and added performance degradation will strongly depend on the tree-search type that these enumeration techniques are implemented upon. In this subsection, several suboptimum enumeration approaches (with respect to the accuracy of the resulting ordered set) are reviewed.

In [Mennenga09] an approach for search sequence determination based on geometrical considerations was presented. The proposed scheme works by dividing the area around the first enumerated value and storing a predefined search sequence for each region. The selection of the enumerated values is then performed by determining the region where the value of z_i lies and retrieving the corresponding search sequence from look-up table (LUT)s. The main drawback of the proposed scheme is its poor scalability, as both the length of the sequence to be stored and the number of differentiated regions grow with the amount of enumerated values. For instance, a total of 71 bounded regions need to be considered to obtain the search sequence of the preferred 16 nodes. Due to the asymmetrical nature of the bounded areas, the selection of the region corresponding to a given value of z_i is a highly complicated matter.

The ideas in [Mennenga09] inspired the work in [Wenk10b], where the incorporation of the approximated norm $l^{\tilde{\infty}}$ provides the system with the scalability and simple architecture that was missing in the original work. In the updated structure, only 8 regions need to be defined regardless of the amount of child nodes to be enumerated. Due to the non-accumulative feature of the approximated norm in use, many of the nodes that share the same real or imaginary component yield an equal $l^{\tilde{\infty}}$ distance with respect to the intermediate point. This represents the main advantage of the proposed enumeration technique, as all the $l^{\tilde{\infty}}$-equidistant nodes form unidimensional subsets that can be easily sequenced following a zig-zag approach.

5.3 The Puzzle Enumerator

Most of the enumeration techniques that have been published so far in the literature follow a sequential scheme, that is, the closest node to z_i is obtained first, then the second closest point is computed and so on. The dependency on the past enumerated values in order to select a certain point within the ordered sequence comes as a logical way of proceeding, but incurs in a great latency of the enumeration algorithm and a significant throughput reduction. Moreover, in high-throughput fully-pipelined architectures, long delays incur in an extended device occupation due to the high amount of required pipeline registers.

With the aim of overcoming the issues of sequential enumeration, a novel complex-plane enumeration technique is presented in this chapter. The proposed puzzle enumerator allows for the independent computation of the nodes within the ordered sequence. Since the information on the previously selected $\rho - 1$ symbols is not required in order to obtain the ρ^{th} most favorable node, the latency of the presented enumeration unit is not increased as the enumeration process progresses. Moreover, the selection of the symbols is performed without the computation of their associated distances, which results in a low computational complexity of the enumeration process.

5.3.1 Initial Considerations

When implementing a non-linear VP system the designer has the flexibility to select the number of lattice elements to be considered during the tree search. If a restricted grid with a minimum of $|\mathcal{L}| = 25$ elements is used, the BER performance of the VP system is barely degraded in the SNR range of interest ($<$20 dB), as has already been shown in Section 4.2. In order to allow for a more simplified structure of the proposed puzzle enumerator, the restricted search set $\mathcal{L}$ used in the previous chapters of this thesis will be extended to yield the more favorable set $\mathcal{L}_{\text{ext}}$. In this section, a simple procedure to determine the minimum cardinality of the extended search set will be presented.

As opposed to the detection problem, where all the constellation symbols are equally probable, those lattice points that are closer to the origin are more likely to be part of the solution vector in VP systems. As a consequence to this, the values of z_i are concentrated around the origin of the lattice forming a circle of radius R_z that expands with the SNR. Figure 5.6 depicts an illustrative example of the intermediate points generated during a SE tree search for different SNR values. As one can notice, the intermediate points are more densely concentrated around the origin of the lattice in the low-SNR range as the probability of having a null perturbation vector is higher in this scenario [Christensen07]. However, as the SNR increases, the values of z_i are more sparsely distributed but still follow a circular pattern.

From the point-of-view of hardware implementation, the value of the R_z radius in the working SNR range is a parameter of major interest as it conditions the quantization accuracy. Therefore, its value will have to be estimated beforehand by means of extensive simulations. Additionally, it can be used to set the minimum dimension of the constrained lattice that will result in an error-free puzzle enumeration.

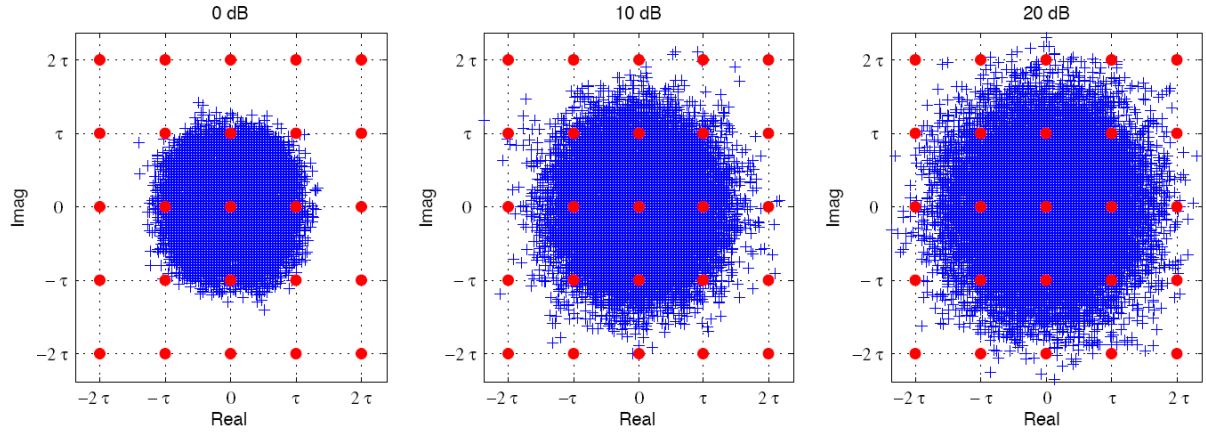

Figure 5.6: Distribution of the intermediate points for several SNR values.

The size of the search set to be used along with the proposed enumerator has to be chosen carefully so that the unidimensional enumeration in the real and imaginary components of the lattice can be performed following a strict zig-zag pattern. If $|\mathcal{L}_{\text{ext}}|$ is too small, there will be $z_i = z_{\Re} + z_{\Im} j$ values lying in the vicinity of the lattice border, and therefore, the local enumeration in the real and/or imaginary components will not result in an ordered zig-zag sequence. This derives in undesired border effects which require extra hardware resources and boundary checks to be solved. Hence, the minimum amount of points of the extended set $\mathcal{L}_{\text{ext}}$ will depend on both R_z and the length of the sequence to be determined (ρ), which gives

$$|\mathcal{L}_{\text{ext}}| = \begin{cases} \left[2\left(\lceil R_z/\tau \rfloor + 1\right) + 1\right]^2 & \text{for } \rho \in [2, 7] \\ \left[2\left(\lceil R_z/\tau \rfloor + 2\right) + 1\right]^2 & \text{for } \rho \in [8, 19] \\ \left[2\left(\lceil R_z/\tau \rfloor + 3\right) + 1\right]^2 & \text{for } \rho \in [20, 37] \\ \vdots & \vdots \end{cases}.$$

5.3.2 Low-complexity Enumeration Procedure

Once the lattice parameters have been set, the proposed low-complexity algorithm proceeds by performing a simple unidimensional enumeration in the real and imaginary axis. This local enumeration procedure is necessary as the elements in the resulting Schnorr-Euchner ordered set are given as a combination of the sequenced real and imaginary symbols, namely $a_i^{(\rho)} = a_{\Re}^p + a_{\Im}^q j$. The procedure of unidimensional enumeration has already been described as part of the node sequencing approach in [Shabany08b][see Figure 5.5]. Nevertheless, even

if a larger search set is considered in the proposed strategy, only a few values need to be sequenced locally as opposed to the full axis sequencing performed in the approach presented in [Shabany08b]. This way, 2 values need to be enumerated for $\rho \in [2,3]$, 3 for $\rho \in [4,7]$, 4 for $\rho \in [8,11]$, and so on.

The initial value of the enumeration process, namely $a_i^{(1)}$, is selected as $a_{\Re}^1 + a_{\Im}^1 j$, as already proposed in [Shabany08b]. In order to the determine the rest of the sequence, a set of simple and pre-defined boundary checks needs to be performed. These border lines are built by pairing-up the possible candidate $a_i^{(\rho)}$ points and spanning the bisector line between them. The group of candidate points is created by selecting only those lattice elements that are feasible values for the current enumerated point. This procedure is performed by discarding the elements that have already been selected as an early solution in the enumerated sequence with all certainty. For example, the child node $a_{\Re}^1 + a_{\Im}^2 j$ will not be considered for the enumeration of the seventh element in the sequence as it must have been selected as the second or third most promising node regardless of the particular value of z_i. A further candidate discarding procedure can be carried out among the remaining elements in $\mathcal{L}_{\text{ext}}$ by considering the nodes $a_{\Re}^x + a_{\Im}^y j$ as unviable solutions if the element $a_{\Re}^z + a_{\Im}^y j$ with $x > z$ is present in the candidate list.

This way, a series of regions are delimited in the area $\mathcal{T} \triangleq \{0 \leq \Delta_{\Re} < \tau/2, 0 \leq \Delta_{\Im} < \tau/2\}$, with $\Delta_{\Re} = |z_{\Re} - a_{\Re}^1|$ and $\Delta_{\Im} = |z_{\Im} - a_{\Im}^1|$, each of which is linked to a certain candidate node. Hence, the selection of the ρ^{th} node in the sequence is performed by identifying the region corresponding to the value of z_i within the puzzle defined in $\mathcal{T}$ for that specific position in the ordered sequence. Note that since the absolute value is taken when computing the unidimensional distance increment in $\Delta_{\Re}$ and $\Delta_{\Im}$, there is no need to map z_i into the first quadrant or the lower triangular part of the first quadrant as performed in [Burg05] and [Mennenga09], respectively.

Figure 5.7 shows the decision regions and the boundary lines (B_x) for the first thirteen enumerated values. Note also that several boundary lines are common for various ρ values, and hence, the computed boundary expressions can be heavily reused during the enumeration process. The mathematical expressions for the border lines shown in Figure 5.7 are depicted in Table 5.1. As one can notice, the simplicity of the boundary expressions is a beneficial fact that will greatly favor the implementation of this enumeration technique.

5.3.3 Unordered Puzzle Enumerator

For the case of an FSE precoder, any of the aforementioned enumerators can be utilized to select the n_i nodes that will be expanded at each level of the tree traversal. However, since no ordering procedure is performed on the selected nodes, it is sufficient to merely identify the set of the n_i closest symbols. This way, a further simplification on the proposed puzzle enumerator can be performed by combining the regions of the puzzles from $\rho = 2$ up to

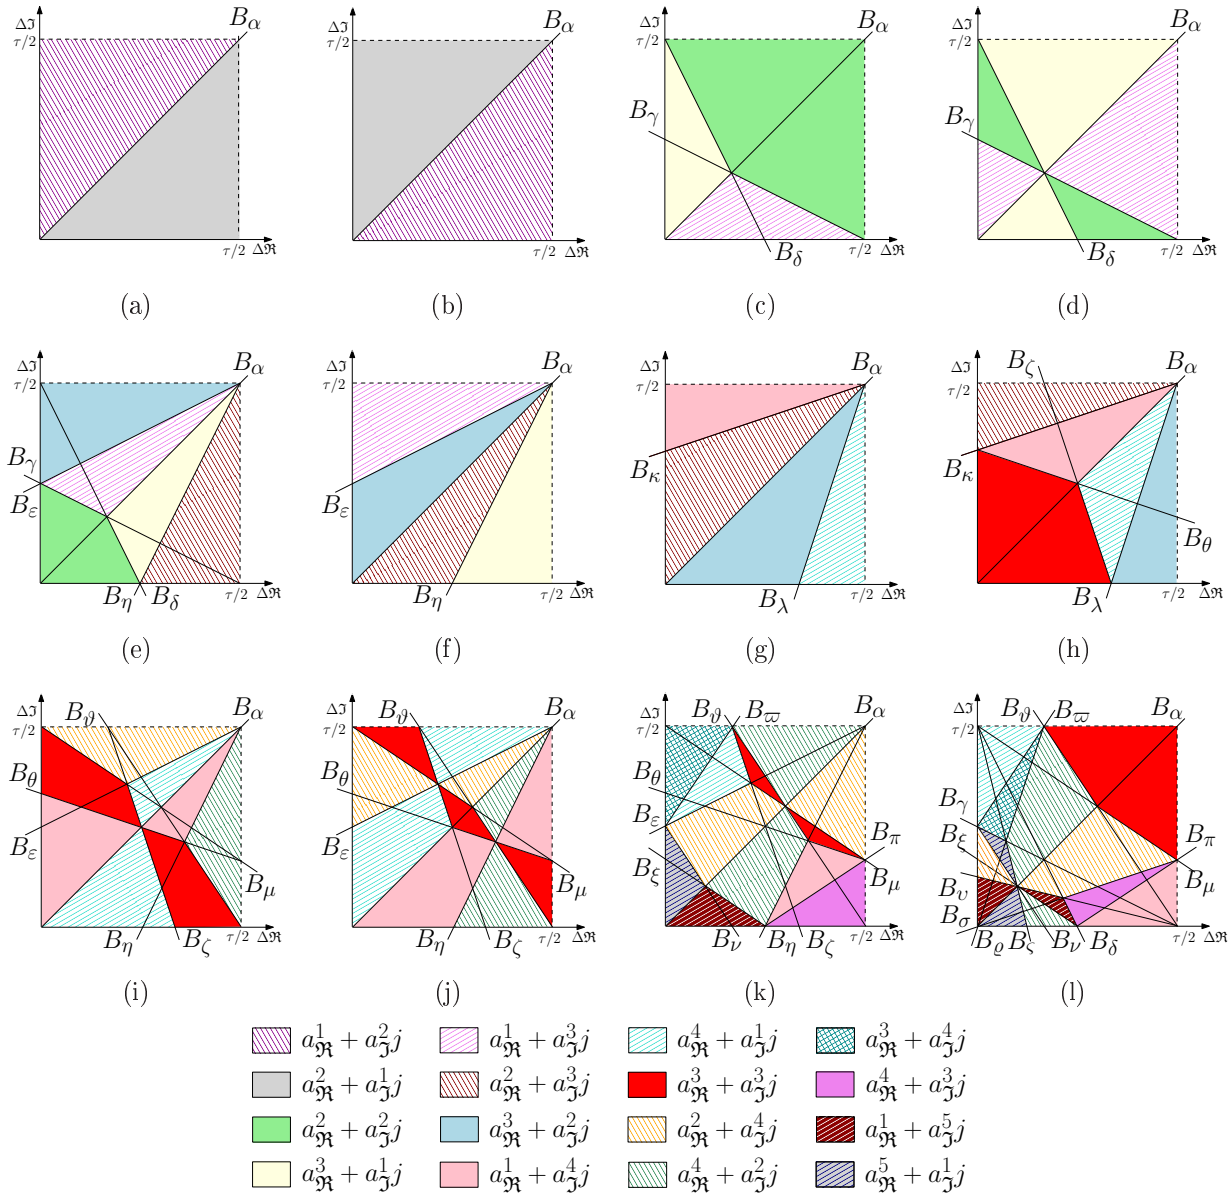

Figure 5.7: Fundaments of the proposed puzzle enumerator for (a) $\rho = 2$, (b) $\rho = 3$,(c) $\rho = 4$,(d) $\rho = 5$,(e) $\rho = 6$,(f) $\rho = 7$, (g) $\rho = 8$, (h) $\rho = 9$, (i) $\rho = 10$, (j) $\rho = 11$, (k) $\rho = 12$ and (l) $\rho = 13$.

$\rho \leq 3$	$B_\alpha : \Delta\Re = \Delta\Im$
$\rho \leq 5$	$B_\delta : 2\ \Delta\Re + \Delta\Im = \tau/2$ $B_\gamma : 2\ \Delta\Im + \Delta\Re = \tau/2$
$\rho \leq 7$	$B_\varepsilon : 2\ \Delta\Im - \Delta\Re = \tau/2$ $B_\eta : 2\ \Delta\Re - \Delta\Im = \tau/2$
$\rho \leq 8$	$B_\kappa : 3\ \Delta\Im - \Delta\Re = \tau$ $B_\lambda : 3\ \Delta\Re - \Delta\Im = \tau$
$\rho \leq 9$	$B_\zeta : 3\ \Delta\Re + \Delta\Im = \tau$ $B_\theta : 3\ \Delta\Im + \Delta\Re = \tau$
$\rho \leq 11$	$B_\vartheta : 3\ \Delta\Re + 2\ \Delta\Im = 3\tau/2$ $B_\mu : 3\ \Delta\Im + 2\ \Delta\Re = 3\tau/2$
$\rho \leq 12$	$B_\nu : 3\ \Delta\Re + 2\ \Delta\Im = \tau/2$ $B_\xi : 3\ \Delta\Im + 2\ \Delta\Re = \tau/2$ $B_\pi : 2\ \Delta\Re - 3\ \Delta\Im = \tau/2$ $B_\varpi : 2\ \Delta\Im - 3\ \Delta\Re = \tau/2$
$\rho \leq 13$	$B_\varrho : 3\ \Delta\Re = \Delta\Im$ $B_\sigma : 3\ \Delta\Im = \Delta\Re$ $B_\varsigma : 4\ \Delta\Re + \Delta\Im = \tau/2$ $B_\upsilon : 4\ \Delta\Im + \Delta\Re = \tau/2$

Table 5.1: Boundary lines to be evaluated for the first 13 nodes in the Schnorr-Euchner sequence.

$\rho = n_i$. If the union of the $n_i - 1$ regions that correspond to a particular symbol spans over the whole $\mathcal{T}$ area, that symbol is unequivocally part of the set of best n_i nodes. Otherwise, a boundary check procedure will have to be carried out to determine whether that certain symbol belongs to the set of preferable nodes.

The set of *unordered* enumerated values is shown in Table 5.2 for several values of n_i. The candidate nodes yet to be determined, e.g. C_i^x, can be obtained by checking the boundary conditions within a single puzzle. The bounded regions for C_i^x with $i \in \{3, \cdots, 11\}$ are depicted in Figure 5.8.

5.4 Implementation of Complex-plane Enumerators

Several complex-plane algorithms for the precoding scenario have been reviewed and proposed in the previous section. Despite the fact that all of the studied approaches yield the Schnorr-Euchner ordered sequence of nodes, there are great differences between the enumeration approaches in terms of computational effort and resource occupation of the hardware architecture. With the aim of assessing the amount of hardware resources dedicated to the task of enumeration in a high-throughput system, the implementation of the most relevant enumeration algorithms has been performed following a rapid-prototyping methodology. Specifically, fully-pipelined architectures of the full-sort enumerator, the ap-

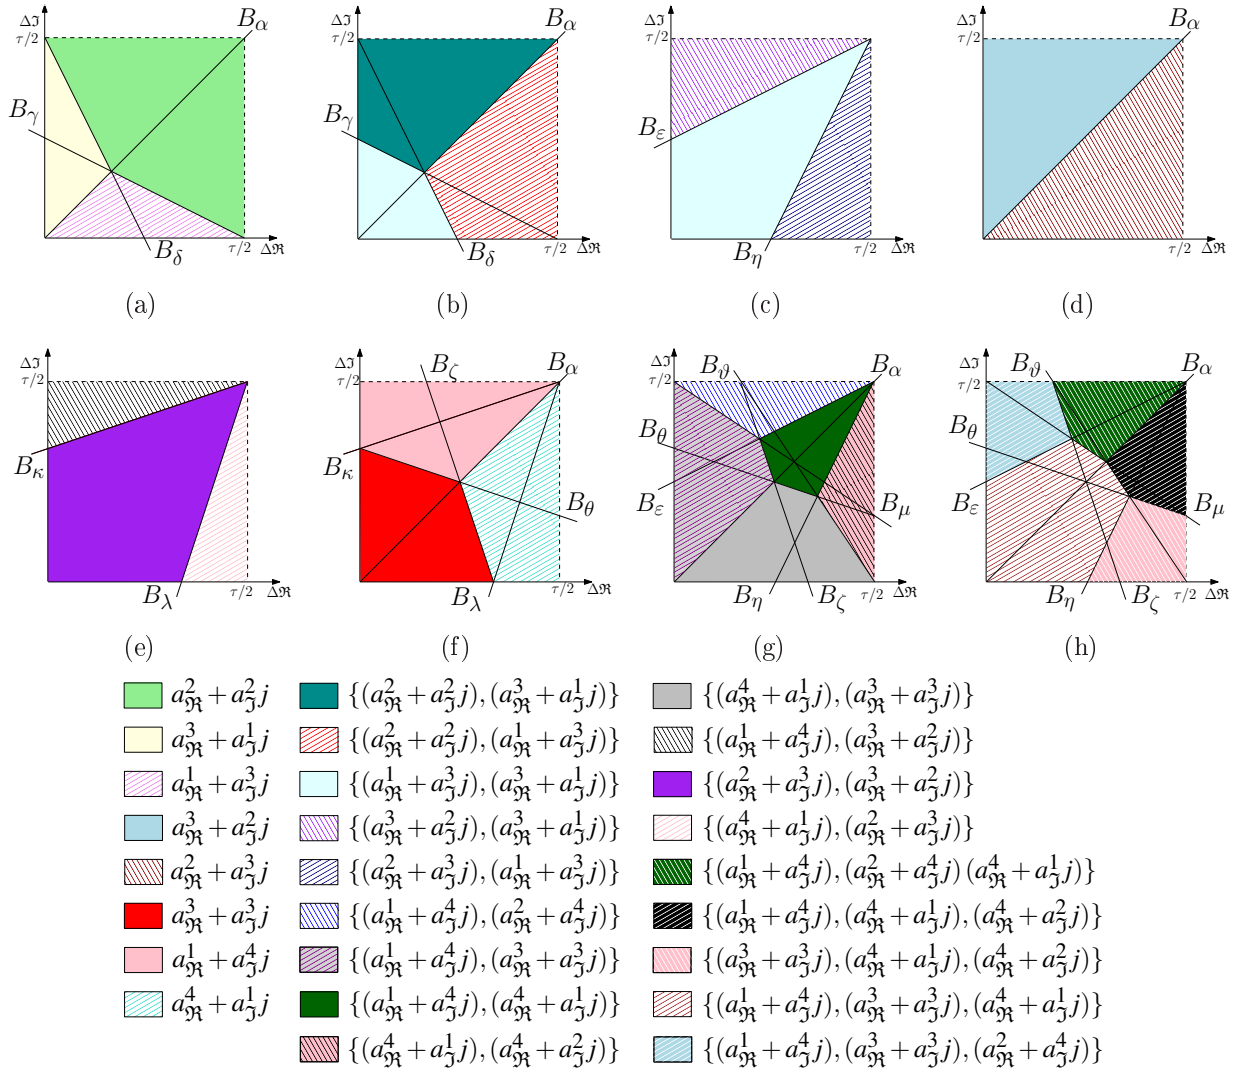

Figure 5.8: Fundaments of the proposed unordered puzzle enumerator to be used along with the FSE for (a) $n_i = 4$, (b)$n_i = 5$,(c) $n_i = 6$,(d) $n_i = 7$,(e) $n_i = 8$,(f) $n_i = 9$, (g) $n_i = 10$ and (h) $n_i = 11$.

n_i	Set of most favorable nodes (in no particular order)
3	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j)\}$
4	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), C_4^1\}$
5	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), C_5^1, C_5^2\}$
6	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), C_6^1, C_6^2\}$
7	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), (a_{\Re}^1 + a_{\Im}^3 j), (a_{\Re}^3 + a_{\Im}^1 j), C_7^1\}$
8	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), (a_{\Re}^1 + a_{\Im}^3 j), (a_{\Re}^3 + a_{\Im}^1 j), C_8^1, C_8^2\}$
9	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), (a_{\Re}^1 + a_{\Im}^3 j), (a_{\Re}^3 + a_{\Im}^1 j), (a_{\Re}^3 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^3 j), C_9^1\}$
10	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), (a_{\Re}^1 + a_{\Im}^3 j), (a_{\Re}^3 + a_{\Im}^1 j), (a_{\Re}^3 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^3 j), C_{10}^1, C_{10}^2\}$
11	$\{(a_{\Re}^1 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^1 j), (a_{\Re}^2 + a_{\Im}^2 j), (a_{\Re}^1 + a_{\Im}^3 j), (a_{\Re}^3 + a_{\Im}^1 j), (a_{\Re}^3 + a_{\Im}^2 j), (a_{\Re}^2 + a_{\Im}^3 j), C_{11}^1, C_{11}^2, C_{11}^3\}$

Table 5.2: Unordered sequence of preferred children for different values of n_i.

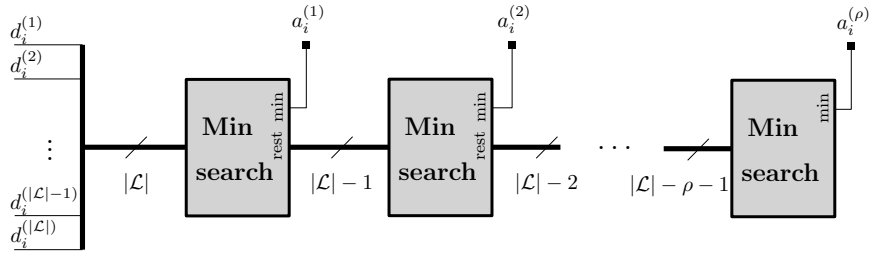

Figure 5.9: Block diagram of the full-sort enumerator.

proaches in [Burg05], [Hess07], [Shabany08b] and the proposed puzzle enumerator have been implemented for the sequencing of the 8 most favorable nodes.

As has been previously stated, all the enumerators with the exception of the proposed puzzle enumerator, base the selection of a certain enumerated symbol on the PEDs of the symbols in a candidate set. Even if the purpose of an enumerator is solely to identify a certain set of points in the vicinity of z_i, some of the calculated PEDs can be later reused by the tree-search algorithm. Therefore, for the sake of a fair comparison between enumeration units, $\rho = 8$ metric computation units (MCUs) have been added to the output of the puzzle enumerator. This way, the occupation and delay results shown in this section reflect the implementation of enumeration units that output both the selected symbols and their associated distance increments.

The block diagram of the fully-pipelined full-sort enumerator is shown in Figure 5.9. The architecture is mainly composed of MSUs that select the node associated with the minimum PED. The set of distance increments is reduced by a unit at every stage of the algorithm, being it composed of the PEDs of all the nodes in $\mathcal{L}$ at the initial stage, namely $\{d_i^{(1)}, d_i^{(2)}, \ldots, d_i^{(|\mathcal{L}|)}\}$.

The rest of the enumeration approaches, with the exception of the puzzle enumerator, perform the sequencing of the nodes by selecting a few candidate nodes from $\mathcal{L}$ and searching for the desired enumerated value within this set ($\mathcal{S}$). Every time a node is selected, it

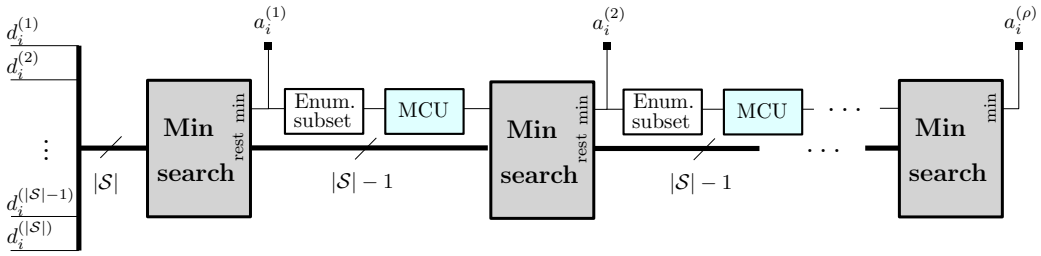

Figure 5.10: Block diagram of the enumeration strategies where the sequencing of the nodes is based on the management of a candidate set $\mathcal{S}$.

is extracted from the candidate set $\mathcal{S}$ and replaced by one or two other candidate nodes depending on the enumeration approach in use. This process is depicted in the block diagram in Figure 5.10. As is shown in this figure, the selection of a certain enumerated value $a_i^{(\rho)}$ is followed by the identification of the next node to enter the set of candidate nodes and the computation of its corresponding PED. Note that, for the approaches in [Burg05] and [Hess07], the number of elements in the candidate subset $\mathcal{S}$ is constant, whereas a variable length candidate list is used for the enumeration strategy in [Shabany08b].

Note also that it is a common practice to scale the lattice elements and the modulation constellation symbols so that a smaller quantization error can be achieved. Thus, after the application of a $2/\tau$ normalization factor, the normalized constrained search set will be of the form $\tilde{\mathcal{L}} = \tilde{\tau}\mathbb{C}\mathbb{Z}$, where $\tilde{\tau} = 2$ represents the modulo constant in this equivalent system. Following from this, the real and imaginary components of the constrained lattice of $|\tilde{\mathcal{L}}| = 25$ symbols will be extracted from the more convenient $\tilde{\chi} = [-4, -2, 0, 2, 4]$ set, without having any effect whatsoever on the final performance of the tree-search algorithm. The application of the normalization factor also affects other values involved in the tree traversal. In the remainder of this section, the distance increments, intermediate points and other parameters that have been computed using the normalized lattice will be identified with the mark ($\tilde{\cdot}$).

5.4.1 Efficient Distance Computation in State-of-the-art Enumerators

Depending on the architecture of the enumerator, different strategies can be implemented in order to reduce the number of costly operations when computing the distance increments. For the full-sort and the precoding version of the algorithm in [Burg05], the resource sharing distance computation procedure described in the latter will be used. This way, the PED computation in Equation (4.1) will be reformulated as a function of the arc ν and the real and imaginary components of the lattice symbols, which gives:

$$\tilde{d}_i(\nu, \tilde{\chi}_p, \tilde{\chi}_q) = u_{i,i}^2 \left[|\tilde{z}_i|^2 + R_\nu^2 + 2(\tilde{z}_{\Re}\tilde{\chi}_p + \tilde{z}_{\Im}\tilde{\chi}_q) \right], \tag{5.2}$$

where R_ν represents the radius of arc ν and $\tilde{\chi}_p$ stands for the p^{th} element of the set $\tilde{\chi}$. Following the approach in (5.2), the amount of required multipliers at the MCUs can be reduced when compared to the independent distance computation procedure, as $|\tilde{z}_i|^2$ needs to be computed only once and the values for the different R_ν^2 can be stored as constants. Moreover, due to the characteristics of the elements in $\tilde{\chi}$, the result for the multiplications of the form $\tilde{z}_{\mathfrak{R}}\tilde{\chi}_x$ can be performed by means of inexpensive bit shifting and negators.

The resource sharing approach in (5.2) can be slightly reformulated to better fit the requirements of the MCUs for the enumerators in [Hess07] and [Shabany08b], where the nodes are arranged in columns and rows instead of in concentric circles. The distance increments are therefore given by the elements in the real and imaginary axis, namely:

$$\tilde{d}_i(\tilde{\chi}_p, \tilde{\chi}_q) = u_{i,i}^2 \left[(\tilde{z}_{\mathfrak{R}}^2 + \tilde{\chi}_p^2 + 2\tilde{\chi}_p\tilde{z}_{\mathfrak{R}}) + (\tilde{z}_{\mathfrak{I}}^2 + \tilde{\chi}_q^2 + 2\tilde{\chi}_q\tilde{z}_{\mathfrak{I}}) \right]. \tag{5.3}$$

In this case too, the values of $\tilde{z}_{\mathfrak{R}}^2$ and $\tilde{z}_{\mathfrak{I}}^2$ are only computed once per enumerated sequence, while the squared values of the elements in $\tilde{\chi}$ are stored as constants.

5.4.2 Implementation of the Puzzle Enumerator

As has already been discussed, the initial step in the setup of the puzzle enumerator is the determination of the cardinality of the set $\tilde{\mathcal{L}}_{\text{ext}}$. Guided by the simulation results on the z_i values provided in Figure 5.6, a value of $\tilde{R}_z = 5$ has been considered sufficient for the SNR range of interest (<20 dB). This results in a minimum size of the extended lattice of $|\tilde{\mathcal{L}}_{\text{ext}}| = 81$, where the extended real and imaginary axis are composed of elements from the set $\tilde{\chi}_{\text{ext}} = \{-8, -6, -4, \ldots, 6, 8\}$.

The ability of the proposed algorithm to select the most favorable nodes in an independent way reflects positively on the simplicity of its implementation. As already stated in the description of the puzzle enumerator, the evaluation of the boundary lines summarized in Table 5.1 represents the only computational effort of the proposed node sequencing strategy. These boundary conditions, which are structured as $A\ \tilde{\Delta}_{\mathfrak{R}} \pm B\ \tilde{\Delta}_{\mathfrak{I}} = C\tilde{\tau}$, can be easily computed by means of adders and constant multiplication modules. What is more, in the case of $A = 2^x$ or $B = 2^x$ with $x \in \mathbb{Z}$, it is possible to equivalently replace the constant multiplication module by a computationally-inexpensive bit shifting approach. Additionally, the computation of the absolute distance increments $\tilde{\Delta}_{\mathfrak{R}}$ and $\tilde{\Delta}_{\mathfrak{I}}$ incurs in virtually no hardware usage as they can be easily obtained by means of simple signal slicing if $2k < |\tilde{z}_{\mathfrak{R}}| < 2k+1$ with $k \in \mathbb{Z}$, and by negation and subsequent slicing if $2t+1 < |\tilde{z}_{\mathfrak{R}}| < 2t$ where $t \in \mathbb{Z}$. This property is applicable as the minimum distance between elements in the real and imaginary axes after the normalization procedure equals $\tilde{\tau} = 2$.

The so-called boundary bits identify each one of the two regions separated by their corresponding boundary line. For example, the boundary bit b_α associated with the boundary

line B_α will be set to '1' when $\tilde{\Delta}_\Re > \tilde{\Delta}_\Im$ and '0' otherwise. In order to determine which puzzle region corresponds to a certain value of $\tilde{z}_i$, it is sufficient to logically combine the required boundary bits and to select the candidate lattice point accordingly by means of a multiplexer. Examples of the circuitry employed to determine the sequence of most favorable nodes can be seen in Figure 5.11. As one can notice, the selection of all the depicted enumerated values can be performed simultaneously, as there is no data dependency between them.

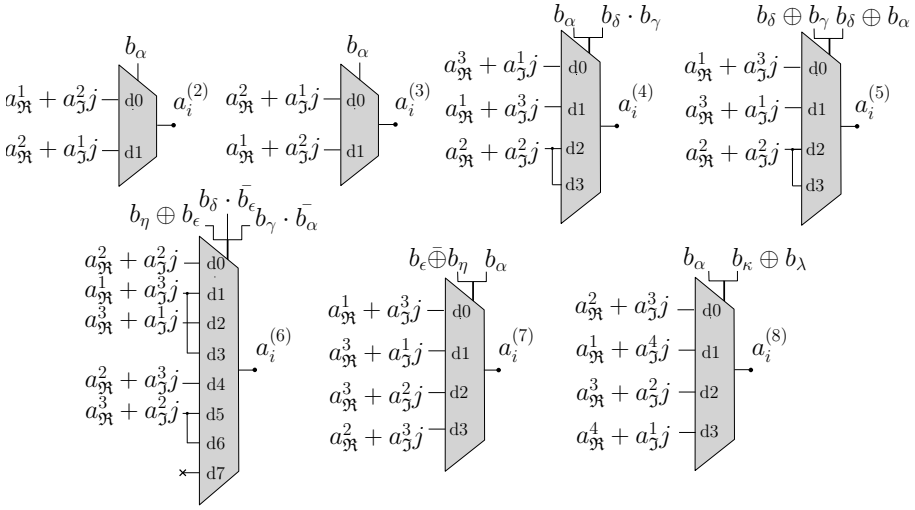

Figure 5.11: Hardware implementation model of the proposed complex enumerator for the first eight enumerated values.

As for the distance computation model used for the MCUs, it should be noted that the previously presented resource sharing approaches yield an excessive computational cost when implemented on the extended grid $\tilde{\mathcal{L}}_{\text{ext}}$. This is due to, on one hand, the increased number of arcs that can be defined based on the model in (5.2), and on the other hand, the higher amount of $\tilde{\chi}^p_{\text{ext}}\tilde{z}_\Re$ and $\tilde{\chi}^q_{\text{ext}}\tilde{z}_\Im$ combinations that need to be computed for both distance calculation approaches in (5.2) and (5.3). These facts have motivated the design of an incremental distance computation model that will be used along with the puzzle enumerator. This way, the distance to a certain lattice point $\tilde{a}^p_\Re + \tilde{a}^q_\Im j$ will be given in terms of the unidimensional distance increment with respect to the first enumerated value in the real and imaginary axis. Two new parameters are introduced, namely $\Xi_\Re = (\tilde{a}^p_\Re - \tilde{a}^1_\Re)$ and $\Xi_\Im = (\tilde{a}^q_\Im - \tilde{a}^1_\Im)$ with $p, q \in \{2, 3, 4\}$ as only 4 values need to be locally enumerated in the real and imaginary axis in order to determine the ordered sequence of 8 nodes. The value of $|\Xi_\Re|$ equals 2 for $p \in \{2, 3\}$, whereas its value is set to 4 whenever $p = 4$. However, the sign of $\Xi_\Re$ will vary depending on the particular location of $\tilde{z}_\Re$ within the real axis $\tilde{\chi}_{\text{ext}}$, as is shown in Figure 5.12. The same rationale applies to the increment in the imaginary axis.

Hence, the PED corresponding to the first enumerated value is computed following the approach in (5.3), whereas the rest of the distances are calculated as

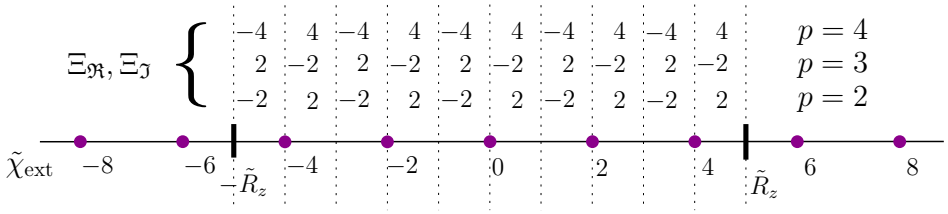

Figure 5.12: Values of the $\Xi_\Re$ and $\Xi_\Im$ increments for different regions of the $\tilde{\chi}_{\text{ext}}$ axis for the 2^{nd}, 3^{rd} and 4^{th} closest nodes in the real axis.

$$\tilde{d}_i(\Xi_\Re, \Xi_\Im) = u_{i,i}^2 \left\{ \left[\delta_\Re + \Xi_\Re^2 + 2\Xi_\Re \left(\tilde{a}_\Re^1 + \tilde{z}_\Re \right) \right] + \left[\delta_\Im + \Xi_\Im^2 + 2\Xi_\Im \left(\tilde{a}_\Im^1 + \tilde{z}_\Im \right) \right] \right\}, \tag{5.4}$$

with $\delta_\Re = (\tilde{z}_\Re + \tilde{a}_\Re^1)^2$ and $\delta_\Im = (\tilde{z}_\Im + \tilde{a}_\Im^1)^2$. This distance computation structure reduces the complexity of the PED calculation for the nodes $\rho > 2$ as most of the terms in (5.4) are either reused from previous computations or are stored as constants. For example, the values for $\delta_\Re$ and $\delta_\Im$ have already been calculated for the PED of $a_i^{(1)}$, namely $u_{i,i}^2(\delta_\Re + \delta_\Im)$. Additionally, the multiplications of the form $2\Xi_\Re x$ can be carried out by means of simple bit shifting and negators due to $|\Xi_\Re| \in (2, 4)$ [see Figure 5.12], whereas $\Xi_\Re^2$ and $\Xi_\Im^2$ can be stored as constants.

Additionally, note that the enumeration areas depicted in Figure 5.7 are symmetrical with respect to the diagonal $\Delta_\Re = \Delta_\Im$. Hence, it is possible to reduce the amount of decision regions and border lines by mapping the intermediate point z_i to the region $\Delta_\Re > \Delta_\Im$. Nevertheless, exploiting the symmetry of the puzzles to reduce the enumeration process to the lower-most triangle is a counterproductive measure for small values of ρ, as the gain derived from having a smaller amount of decision regions does not compensate for the hardware required to map the signal into the desired region and subsequently demap the result of the enumeration process. Consequently, the puzzles shown in Figure 5.7 have been used as originally described for the current implementation of a $\rho = 8$ node enumerator.

5.5 Comparative Analysis

Several features of the complex-plane enumerators under study will be analyzed in this section. First, the complexity in terms of average number of PED computations of the SE tree search when incorporating different enumeration approaches will be assessed. Next, the effect of implementing a different enumeration approach on the candidate node memory size during a SE tree traversal will be analyzed. Finally, the efficiency of the fully-pipelined implementation of the complex-plane enumerators will be studied, where the main focus will be on the latency of the enumeration process and the resource occupation on the targeted FPGA.

5.5.1 Average Number of PED Computations

The amount of distance computations performed by the enumerator is a characteristic of major importance as it can greatly impact the efficiency of the tree-search architecture. For example, in a bounded breadth-first system, a high amount of distance calculations during the enumeration process derives in an exorbitant cost in hardware resources. On the other hand, in a depth-first system such as the SE, the problem of an increased amount of calculations does not affect the area occupation so severely, as the implemented architecture is reused at every iteration. However, the longer run-time of the SE in the high-SNR range causes a considerable increase in the number of unnecessary computations performed by the enumerator, i.e. distance increments of nodes that are not visited during the tree traversal, which derives in an excessive power consumption.

This effect is shown in Figure 5.13, where the average number of distance computations performed by several enumeration techniques during an 8×8 SE tree traversal is depicted. The data in this figure reflects the overwhelming complexity in terms of computed nodes of the full-sorting approach, which is considerably higher and scales faster with the SNR than the rest of the enumeration techniques. The computational saving achieved by reducing the number of candidate subsets from $P_{\mathcal{S}} = 6$ [Burg05] to $P_{\mathcal{S}} = 5$ [Hess07] is more noticeable in the high-SNR range. This comes as a consequence of the null vector being a more probable solution at low SNRs [Christensen07], which results in the early emptying of the subset containing solely the 0 element in [Burg05]. Therefore, the number of effective subsets in the circular enumerator is reduced to 5 during most of the node sequencing process, and consequently, the number of distance computations is virtually identical to that of the approach in [Hess07]. The other two enumeration strategies under study, namely the neighbor expansion technique and the proposed puzzle enumerator, require a considerably smaller amount of distance computations to perform the sequencing of the most favorable nodes. The data displayed in this figure shows that the proposed puzzle enumerator requires the smallest amount of distance computations among all the enumeration approaches under study.

5.5.2 PED Cache Memory Size in Depth-first Systems

The efficient pruning strategy of the SE algorithm is based on the tracking of potentially valid nodes at each level of the tree search. From the point of view of hardware implementation, this entails the incorporation of a PED cache, whose size will depend strongly on the enumeration strategy to be used. If a full-sorting enumeration approach is followed, a $(N-1) \times (|\mathcal{L}|-1)$ memory is required to enable an efficient SE tree traversal, while only $(N-1) \times P_{\mathcal{S}}$ candidate PEDs need to be stored if the subset enumeration approaches in [Burg05] and [Hess07] are implemented. The variable length of the candidate list $\mathcal{C}$ em-

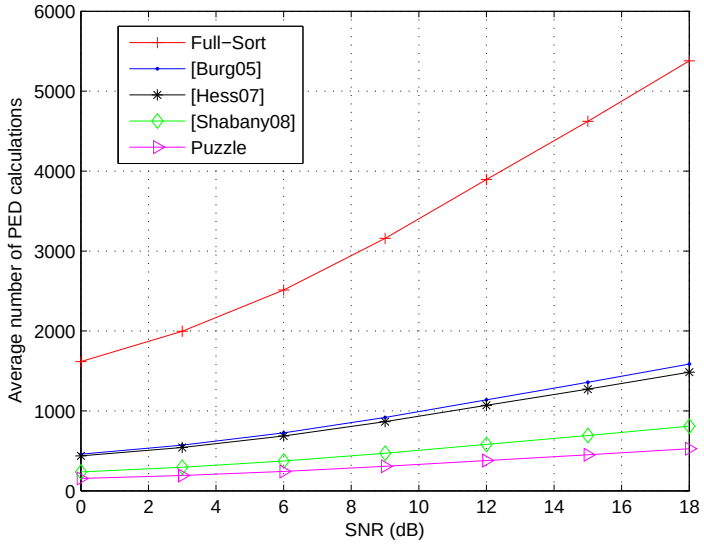

Figure 5.13: Average number of distance computations required by different enumeration techniques in an 8×8 SE tree traversal.

ployed by the approach in [Shabany08b] greatly hinders the sizing of the required PED cache. If an optimum performance of the SE is to be guaranteed, the safest choice is to dimension the memory size for the worst-case scenario. This results in the implementation of a $(N-1) \times (|\mathcal{L}|-1)$ candidate PEDs memory, as setting a limit on the length of $\mathcal{C}$ may result in a faulty enumeration sequence. On the contrary, if the SE implementation is performed using the proposed puzzle enumerator, a reduced memory with only $(N-1) \times 1$ stored candidate PEDs is required.

5.5.3 Latency of the Enumeration Process

The data in Table 5.3 shows the latency resulting from the enumeration process when implemented on an FPGA device. Specifically, the amount of clock cycles (φ) required in order to select the ρ^{th} enumerated value and compute its corresponding PED are depicted.

The sequential nature of the full-sort scheme and the algorithms in [Burg05],[Hess07] and [Shabany08b] is reflected in the provided results. Usually, an initial delay of around 10 cycles is required to select the candidate nodes and compute their corresponding distances. After the selection of the initial value of the enumeration, the rest of the nodes in the sequence and their PEDs are outputted a few clock cycles apart. For the full-sorting scheme, for example, 5 clock cycles are necessary to perform the *compare-and-select* stages required to find the node with the minimum PED among the $|\mathcal{L}|-\rho$ candidate nodes after the $(\rho-1)^{\text{th}}$ value has been identified. The schemes in [Burg05] and [Hess07] manage a considerably smaller list of candidates, which allows for the minimum-search process to be performed in just 3 clock cycles. Nevertheless, the selection and PED computation of the next best sibling node

derives in 2 extra clock cycles per enumeration stage for the approach in [Burg05], whereas only a single cycle is required for the unidimensional subset enumerator due to the simpler sequencing of the nodes within each candidate subset. Additionally, note that $\varphi(\rho)-\varphi(\rho-1)$ gradually increases with ρ in the neighbor expansion model in [Shabany08b]. This is a consequence of the variable length of the candidate list $\mathcal{C}$, which has been dimensioned to fit the worst-case scenario for the current implementation. This way, every time a node is extracted from the candidate list $\mathcal{C}$, two additional nodes are considered to be included. The higher dimension of the candidate list at the later stages of the enumeration algorithm derives in more clock cycles being required in order to perform the search of the minimum PED.

The latency of the enumeration procedure for the proposed puzzle enumerator shown in Table 5.3 reflects the non-sequential nature of the algorithm. The constant latency achieved by the proposed approach is considerably smaller than that of the other enumeration schemes. Additionally, note that most of the latency is due to the PED computation as the sole enumeration procedure can be carried out in just two clock cycles.

	Enumeration + MCUs					Enum. only
	Full sort	**[Burg05]**	**[Hess07]**	**[Shabany08b]**	**Puzzle**	**Puzzle**
$\varphi(1)$	14	12	11	8	8	1
$\varphi(2)$	19	17	15	9	9	2
$\varphi(3)$	24	22	19	12	9	2
$\varphi(4)$	29	27	23	16	9	2
$\varphi(5)$	34	32	27	21	9	2
$\varphi(6)$	39	37	31	26	9	2
$\varphi(7)$	44	42	35	31	9	2
$\varphi(8)$	49	47	39	37	9	2

Table 5.3: Amount of clock cycles required to obtain the ordered sequence of child nodes.

5.5.4 Resource Occupation

The fully-pipelined architectures of the proposed enumerator along with the state-of-the-art enumerators reviewed in this chapter have been implemented on a Virtex-5 XC5Vlx30-3 FPGA. The device occupation results and the maximum clock frequency value shown in Table 5.4 for the different designs have been obtained by means of the *place and route* tool included in the *System Generator for DSP* software.

The considerably smaller hardware resource usage of the proposed scheme is noticeable from the data shown in this table. The total gate equivalent (GE) count of the puzzle enumerator is roughly 15 times smaller than that required by the approaches in [Burg05], [Hess07] and [Shabany08b], while the slice occupation is around 5 times smaller. When compared to the more straightforward full-sorting approach, the proposed enumeration unit yields a notably lower slice occupation and a 30 times smaller GE count. As for the un-

	Enumeration + MCUs					
	Full sort	**[Burg05]**	**[Hess07]**	**[Shabany08b]**	**Puzzle**	**Puzzle$^{\theta}$**
Number of occupied slices	4243	2198	1735	1922	455	297
Number of slice LUTs	11652	5602	4295	4883	900	696
Number of DSP48Es	10	10	10	10	10	10
Total gate equivalent count	573 K	351 K	268 K	288 K	18 K	13 K
Maximum clock frequency(MHz)	300	200	278	270	300	370

Table 5.4: Device occupation and maximum achievable frequency for different complex enumerators with MCUs in a fully-pipelined scheme.

ordered enumeration unit presented in Section 5.3.3 (Puzzle$^{\theta}$), an additional resource saving is achieved due to its simplified structure.

Additionally, the device occupation results for the puzzle enumerator core without the MCUs are provided in Table 5.5 for completion. Resource usage results show the low complexity of the proposed enumerator, which can be implemented with just 3K GE and without the allocation of any costly multiplier units.

	Enumeration only	
	Puzzle	**Puzzle$^{\theta}$**
Number of occupied slices	111	71
Number of slice LUTs	299	209
Number of DSP48Es	0	0
Total gate equivalent count	3 K	2 K
Maximum clock (MHz)	360	385

Table 5.5: Device occupation and maximum achievable frequency for the proposed complex enumerators in a fully-pipelined scheme.

5.6 Chapter Summary

This chapter has dealt with the design and implementation of a parallel complex-plane Schnorr-Euchner enumerator for its application to VP systems. As opposed to other enumeration algorithms found in the literature, where the symbols are selected in a sequential fashion, the proposed puzzle enumerator performs an independent identification of the symbols within the ordered sequence. This way, the selection of the first ρ symbols according to the Schnorr-Euchner enumeration can be performed simultaneously, which reduces the delay of the enumeration process to a great extent, specially for high values of ρ. When used in a SE tree traversal, the ability of the proposed scheme of independently selecting the values in the enumerated sequence derives in a reduced amount of distance calculations and a smaller size of the PED candidate cache when compared to other enumeration approaches.

Additionally, the fully-pipelined and high-throughput implementation of the proposed algorithm and other state-of-the-art complex-plane enumerators has been carried out. Comparative results on the device occupation, multiplier utilization, delay and maximum achiev-

able frequency of the different enumerators has shown that the proposed approach is able to perform the enumeration task at a higher symbol frequency with minimum latency and hardware resource occupation.

The puzzle enumerator is specially useful for fully-pipelined high-throughput precoders where a fixed amount of symbols need to be enumerated. Nevertheless, the proposed enumeration unit can also be used in recursive systems, such as the SE, if a limited enumeration sequence is employed.

Chapter 6

Implementation of Fixed-complexity Algorithms for VP

6.1 Introduction

As it has been detailed previously in Chapter 4, several algorithms have been proposed for the closest-point lattice search problem in VP systems hitherto. Nevertheless, the optimality of these algorithms has been assessed mainly in terms of error-rate performance and computational complexity, leaving the hardware cost of their implementation an open issue. In this chapter, the problem of designing and implementing a high-throughput tree-search algorithm for VP systems will be addressed.

So far, the amount of published work related to the implementation of precoding algorithms has been very limited. The first very-large-scale integration (VLSI) implementation of a lattice reduction-aided VP system on an application-specific integrated circuit (ASIC) is described in [Burg07]. The implemented algorithm avoids the search for the optimum perturbation vector by performing a simple rounding operation on an equivalent lattice-reduced version of the precoding matrix. However, the error introduced by the quantization and rounding operations leads to a suboptimum BER performance. Another noteworthy research work on the field of implementation was presented in [Bhagawat08], where the first hardware implementation of a simple variant of DPC on a Xilinx Virtex-II FPGA device is described. Despite the novelty of the proposed precoding structure, the throughput of the algorithm is of just 51 Mbps for a considerable resource usage. Additionally, the implementation design of a THP precoder used to tackle multiple-access interference in multi carrier code division multiple access schemes has been described in [Lin06]. However, implementation results on low-complexity and high-throughput vector precoders are still missing in the open literature.

Despite the lack of published research in the area of hardware architectures for precoding algorithms, the implementation issues of tree-search schemes in MIMO detection scenarios have been widely studied. For example, the FPGA implementation of the FSD

detector has been analyzed in [Barbero06d, Barbero06c, Barbero08], whereas the hardware architecture of the K-Best tree search considering a real equivalent model was researched in [Wong02, Wenk06, Li06, Guo06, Shabany08a, Shen10]. Moreover, the implementation of a K-Best detector with suboptimum complex-plane enumeration was performed in [Chen07, Mahdavi10]. The adaptation of these tree-search schemes to precoding systems implies many variations with respect to the original description of the algorithms. Even if many lessons can be learned from the hardware architecture of tree-search techniques for point-to-point MIMO systems, the peculiarities of the precoding scenario render the results of the aforementioned publications inadequate for the current research topic.

Thus, this chapter will address the high-throughput implementation of fixed-complexity algorithms for VP systems. The K-best and FSE tree-search algorithms for the 4×4 antennas setup will be implemented on a Xilinx Virtex VI FPGA following a rapid-prototyping methodology. In order to comply with the high-throughput requirement, both schemes will operate in the complex-plane and will be implemented in a fully-pipelined fashion providing one output per cycle. In the case of the K-Best, this entails the utilization of the puzzle enumerator introduced in Chapter 5.

6.2 General Architecture Overview

Both tree-search schemes share the same general distance computation structure, as can be seen in Figure 6.1. The lack of loops in the hardware architecture of the fixed-complexity tree-search techniques enables a high-throughput and fully-pipelined implementation of the data perturbation process.

The AEDs of the candidate branches are computed by accumulating the PEDs calculated at the local distance processing units (DPUs) to the AEDs of the previous level. This way, the AEDs down to level i corresponding to the considered candidate branches, namely $\left[D_i^{(1)}, \ldots, D_i^{(\psi_i)}\right]$, are passed on from DPU i to DPU $i-1$. The parameter ψ_i stands for the number of candidate branches at each level of the tree search, being it $\psi_i = K \ \forall i$ for the K-Best and $\psi_i = \prod_{j=i}^{N} n_j$ for the FSE model.

Two input memory blocks, named *Data Memory* and *Channel Memory*, have been included to store the data symbols and the values of the triangular matrix $\boldsymbol{U}$, respectively. The off-diagonal matrix coefficients are stored as u_{ij}/u_{ii}, whereas the diagonal values are in the form of u_{ii}^2 to simplify the calculation of (4.1) and (4.2). Note that the matrix preprocessing stage required by the FSE and K-Best approaches has not been included in the hardware design. The computation of the intermediate points z_i requires the values of all previous $\delta_j = a_j + s_j$. To avoid redundant calculations, the set of values $\left[\delta_j^{(1)}, \ldots, \delta_j^{(\psi_j)}\right]$ for all $j > i$ is transferred to DPU i, as is shown in Figure 6.1.

The hardware architecture of the first DPU is common in both schemes. The computation

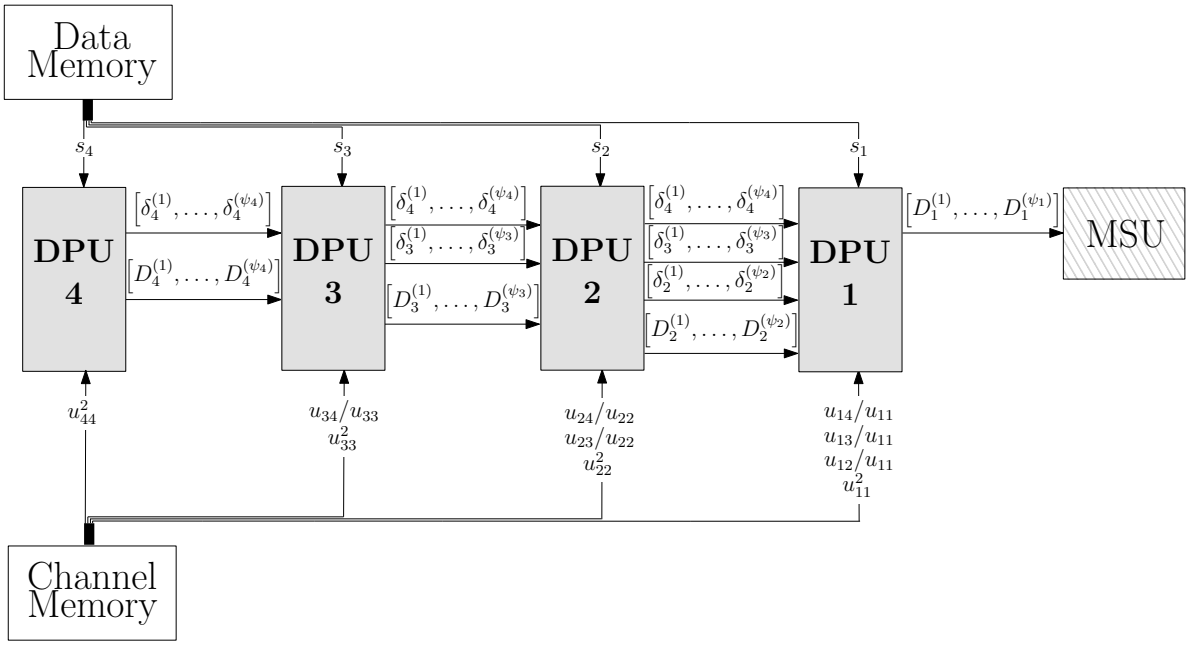

Figure 6.1: General hardware architecture of the fixed-complexity tree-search techniques for an $N = 4$ user system.

of the Euclidean distances in this level does not involve any data from previous levels, and therefore, the only operation to be performed is to select the ψ_N lattice values closest to s_N and to compute the corresponding PEDs. Given that the position of the modulation's constellation within the complex lattice is known beforehand, and considering the symmetries of the complex lattice, it is possible to select the nodes to be passed on to the next level without performing any extra distance calculations and sorting procedures. Additionally, the hardware structure of the last DPU is also equal for both algorithms, as only the most favorable child node that stems from each one of the ψ_2 parent nodes needs to be expanded at this level. Such a task can be performed by simply rounding the value of z_1 to the position of the nearest lattice point.

The main and crucial differences between the FSE and K-Best tree-search algorithms rely on the DPUs of levels $1 < i < N$. In the following sections, the hardware architecture of the DPUs for the K-Best and FSE schemes will be studied.

6.3 DPU for the K-Best

The difficulty of performing the sorting procedure in the complex plane, where the amount of nodes to be considered is higher, and the intricacy of complex-plane enumeration have led to the dominance of RVD as the preferred technique when implementing the K-Best tree search. Nevertheless, direct operation on the complex signals is preferred from an implementation point of view as the length of the search tree is halved, and hence, the latency and critical

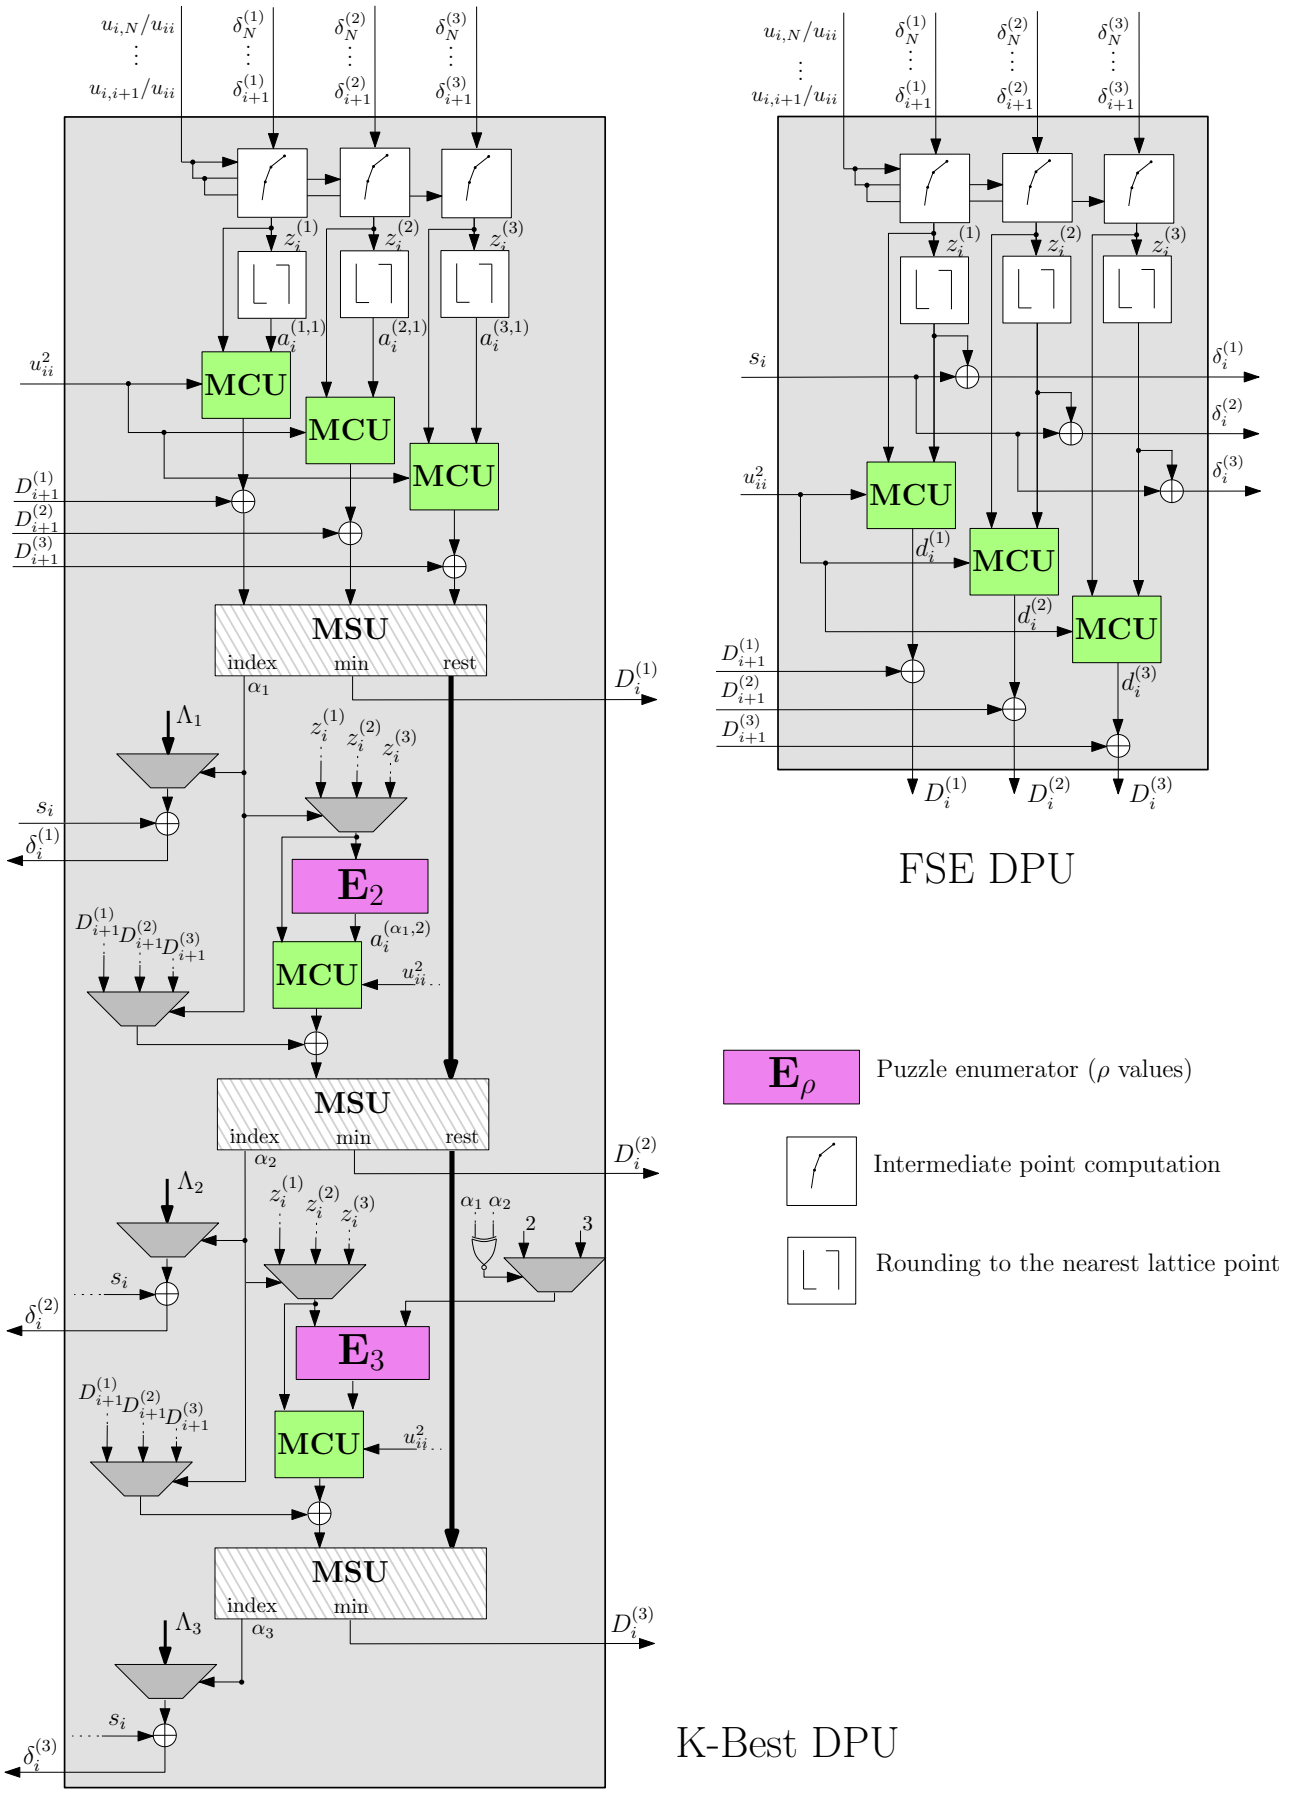

Figure 6.2: Block diagram of the K-Best and FSE DPU.

path of the design can be shortened.

6.3.1 Structure of the Sorting Stage

Regardless of the domain of the signals to be used, the bottleneck in this type of systems is usually the sorting stage performed at each tree level. The number of child nodes that stem from the same parent node will be defined as B, being its value $B = |\mathcal{L}|$ for the complex-plane model, whereas $B = \sqrt{|\mathcal{L}|}$ will be required for the RVD scheme. The PED calculation and subsequent sorting procedure on the KB child nodes at each level is a computationally expensive process that compromises the throughput of the whole system. With the aim of alleviating the burden of the sorting stage, the use of the Schnorr-Euchner ordered sequence of child nodes and the subsequent merging of the sorted sublists is proposed in [Wenk06]. Even if the proposed scheme is implemented on an RVD model due to the simplicity of the local enumeration, it is possible to extend it to the complex-plane if a low-complexity enumerator, such as the puzzle enumerator presented in Chapter 5, is utilized. Additionally, a fully-pipelined RVD architecture of the sorted sublists algorithm is proposed in [Tsai10] for high-throughput systems. By dividing the real axis into $2B$ regions and storing the corresponding enumeration sequences in LUTs, the algorithm is able to determine the child node order by means of a simple slicing procedure. Nevertheless, this technique is advantageous only when operating with RVD symbols as the amount of data to be stored and the quantity of non-overlapping regions grow remarkably when complex-valued symbols are utilized. In any case, the use of the sublist merging approach reduces the amount of PED computations at each level to $K^2 \leq KB$.

6.3.1.1 The Winner Path Extension Algorithm

The number of costly distance computations can be further reduced by implementing the winner path extension (WPE) selection approach presented in [Mondal08]. The proposed scheme selects the K most favorable branches in K iterations by performing just $2K - 1$ PED computations. An illustrative example of the WPE algorithm is depicted in Figure 6.3 for a system with $B = 4$ and $K = 3$. The child nodes at a certain tree level i are tagged as $a_i^{(x,y)}$, where x represents the index of the parent node and y denotes the position of that certain node within the ordered Schnorr-Euchner sequence of child nodes.

The WPE sorting procedure is based on the generation and management of a node candidate list $\mathcal{A}$. This way, the child node corresponding to the k^{th} most favorable branch is extracted from the candidate list of the k^{th} sorting stage $\mathcal{A}_k$. The initial values in the candidate list are comprised of the AEDs down to level i of the best child nodes that stem from each one of the K parent nodes, which gives $\mathcal{A}_1 = \left\{a_i^{(1,1)}, \ldots, a_i^{(K,1)}\right\}$. The *winner* branch in the initial sorting stage, or equivalently the first of the K most favorable branches,

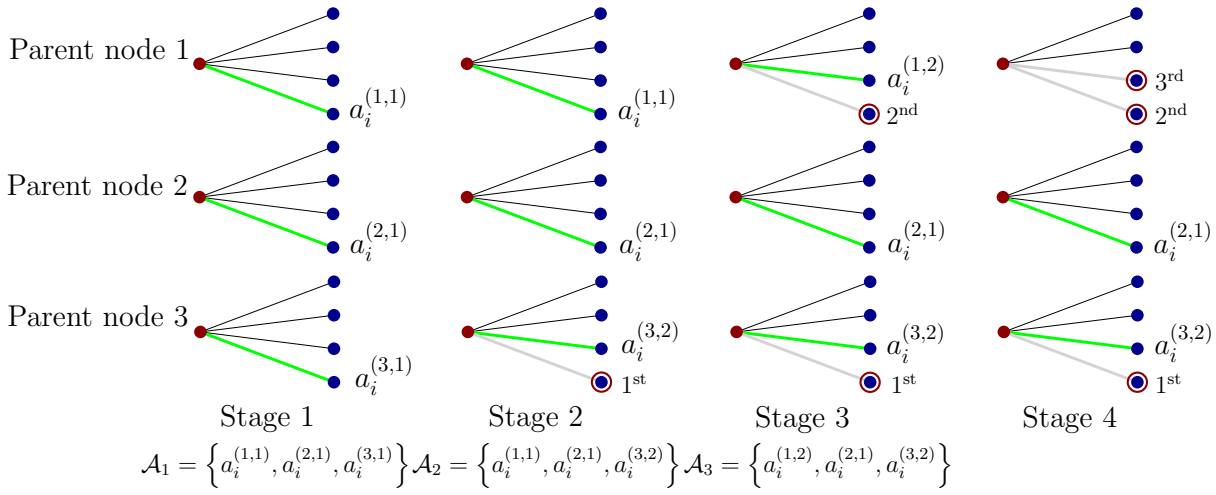

Figure 6.3: Illustrative example of the WPE distributed sorting algorithm for a system with $K = 3$ and $B = 4$.

is selected as the branch with the smallest AED within $\mathcal{A}_1$. The PED of the second most favorable child node that stems from the same parent node as the latest appointed *winner* branch is computed ($a_i^{(3,2)}$ in the example illustrated in Figure 6.3) and the AED of the resulting tree branch is added to the candidate list $\mathcal{A}_2$. The algorithm proceeds accordingly until the K required branches have been identified.

6.3.2 Structure of the K-Best DPU

The structure of the proposed K-Best DPU is depicted in Figure 6.2 for a system with $K = 3$. The branch selection procedure is carried out in K fully-pipelined sorting stages following a modified version of the WPE algorithm presented in [Mondal08],[Tsai10]. First of all, the computation of the intermediate points is performed for each one of the K branches that are passed on from the previous level. The set of best child nodes that stem from each parent node can be computed by simply rounding off the value of the intermediate point to the nearest lattice point. The distance increments for those K best children are computed by K MCUs and are accumulated with their corresponding D_{i-1} values. These distance values and their corresponding branches comprise the candidate list Λ_1. The minimum AED within Λ_1 is found at the MSU by simple concatenation of compare-and-select blocks. The MSU also outputs the index of the first *winner* branch $\alpha_1 \in \{1, \ldots, K\}$ so that the appropriate value of z_i can be selected for the local enumeration procedure.

At the second stage of the sorting procedure, the $a_i^{(\alpha_1,2)}$ node needs to be identified for any parent node index α_1. This task is performed by the E_2 block, which comprises a puzzle enumerator that outputs the second most favorable node given a certain value of z_i. However, in the subsequent stages of the algorithm, the enumeration procedure will depend on the index of the previously appointed *winner* branches. Hence, if $\alpha_1 = \alpha_2$, the third most

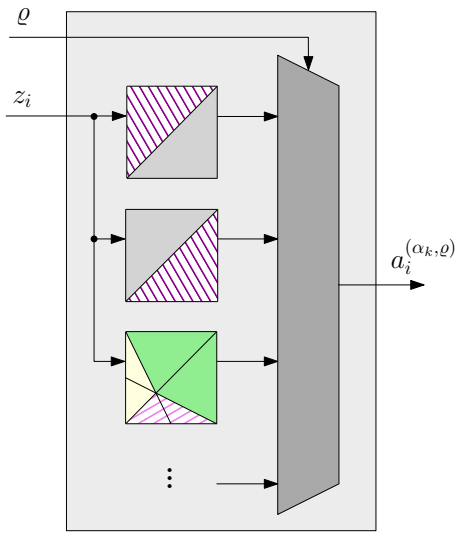

Figure 6.4: Puzzle enumeration unit with selectable sequence index ϱ.

promising child node will need to be expanded, namely $a_i^{(\alpha_1,3)}$, whereas the second most favorable node in the α_2 branch ($a_i^{(\alpha_2,2)}$) will be required if $\alpha_1 \neq \alpha_2$. Consequently, the new candidate branch to be included in the $\mathcal{A}_k$ candidate list at the k^{th} sorting stage, will require the expansion of the ϱ^{th} most favorable child node, where ϱ may take any value within the set $\{2, 3, \ldots, k\}$.

The enumeration approach at each sorting stage k has been carried out by means of a puzzle enumerator unit capable of ascertaining the sequence of the first k child nodes. The desired child node in the ordered sequence is determined by an additional input variable which keeps track of the amount of already expanded child nodes for each parent node. The structure of the proposed enumeration unit is depicted in Figure 6.4. The puzzle enumerator has been selected as the enumeration scheme to be used along with the WPE due to its lower hardware resource demand and non-sequential nature, as already discussed in Section 5.5.

6.4 DPU for the FSE

The intricate node ordering and selection procedure required by the K-Best algorithm is replaced by a simple Schnorr-Euchner enumerator in the FSE tree-search model. This derives in a considerably simpler DPU architecture of the FSE scheme.

Figure 6.2 depicts the structure of the FSE DPU, where the block diagram for $n_T^{(i)} = \prod_{j=1}^{i-1} n_j = 3$ and $n_i = 1$ is represented. First of all, the data of the $\{\delta_{i+1}, \ldots, \delta_n\}$ values transferred from level $i+1$ are used to compute the intermediate value z_i for each one of the parent nodes. Afterwards, the node selection procedure is performed by means of a simple rounding operation when $n_i = 1$, as depicted in the illustrative example in Figure 6.2, or by means of the unordered puzzle enumerator presented in Chapter 5 for the cases where

$n_i > 1$. The PEDs of the selected nodes are then computed by $n_T^{(i)}$ MCUs and accumulated to the AEDs from the previous level.

6.5 Design Considerations

This section addresses the design parameter selection for the fixed-complexity algorithms to be implemented in hardware. Additionally, the impact of applying an approximate norm for the computation of the distance increments is studied from an error-rate performance point of view.

6.5.1 Choice of the Design Parameters

The configurable parameters K and $\boldsymbol{n}$ offer a flexible trade-off between performance and complexity for the K-Best and FSE encoders, respectively. These configuration parameters establish the shape of the search tree, which in turn determines the amount of hardware multipliers required for its implementation. Embedded multipliers are scarce in FPGA devices and are considered an *expensive* resource in ASIC designs. Thus, the number of multiplication units required by the tree-search algorithm has been regarded as the critical factor in the current hardware architecture design. For the sake of a fair comparison, the configuration parameters of the fixed-complexity tree-search methods have been selected so as to yield a similar amount of allocated embedded multipliers.

The amount of embedded multipliers required for the K-Best tree search with a WPE sorting scheme depends strongly on the enumeration approach used to create the ordered sequence of child nodes. Despite the fact that the puzzle enumerator has already been selected as the preferred node sequencing scheme for the K-Best structure due to its low complexity, the current analysis on the multiplier allocation volume will be performed on a more general scope. The rationale for this is twofold; on one hand, it will enable the study of the amount of allocated multipliers in a fixed-complexity tree-search structure with different enumeration techniques, and on the other hand, the provided analysis will better showcase the benefits of the proposed enumerator when implemented upon a non-iterative tree-search scheme.

Considering that 3 multipliers are used for the multiplication of two complex terms, the number of multiplication units required for the K-Best tree-search structure can be computed as

$$N_{\mathrm{MUL,KB}} = 6K + 3(N-2)\left[P_{\mathcal{S}}K + P_{\mathcal{C}}(K-1)\right] + 3K\left\lceil\frac{N(N-1)}{2}\right\rceil,$$

where the parameters $P_{\mathcal{S}}$ and $P_{\mathcal{C}}$ are related to the enumeration strategy used to select the child nodes to be expanded. The number of locally-enumerated subsets that are defined during the enumeration procedure is represented by the parameter $P_{\mathcal{S}}$. This way, the puzzle

enumerator and the neighbor expansion scheme [Shabany08b] require a value of $P_{\mathcal{S}} = 1$, whereas the circular subset scheme [Burg05] and the unidimensional subset enumerators [Hess07] work with values of $P_{\mathcal{S}} = 6$ and $P_{\mathcal{S}} = 5$, respectively. The parameter $P_{\mathcal{C}}$, on the other hand, represents the amount of nodes that are added to the candidate list after enumerating the latest child node. The puzzle enumerator, the circular subset enumeration approach and the unidimensional subset enumerator give values of $P_{\mathcal{C}} = 1$. For the neighbor expansion scheme however, the worst-case scenario has been contemplated, which gives $P_{\mathcal{C}} = 2$.

The effect of the enumeration unit on the structure of the FSE is more reduced as most of the tree levels consider the expansion of a single node. Nevertheless, for those levels $i \neq N$ with $n_i > 1$, the choice of complex-plane enumerator will determine the amount of allocated hardware resources. The total amount of embedded multipliers for an FSE tree structure can be computed as

$$N_{\mathrm{MUL,FSE}} = \sum_{i=1}^{N} N_{\mathrm{MUL,FSE}}^{(i)},$$

where the number of multipliers for each individual tree search level is given by

$$N_{\mathrm{MUL,FSE}}^{(i)} = \begin{cases} 3\sum_{j=i+1}^{N} n_T^{(j)} + 3n_T^{(i)} & \text{for } i = N \text{ or } n_i = 1 \\ \\ 3\sum_{j=i+1}^{N} n_T^{(j)} + 3n_T^{(i+1)}\left[P_{\mathcal{S}} + P_{\mathcal{C}}(n_i - 1)\right] & \text{otherwise.} \end{cases} \tag{6.1}$$

The number of required embedded multipliers for the K-Best and FSE tree-search techniques is shown in Figure 6.5 for a system with $N = 4$ single-antenna users. The amount of multiplier units is given as a function of the number of candidate branches, namely K and n_T for the K-Best and FSE approaches, respectively. The data in this figure show that the amount of hardware resources in the K-Best tree-search model grows linearly with K. The multiplier demand growth rate is more elevated for the circular and unidimensional subset enumeration approaches due to the higher number of defined candidate subsets $P_{\mathcal{S}}$. The smaller amount of initial PED computations required in the single-subset schemes ($P_{\mathcal{S}} = 1$) derives in a smaller multiplier occupation for the neighbor expansion and puzzle enumeration approaches. Among the latter, the puzzle enumerator is the node sequencing scheme that requires the smallest multiplier count as only a single new candidate node is considered for each enumeration stage, i.e. $P_{\mathcal{C}} = 1$.

The complexity of the tree search in terms of allocated multipliers does not grow linearly with n_T for the FSE case, as is shown in Figure 6.5. This is due to the n_i values being differently distributed through the tree configuration vector depending on the divisibility of n_T. Hence, for a prime n_T value, there is no need for the implementation of a complex-plane enumerator as the only tree level with $n_i > 1$ is located at the top of the tree, where the nodes can be sequenced straightforwardly. This is reflected in the multiplier count expression

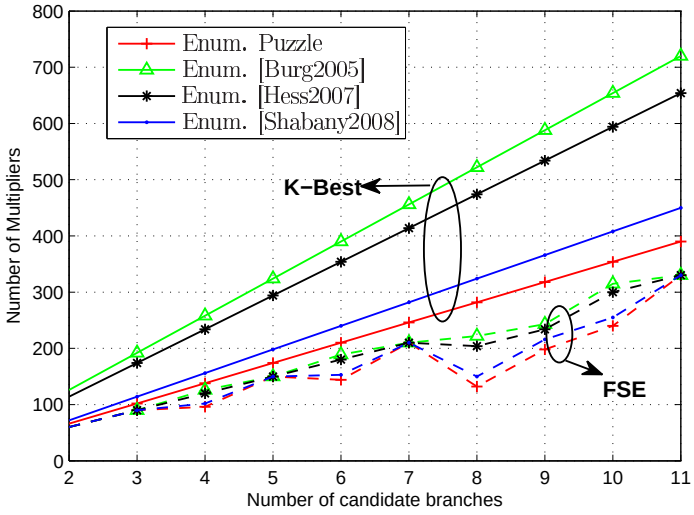

Figure 6.5: Number of required multipliers for the K-Best and FSE tree-search techniques as a function of the number of candidate branches.

in (6.1), where the amount of embedded multipliers for systems with a prime n_T value can be computed without considering the enumerator dependant $P_{\mathcal{C}}$ and $P_{\mathcal{S}}$ parameters. For the rest of the n_T values, namely divisible values of n_T, the effect of implementing different enumeration techniques is not very significant either, as is shown in Figure 6.5. The reason for this is that the effect that a certain enumerator has on the hardware complexity of the FSE tree search is only limited to those tree levels $i \neq N$ where $n_i > 1$.

In order to assess the hardware resource occupation required for the the implementation of the different tree-search algorithms, the design parameter values $K = 7$ and $n_T = 10$ have been selected for the K-Best and FSE models, respectively. This choice of parameters ensures a similar multiplier occupation for both schemes and offers a significantly better error-rate performance than other lower K and n_T pairs, e.g. $K = 6$ and $n_T = 9$.

The BER vs SNR curves of the implemented fixed-complexity schemes are depicted in Figure 6.6. As one can notice, the error-rate performance of the implemented models is close to the optimum set by the SE in the low-to-mid SNR range. However, a slight performance degradation of 0.5 dBs is noticeable for the FSE model at high-SNRs, whereas the performance gap of the K-Best structure increases with the SNR, reaching up to 3 dBs at a BER of 10^{-6}.

6.5.2 Implementation of an Approximate Norm

A significant portion of the hardware resources in the implementation of any tree-search algorithm is dedicated to computing the ℓ^2 norms required by the cost function in (2.8) (or equivalently in (4.3) for ordered tree traversals). Additionally, the long delays associated with squaring operations required to compute the PEDs account for a significant portion

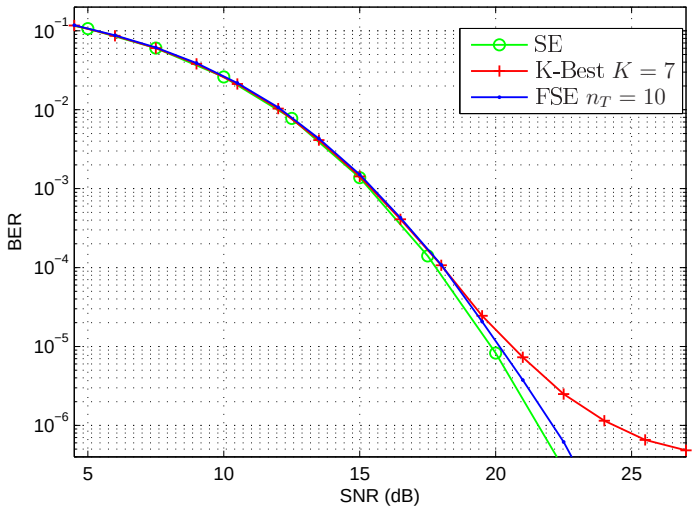

Figure 6.6: BER performance of the implemented FSE and K-Best tree-search structures.

of the latency of the fixed-complexity tree-search architectures. It is possible to overcome these problems by using an alternative norm that prevents the use of the computationally-expensive squaring operations.

The application of the modified-norm algorithm (MNA) [Burg04] entails two main benefits: on one hand, a simplified distance computation scheme that immediately reduces silicon area and delay of the arithmetic units can be performed, and on the other hand, a smaller dynamic range of the PEDs is achieved. The key point of the MNA is to compute the square root of the accumulated and partial distance increments, namely $E_i = \sqrt{D_i}$ and $e_i = \sqrt{d_i}$, respectively. Hence, the accumulation of the distance increments in this equivalent model gives $E_i = \sqrt{E_{i+1}^2 + e_i^2}$. An approximate norm can now be applied to get rid of the computationally-expensive squaring and square root operations, such that $E_i \approx f(|E_{i+1}|, |e_i|)$. This way, the distance accumulation and computation in (4.1) can be reformulated as

$$E_i = E_{i+1} + e_i, \tag{6.2}$$

with

$$e_i = u_{i,i} \left(|\Re(a_i + z_i)| + |\Im(a_i + z_i)| \right) \tag{6.3}$$

for the $\ell^{\tilde{1}}$-norm variant of the algorithm. The norm approximation can also be performed following the $\ell^{\tilde{\infty}}$-norm simplified model, in which case the following expressions should be considered

$$E_i = \max\left(E_{i+1}, e_i\right), \tag{6.4}$$

with

$$e_i = u_{i,i} \left[\max\left(|\Re(a_i + z_i)|, |\Im(a_i + z_i)| \right) \right].$$

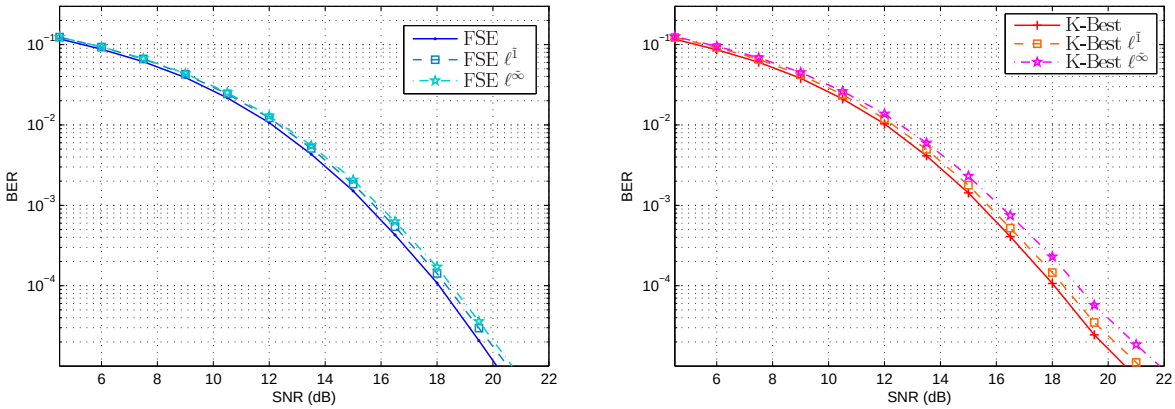

(a) FSE (b) K-Best

Figure 6.7: BER performance degradation introduced by approximating the ℓ^2 norm by the simplified $\ell^{\tilde{1}}$ and $\ell^{\tilde{\infty}}$ norms in the FSE ($n_T = 10$) (a) and K-Best ($K = 7$) (b) tree-search approaches.

The implementation of an approximate norm impacts the error-rate performance of the VP system differently depending on the tree-search strategy used in the perturbation process. This fact is shown in Figure 6.7, where the BER performance degradation introduced when approximating the ℓ^2 norm by the suboptimum $\ell^{\tilde{1}}$ and $\ell^{\tilde{\infty}}$ norms is depicted for the FSE and K-Best tree-search approaches. For the FSE case depicted in Figure 6.7(a), the use of an approximate norm only affects the accumulated distances related to the candidate branches, but not the branches themselves. This is due to the fact that the nodes expanded at each level where $n_i \leq 2$ are the same regardless of the norm used to compute the distance increments to z_i. In the K-Best model, on the other hand, the node selection procedure is solely based on previously computed distances, and therefore, the introduction of a simplified norm will noticeably alter the structure of the candidate branches. Consequently, a higher error-rate performance degradation of the K-Best algorithm with an approximate norm can be expected when compared to the norm-simplified FSE model.

The implementation of the approximate $\ell^{\tilde{1}}$ norm yields a high-SNR performance loss of 0.22 dB and 0.25 dB for the FSE and K-Best fixed-complexity algorithms, respectively. Due to the worse approximation of the Euclidean distances performed by the suboptimum $\ell^{\tilde{\infty}}$ norm, the performance gap with respect to the optimum FSE and K-Best structures is widened in this case. This way, a performance loss of 0.45 dBs is experienced by the simplified $\ell^{\tilde{\infty}}$-FSE model, whereas an error-rate degradation of 0.85 dBs is suffered by the K-Best in the high-SNR regime. In any case, the implementation of an alternative norm does not alter the diversity order of the VP scheme.

The computational complexity reduction yielded by both norm approximation approaches is similar, whereas the performance is slightly better for the $\ell^{\tilde{1}}$ norm approximation. Consequently, the $\ell^{\tilde{1}}$ norm-simplified model will be considered for hardware implementation.

	K-Best	FSE	K-Best $\ell^{\tilde{1}}$	FSE $\ell^{\tilde{1}}$
Number of Occupied Slices (39,360)	25 %	10%	26%	10 %
Number of Slice Registers (314,880)	10 %	5 %	10%	4%
Number of Slice LUTs (157,440)	20 %	7 %	20%	7%
used as logic	11 %	6 %	11%	6 %
used as memory	27 %	2 %	27%	2%
Number of DSP48e1s (576)	40 %	40 %	31 %	29%
Q (Gbps)	5.52	5.63	5.34	5.62

Table 6.1: Hardware resource occupation and throughput of the tree-search architectures under study.

6.6 Implementation Results

The proposed tree-search architectures have been implemented on a Xilinx Virtex VI FPGA (XC6VHX250T-3). The occupation results have been obtained by means of the *place and route* tool included in the *System Generator for DSP* software.

Table 6.1 depicts the device occupation summary of the implemented vector precoders. Even if the FSE and K-best models use a similar amount of embedded multipliers (DSP48e1), the device occupation in terms of slices is considerably higher for the latter. This is due to the longer latency of the K-best architecture caused by the distributed sorting procedure, which ultimately results in a great amount of data being stored in several pipeline stages. As a consequence to this, around 27% of the slice LUTs are used as memory in the K-best implementation, as opposed to the 2% utilized by the FSE for the same purpose. As already anticipated, the utilization of the approximate $\ell^{\tilde{1}}$ norm yields a notable reduction in the amount of allocated embedded multipliers for both fixed-complexity tree-search models.

The maximum throughput of the implemented architectures in terms of processed gigabits per second is also shown in Table 6.1. For a system with N users and a constellation of P elements, the throughput for fully-pipelined architectures can be computed as

$$Q = N f_{\text{clock}} \log_2(P),$$

where f_{clock} represents the maximum working frequency of the design as given by the *Post-Place & Route Static Timing Report*. Both tree-search algorithms achieve a very high data-processing throughput (in the range of 5 Gbps) due to the loop-less structure that enables the processing of a new data vector at every clock cycle.

6.7 Chapter Summary

This chapter has addressed the issues of a fully-pipelined implementation of the FSE and K-Best tree-search approaches for a 4×4 VP system. The sorting stages required by the K-best scheme have been performed by means of the WPE distributed sorting strategy along

with the puzzle enumerator presented in Chapter 5. The latter has also been incorporated into the FSE structure to determine the child nodes to be expanded in those tree levels $i < N$ where $n_i > 1$. The design parameters that establish the performance-complexity trade-off of these non-recursive tree-search approaches have been set so as to yield a similar count of allocated embedded multipliers. Additionally, the use of an approximate norm to reduce the computational complexity of the PED calculations has been contemplated. The complexity reduction achieved by both norm simplification models $\ell^{\tilde{\infty}}$ and $\ell^{\tilde{1}}$ is similar, whereas the achieved error-rate performance is slightly better for the latter. Consequently, the $\ell^{\tilde{1}}$ approximate norm model has been considered for hardware implementation.

Provided performance results have shown a close-to-optimal performance and very high achievable throughput for both techniques. Nevertheless, the error-rate performance of the FSE has been shown to considerably outperform the K-Best in the high-SNR range. Additionally, the provided FPGA resource occupation results have demonstrated the greater efficiency of the FSE architecture when compared to the K-Best fixed-complexity structure.

Due to the good performance, occupation results and simplicity of implementation, it is concluded that the FSE is best suited for the practical implementation of fixed-complexity and high-throughput vector precoders.

Chapter 7

Low-complexity Sequential Tree-search Algorithms for VP

7.1 Introduction

The so-far addressed K-Best and FSE bounded breadth-first techniques represent a low-complexity alternative to the widely-used SE algorithm. The advantages of a non-iterative tree traversal, such as the fixed-complexity of the tree search and the high data processing throughput, have been presented and analyzed in previous chapters of this thesis. Nevertheless, the branch pruning strategy in bounded breadth-first schemes also presents other shortcomings.

Given that the selection of the nodes at each stage is performed based on the PEDs or AEDs up to that level and not the Euclidean distance associated with the entire branch, the validity of the pruning is not certain until the bottom of the tree is reached. This speculative pruning strategy results in a great amount of unnecessary distance computations and an excessive power consumption. Consequently, many of the PED calculations performed during a bounded breadth-first tree search, along with the corresponding allocated hardware resources and power usage, can be spared by traversing the tree in both forward and backward directions.

The depth-first model of an originally breadth-first algorithm is presented in this chapter for the computation of the perturbation vector in VP systems. The tree-search structure of the proposed sequential best-node expansion (SBE) scheme is inspired by the so-called conditioned ordered successive interference cancelation (COSIC) detector introduced in [Hess07][Wenk10a]. The presented tree-search approach also features a novel user ordering strategy that provides a good trade-off between error-rate performance and complexity of the algorithm in terms of amount of visited nodes.

With the aim of further reducing the complexity of the SBE tree search, a simplified version of the proposed tree traversal strategy is presented. The low-complexity SBE (LC-SBE) algorithm uses an innovative system of flexible run-time constraints along with an

approximate distance computation model to allow for an efficient implementation of a vector precoder at the cost of a negligible performance loss.

7.2 The SBE Algorithm

The proposed scheme features two main concepts: a tree traversal architecture where only the most promising nodes are expanded and an innovative matrix preprocessing strategy. The former is based on the COSIC tree search described for signal detection in MIMO scenarios in [Hess07][Wenk10a], while the latter has been specially designed to yield a good trade-off between performance and run-time of the SBE algorithm.

7.2.1 SBE Tree-search Architecture

The tree structure of the SBE is a depth-first implementation of an originally breadth-first algorithm, which enables the usage of a sphere constraint for additional pruning. This way, the Euclidean distance related to the first computed branch is considered as the initial radius R, which is updated every time a leaf node with a smaller AED is reached.

The main motivation for the change of tree traversal strategy is the reduction of the unnecessary distance computations performed in the fully-parallel scheme. As a consequence of adopting a sequential scheme, the complexity of the tree search is no longer fixed. However, by avoiding a speculative node-expansion policy, an overall smaller computational effort is required when compared to parallel branch-processing approaches.

The tree-search architecture under consideration uses different node expansion policies throughout the tree traversal:

1. *On-demand sibling node expansion*

   This node expansion approach is only performed at the root level of the tree, namely $i = N$. In the initial step of the algorithm, the best node according to the Schnorr-Euchner enumeration is selected and passed on to the next level. When the algorithm returns to the root of the tree, the next most favorable node is expanded only if its PED is smaller than R. Note that the ordered child node sequence at this level can be obtained without performing any PED calculations and subsequent sorting procedures.

2. *Single node expansion*

   A single child node is expanded per parent node in levels $i < N$, being a radius check performed at each one of these levels . This way, if a node with $D_i > R$ is reached, all its descendants are discarded and the algorithm returns to the root level. Otherwise, the tree traversal is resumed. When a leaf node with $D_1 < R$ is found, the branch

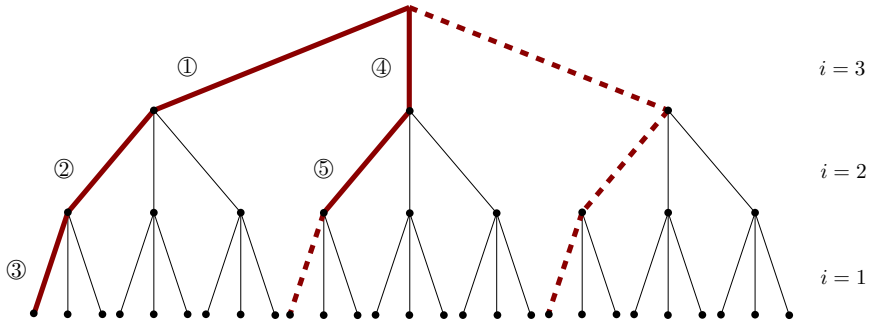

Figure 7.1: Example of a SBE tree search in a 3 user system with $|\mathcal{L}| = 3$.

leading to that node is stored as a candidate solution and the search radius is updated accordingly.

The sequential nature and the radius reduction procedure allow for an efficient pruning of undesired branches. Additionally, the single node expansion policy provides a more regular data path than that of the SE and removes the need for tracking candidate sibling nodes throughout the tree traversal.

An example of a SBE tree search is shown in Figure 7.1 for a 3 user system with a restricted search set of $|\mathcal{L}| = 3$ elements. The branches discarded due to the single node expansion policy are represented by the thinner lines, whereas the branches that have been pruned due to the radius reduction procedure are illustrated in dashed lines. Finally, the straight thick lines depict the tree traversal performed by the SBE algorithm.

7.2.1.1 Implementation Considerations

Following the inherent sequential nature of the SBE algorithm, the implementation of the tree search can be carried out by concatenating N single-branch DPUs and iterating through the architecture as necessary. This implementation approach derives in a modest hardware usage and a more efficient pruning, but also yields a reduced and variable throughput. Nevertheless, the former disadvantage can be alleviated by running several DPU instances in parallel, hence processing various vector symbols at the same time.

The data processing throughput of the SBE algorithm can be enhanced by relaxing the premise of an strictly sequential tree traversal. This way, the architecture of the SBE tree can be parallelized to some degree by employing several single-branch DPUs to simultaneously process various branches within the same search tree. This approach results in a higher data processing throughput derived from an earlier shrinking of the sphere radius, but also involves a less computationally-efficient tree search.

$\boldsymbol{\Psi} = \left(\boldsymbol{H}\boldsymbol{H}^H + \xi \boldsymbol{I}_N\right)^{-1}$
$\boldsymbol{\Pi} = \boldsymbol{I}_N$
for $i = 1, \ldots, N$
    if $i = N-1$
        $q = \underset{\hat{q} \in \{i,\ldots,N\}}{\operatorname{argmax}} \boldsymbol{\Psi}(\hat{q}, \hat{q})$
    else
        $q = \underset{\hat{q} \in \{i,\ldots,N\}}{\operatorname{argmin}} \boldsymbol{\Psi}(\hat{q}, \hat{q})$
    end
    $\boldsymbol{\Xi} = \tilde{\boldsymbol{I}}_N^{(i,q)}$
    $\boldsymbol{\Pi} = \boldsymbol{\Xi}\boldsymbol{\Pi}$
    $\boldsymbol{\Psi} = \boldsymbol{\Xi}\boldsymbol{\Psi}\boldsymbol{\Xi}^T$
    $\boldsymbol{D}(i,i) = \boldsymbol{\Psi}(i,i)$
    $\boldsymbol{\Psi}(i:N,i) = \boldsymbol{\Psi}(i:N,i)/\boldsymbol{\Psi}(i,i)$
    $\boldsymbol{\Psi}(i+1:N,i+1:N) = \boldsymbol{\Psi}(i+1:N,i+1:N)$
        $-\boldsymbol{\Psi}(i+1:N,i)\boldsymbol{\Psi}(i+1:N,i)^H\boldsymbol{D}(i,i)$
end
$\boldsymbol{L}$ = lower triangular part of $\boldsymbol{\Psi}$
$\boldsymbol{U} = \boldsymbol{D}^{1/2}\boldsymbol{L}^H$

Table 7.1: Computation of the upper-triangular matrix $\boldsymbol{U}$ with $\mathfrak{O}_{\text{SBE}}$ ordering.

7.2.2 Matrix Preprocessing for the SBE

The rearrangement of the user streams in the search tree has a great impact on the amount of nodes that are visited during an iterative tree traversal. In systems where there is some sort of limitation on the number of nodes to be processed, such as the K-Best or the FSE, the procedure of user ordering solely affects the error-rate performance of the algorithm. Due to the special pruning strategy performed in the SBE algorithm, the rearrangement of the user streams will impact both its performance and run-time.

As is the case with the FSE, the first evaluated branch during the SBE tree search, which is associated with the initial sphere constraint, corresponds to the solution of the successive precoding procedure performed by THP. Consequently, the ordering strategy used in this type of systems will be the starting point for the design of the SBE matrix preprocessing approach.

When generating the upper-triangular matrix following a *best-last* ordering strategy, higher values of the $u_{N,N}$ element are obtained when compared to the unordered matrix decomposition due to the worst stream being precoded at the last stage. This results in higher PEDs of the nodes at the root level, which in turn derives in very few nodes satisfying the radius constraint ($d_i < R$). Therefore, most of the root nodes are pruned and the SBE tree search is finished after just a few iterations. Certainly, a reduced run-time is a

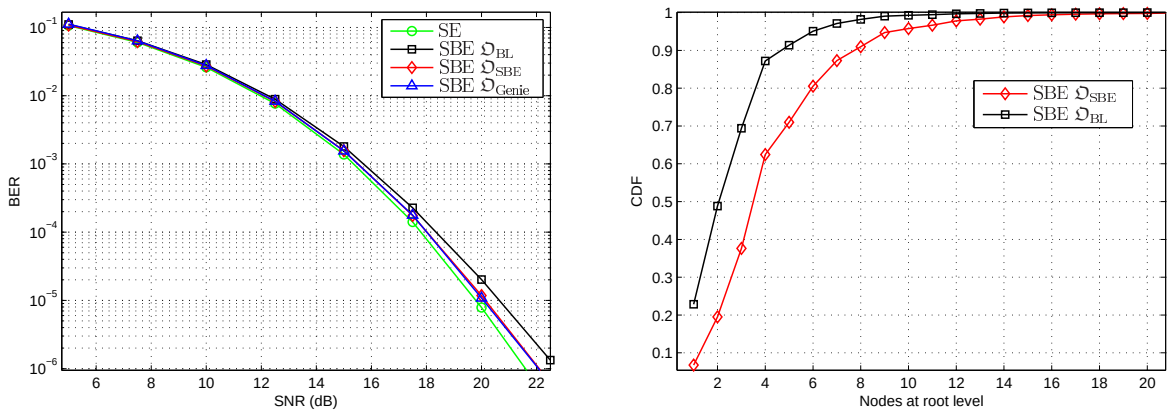

(a) (b)

Figure 7.2: BER performance of the SBE with different ordering strategies (a) and the CDF of the nodes expanded at the root level during a SBE tree traversal (b).

beneficial property for any sequential algorithm. However, by intentionally increasing the amount of evaluated nodes for the first precoded stream, it is possible to find more suitable perturbation elements for the rest of the streams.

With this statement in mind, a new ordering approach is proposed to be used along with the SBE tree traversal ($\mathfrak{O}_{\mathrm{SBE}}$). The novel ordering strategy consists of following a *best-last* approach when determining the user order in levels $i < N-1$. Therefore, the ordering procedure starts by selecting the best data stream at the bottom level of the tree, namely $i = 1$, and proceeds equally through the rest of the levels until $i = N-2$. Once the second level of the tree is reached ($i = N-1$), the worst stream is selected instead. By shifting the position of the worst user from the top of the tree to the next level, smaller PEDs associated with the root nodes can be attained since $u_{N,N}^{\mathfrak{O}_{\mathrm{BL}}} > u_{N,N}^{\mathfrak{O}_{\mathrm{SBE}}}$. This allows for the SBE to perform a more lenient pruning at the top level of the tree, which leads to a slightly longer run-time of the algorithm along with a better error-rate performance.

The pseudocode for the joint user order determination and triangular matrix computation of the $\mathfrak{O}_{\mathrm{SBE}}$ strategy, shown in Table 7.1, has been obtained by appropriately modifying the algorithm description given in Table 4.1 for the $\mathfrak{O}_{\mathrm{BL}}$ approach.

7.2.3 Performance Results

Figure 7.2(a) depicts the BER performance of the SE and the SBE algorithms for a 4×4 antenna system with 16-QAM modulation. More specifically, the error-rate performance of the SBE tree search with the presented $\mathfrak{O}_{\mathrm{BL}}$ and $\mathfrak{O}_{\mathrm{SBE}}$ ordering strategies is depicted. Additionally, the BER performance of the SBE algorithm with the brute force *genie* ordering $\mathfrak{O}_{\mathrm{Genie}}$ is shown. The latter consists of ordering the user data streams according to all possible $N!$ permutations for each channel realization and performing the precoding for each considered

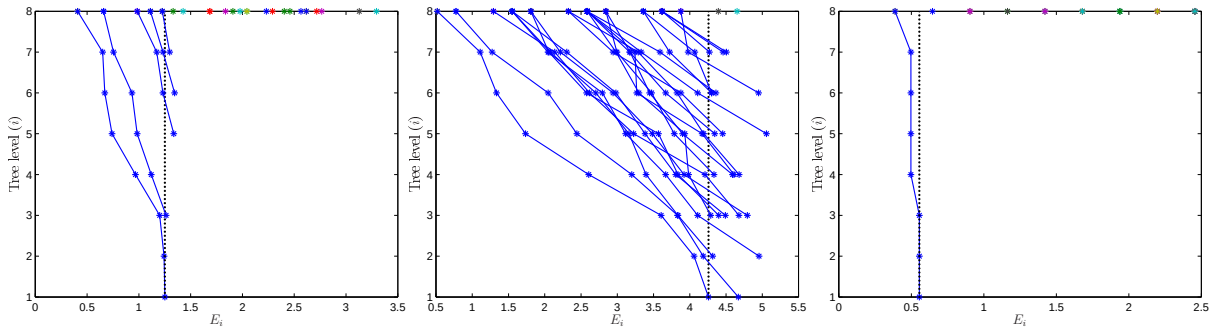

(a) ℓ^2 (b) $\ell^{\tilde{1}}$ (c) $\ell^{\tilde{\infty}}$

Figure 7.3: Accumulated distances of the nodes expanded during a SBE tree traversal with the ℓ^2 norm (a) and with the approximated $\ell^{\tilde{1}}$ (b) and $\ell^{\tilde{\infty}}$ (c) norms.

ordering tuple. The $\mathfrak{O}_{\text{Genie}}$ user order for that particular channel realization is then selected as the ordering sequence that yields the best performance. This ordering approach therefore establishes the upper bound on the error-rate performance of the SBE tree-search structure. The data displayed in this figure shows that the proposed ordering approach clearly outperforms the more traditional $\mathfrak{O}_{\text{BL}}$ ordering. What is more, the proposed $\mathfrak{O}_{\text{SBE}}$ approach yields virtually the same BER performance as the brute force $\mathfrak{O}_{\text{Genie}}$ strategy when applied to the SBE structure. Note also that the error-rate performance of the presented SBE approach is close to the optimum set by the SE algorithm.

As an additional result, the cumulative distribution function (CDF) of the amount of node expansions at the root level of the tree search is depicted in Figure 7.2(b). The longer run-time of the SBE algorithm caused by rearranging the users according to the proposed ordering approach can be observed in this figure.

7.3 The Low-complexity SBE Algorithm

Despite its already simple architecture, the features of the SBE algorithm can be further improved in terms of required hardware resources, delay and run-time. The proposed low-complexity SBE (LC-SBE) aims in this direction by applying a joint strategy of approximate norm computations and flexible run-time constraints. The implementation of the more computationally-efficient LC-SBE yields an insignificant performance degradation with respect to the original algorithm considering the great hardware resource saving and the additional complexity reduction that it is achieved.

7.3.1 SBE with an Approximate Norm

As already discussed in Section 6.5.2, a significant portion of the hardware resources in the implementation of any tree-search algorithm is dedicated to computing the PEDs related to

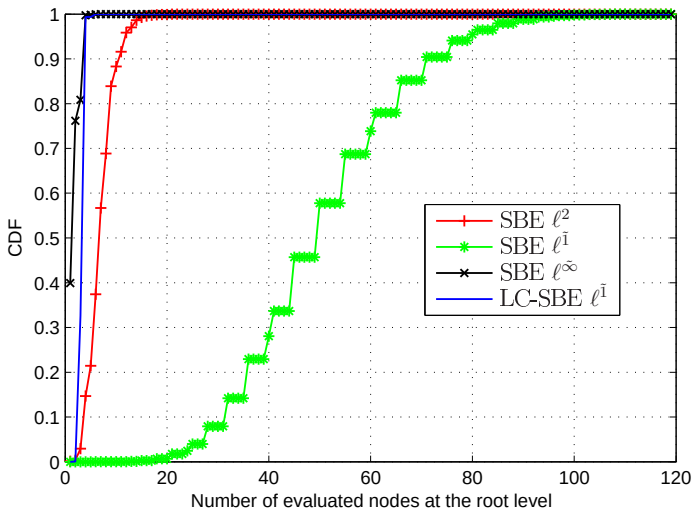

Figure 7.4: CDF of the nodes evaluated at the root level of a SBE and LC-SBE tree traversal in an 8 antenna system.

the nodes in the search tree. Additionally, in iterative systems such as the proposed SBE, the long delays associated with squaring operations required to compute the ℓ^2 norms greatly contribute to the critical path of the hardware architecture. The use of an approximate norm alleviates these problems, but it also results in a slight error-rate performance degradation.

With the introduction of the approximate norms the iteration complexity of the SBE is also altered. The effect on the pruning behavior is illustrated in Figure 7.3, where the AEDs computed during an 8 level SBE tree traversal are shown for different norm implementations. The image depicts the PEDs of all the nodes in the first level (e_N) and the accumulated distances of the branches expanded during the tree traversal in the rest of the levels. The vertical dotted lines represent the initial pruning radius for the different norm models, namely $R = E_1$. For the particular case illustrated in this example, the solution vector given by the SBE algorithm is the same regardless of the norm in use. Nevertheless, the resulting run-time of the SBE varies significantly depending on the implemented distance computation model.

By accumulating the distance increments following the $\ell^{\tilde{1}}$ approach described in (6.2) and (6.3), the square root operation required by the MNA ℓ^2 model is omitted for the sake of computational complexity, namely $E_i = \sqrt{E_{i+1}^2 + e_i^2} \approx E_{i+1} + e_i$. This results in a fast escalation of the AED values in the $\ell^{\tilde{1}}$ tree search, as is shown in Figure 7.3(b). As a consequence to this, many root nodes fulfil $e_N \leq E_1 = R$ and therefore, a very lenient pruning and an extremely long run-time are experienced by the $\ell^{\tilde{1}}$-SBE. The problem of the increased number of iterations in the $\ell^{\tilde{1}}$ variant is aggravated as the tree grows in dimensions, i.e. more users are added to the system.

The behavior of the SBE with the ℓ^{∞} norm is the opposite, since the max operator

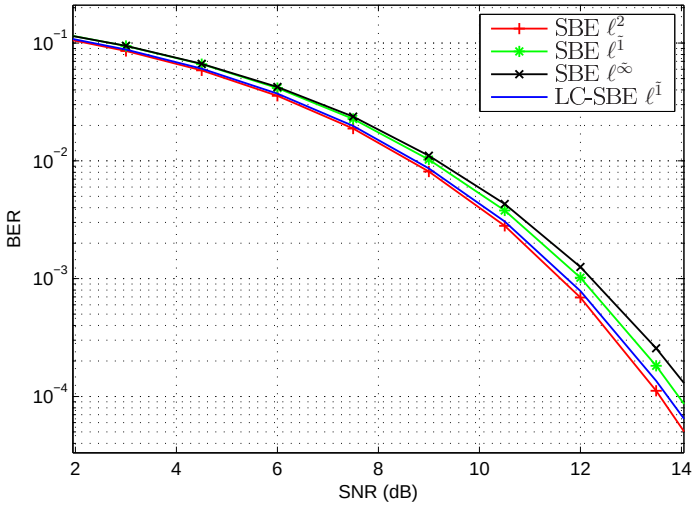

Figure 7.5: BER performance of the LC-SBE and the SBE algorithm with the ℓ^2,$\ell^{\tilde{1}}$ and ℓ^{∞} approaches for distance computation.

in (6.4) causes little variations in the value of the AEDs as the tree is traversed. The combination of this feature with the preferred ordering strategy for the SBE, which derives in the concentration of the largest PEDs at the top levels of the tree, results in a very early pruning of the root nodes, as is shown in Figure 7.3(c).

On a more general scope, the variations on the run-time of the SBE caused by the use of an approximate norm will be further analyzed by means of the CDF of the amount of evaluated root nodes. The data displayed in Figure 7.4 support the conclusions drawn from the illustrative example in Figure 7.3, such as the considerably less aggressive pruning of the $\ell^{\tilde{1}}$-SBE algorithm when compared to the other addressed SBE variants, or the short run-time of the ℓ^{∞}-SBE. To this respect, the statistical analysis on the expanded nodes at the top level concludes that the aggressive pruning strategy performed when applying an ℓ^{∞} distance computation model derives in a single branch computation being performed in roughly 40% of the tree-search instances.

Clearly, the adoption of an approximated norm approach introduces a certain error-rate performance degradation, as is shown in Figure 7.5. The BER performance of the SBE algorithm with the addressed approximate norms in an 8×8 system with 16-QAM modulation is shown in Figure 7.5. Despite having a similar complexity in terms of distance computation effort, the $\ell^{\tilde{1}}$ variant of the simplified-norm SBE achieves a better error-rate performance than the ℓ^{∞} model, mainly due to the more aggressive pruning process performed in the latter. Nevertheless, both approximate distance computation models preserve the diversity order of the ℓ^2-SBE tree search.

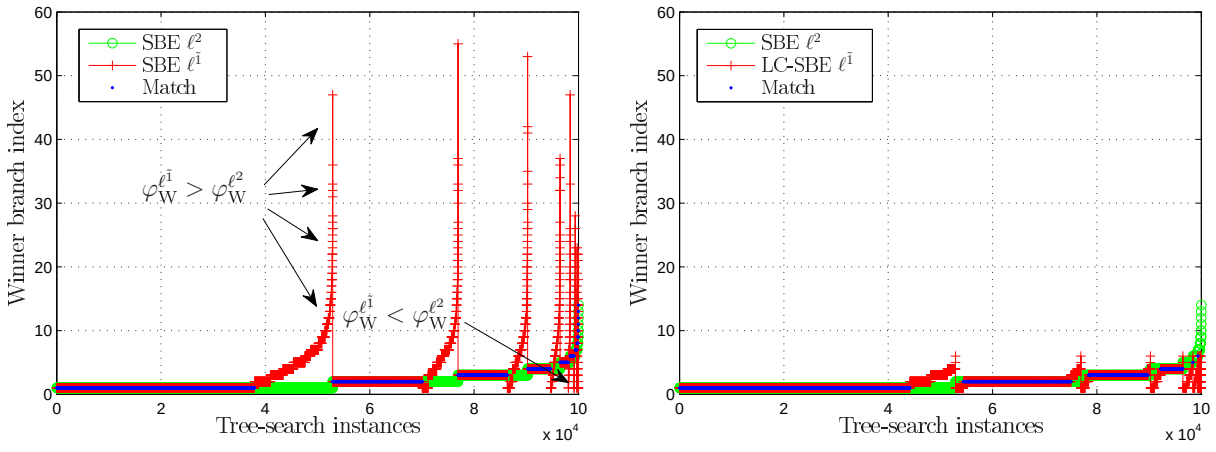

Figure 7.6: Index of the branch selected as the solution vector for the $\ell^{\tilde{1}}$ and ℓ^2-SBE (a), and for the ℓ^2-SBE and the proposed LC-SBE (b).

7.3.2 Flexible Run-time Constraint

The implications of using a low-complexity distance computation model for the SBE tree search were presented in the previous section. As already discussed, the $\ell^{\tilde{1}}$ approximate norm is preferred from a BER performance point of view, whereas the ℓ^{∞} norm is considerably more attractive in terms of run-time of the SBE algorithm. In this section, a low-complexity tree-search algorithm that aims at closing the performance gap with respect to the (ℓ^2-)SBE will be presented. The proposed scheme combines two advantageous features, namely, the close-to-optimum performance of the $\ell^{\tilde{1}}$-SBE and the brief run-time of the ℓ^{∞}-SBE.

To that end, an analysis on the $\ell^{\tilde{1}}$ tree traversal will be performed first, identifying the main causes for the performance degradation suffered when replacing the optimum ℓ^2 distance computation model by this approximate norm. Next, a novel variable constraint algorithm will be presented, which will greatly improve the long run-time issues of the $\ell^{\tilde{1}}$-SBE, while at the same time considerably enhancing its performance. The resulting tree-search algorithm will be denoted as the LC-SBE.

7.3.2.1 Performance Degradation in the $\ell^{\tilde{1}}$-SBE

When approximating the distance computation in (2.8) by a simplified model, the metrics associated with all the branches in the SBE tree are altered, which may lead to selecting a suboptimum branch within the SBE tree. Note that, as already discussed, the nodes at the root level can be sequenced without computing their associated distances and therefore, the node expansion at this level can be performed equally for both ℓ^2 and $\ell^{\tilde{1}}$ variants. Furthermore, given that only the most favorable node is expanded in the subsequent levels,

which results in the same selected child node regardless of the norm in use, we conclude that the set of eligible nodes in the structure of the SBE tree is not modified with the implementation of an approximate norm. Hence, the φ^{th} branch in the SBE tree will represent the same candidate solution vector regardless of the approach used to compute the distance increments.

We shall refer to the branch selected as the optimum within the ℓ^x-SBE tree search as the *winner* branch, whose index will be denoted as $\varphi_{\text{W}}^{\ell^x}$. Due to the metric discrepancies between the ℓ^2 and $\ell^{\tilde{1}}$ models, the perturbation vector selected by the SBE tree search may differ in some cases, i.e. $\varphi_{\text{W}}^{\ell^{\tilde{1}}} \neq \varphi_{\text{W}}^{\ell^2}$. These events, which will be referred to as *false positive* events, represent the only cause for the performance degradation suffered when replacing the optimum ℓ^2 distance computation model by an approximate norm in the SBE tree search.

With the aim of further analyzing the rationale behind the *false positive* events, the variations in the tree traversal of the SBE with different distance computation models have been studied for 100000 channel and symbol vector combinations. For each analyzed search tree, the *winner* branches $\varphi_{\text{W}}^{\ell^2}$ and $\varphi_{\text{W}}^{\ell^{\tilde{1}}}$ have been stored. The data collected for this study, which are displayed in Figure 7.6(a), have been properly arranged by increasing order of $\varphi_{\text{W}}^{\ell^2}$ and $\varphi_{\text{W}}^{\ell^{\tilde{1}}}$ to ease the readability of the results. As is shown in this figure, the great majority of the *false positives* are caused by $\varphi_{\text{W}}^{\ell^{\tilde{1}}} > \varphi_{\text{W}}^{\ell^2}$ events, while only a small part of the errors are due to situations in which $\varphi_{\text{W}}^{\ell^{\tilde{1}}} < \varphi_{\text{W}}^{\ell^2}$. The former and most unfavorable events are triggered by the extended run-time of the $\ell^{\tilde{1}}$-SBE. This is due to the fact that in these occasions the optimum $\varphi_{\text{W}}^{\ell^2}$ branch is visited at an early stage of the $\ell^{\tilde{1}}$-SBE tree traversal and is stored as a candidate solution. Nevertheless, given the nature of the $\ell^{\tilde{1}}$ distance distribution, a great amount of iterations through the SBE tree are performed in the subsequent stages of the algorithm until eventually, a *false positive* branch with a smaller $\ell^{\tilde{1}}$ metric is found.

7.3.2.2 Flexible Run-time Constraint for the $\ell^{\tilde{1}}$-SBE

The problem of the wrong branch selection caused by $\varphi_{\text{W}}^{\ell^{\tilde{1}}} > \varphi_{\text{W}}^{\ell^2}$ events can be overcome by setting an appropriate limit to the amount of iterations to be performed during the $\ell^{\tilde{1}}$-SBE tree traversal. This way, the probability of choosing a suboptimum branch in later iterations through the SBE tree search is reduced. Put in other words, the implementation of a limit in the number of iterations *forces* the $\ell^{\tilde{1}}$-SBE to keep an early correct decision in many occasions, consequently improving the overall system performance. The limit on the maximum allowable number of branch expansions (φ_{TH}) has to be set carefully. If the threshold is set too low, the strict constraint may prevent the evaluation of the optimum branch during the tree traversal, which will result in the selection of an erroneous perturbation vector. However, if the threshold is too lenient, most of the *false positive* events will not be prevented.

In previous sections, the effect of the $u_{N,N}$ parameter on the run-time of the SBE has

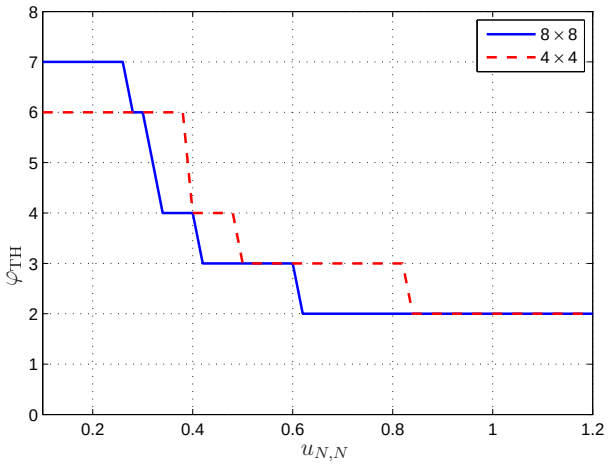

Figure 7.7: Run-time constraint values (φ_{TH}) for 4×4 and 8×8 antenna setups with $p = 0.9$.

been discussed. Clearly, those channel realizations that yield small values of $u_{N,N}$, and hence perform an increased amount of iterations through the search tree, will allow for higher values of $\varphi_{\mathrm{W}}^{\ell^2}$, and viceversa. With this statement in mind, the implementation of a flexible threshold system is proposed, where the pruning constraint of the tree search is updated depending on the particular structure of the triangular matrix $\boldsymbol{U}$.

Hence, the value of φ_{TH} will be selected based on previously gathered statistics on the value of $\varphi_{\mathrm{W}}^{\ell^2}$ subjected to $u_{N,N}$. The gathered data is discretized by dividing the range of captured $u_{N,N}$ values in several intervals. The run-time constraint for a given value of $u_{N,N}$ is then set by analyzing the captured $\varphi_{\mathrm{W}}^{\ell^2}$ values for the corresponding interval and evaluating the amount of expanded root nodes that are necessary to find the optimum solution with a probability of p. The value of p will determine the leniency of the constraint for each discrete range of $u_{N,N}$ values. Usually, a value of $p \in [0.9, 0.99]$ provides good performance results, with the values in the lower range enabling shorter run-times of the algorithm.

The data displayed in Figure 7.6(b) show the index of the $\ell^{\tilde{\mathrm{I}}}$-SBE *winner* branch after the implementation of the proposed flexible threshold approach ($\varphi_{\mathrm{W}}^{\ell^{\tilde{\mathrm{I}}},\mathrm{LC}}$). As one can notice, most of the *false positive* events with $\varphi_{\mathrm{W}}^{\ell^{\tilde{\mathrm{I}}}} \gg \varphi_{\mathrm{W}}^{\ell^2}$ have been corrected. In some cases, the constraint is not restrictive enough and the selected branch in the LC-SBE tree traversal is not the optimum SBE solution, namely $\varphi_{\mathrm{TH}} > \varphi_{\mathrm{W}}^{\ell^{\tilde{\mathrm{I}}},\mathrm{LC}} > \varphi_{\mathrm{W}}^{\ell^2}$. However, even in this cases, a certain performance enhancement is introduced as the index of the suboptimum branch is significantly smaller after the implementation of the flexible threshold, that is $\varphi_{\mathrm{W}}^{\ell^{\tilde{\mathrm{I}}},\mathrm{LC}} \ll \varphi_{\mathrm{W}}^{\ell^{\tilde{\mathrm{I}}}}$. This implies that the new approach favors the selection of those perturbation vectors that are closer to the origin, and hence have a smaller branch index. Due to their proximity to the origin of the lattice, they enable the reduction of the precoded signal power, which reflects positively on the final system performance.

7.3.3 Performance Results

The BER performance of the SBE-based algorithms is shown in Figure 7.5 for an 8 single-antenna user system. The data displayed in this figure show that the *false positive* correction procedure performed on the $\ell^{\tilde{1}}$-SBE tree search bridges the performance gap between the optimum SBE structure (ℓ^2-SBE) and the simplified distance computation model.

Moreover, Figure 7.4 depicts the CDF of the amount of expanded branches for the LC-SBE along with the previously discussed (ℓ^2-)SBE, $\ell^{\tilde{1}}$-SBE and $\ell^{\tilde{\infty}}$-SBE. As one can notice, the run-time of the SBE tree search based on $\ell^{\tilde{1}}$ distances is reduced drastically after the implementation of the variable threshold approach, yielding an iteration effort that is even comparable to that of the computationally-attractive $\ell^{\tilde{\infty}}$-SBE.

7.3.4 Implementation Considerations

The steepness of the LC-SBE CDF curve in Figure 7.4 indicates a low variability on the amount of expanded nodes at the root level. This fact is beneficial if a semi-parallel branch computation approach is considered for hardware implementation. In such a case, the benefit with respect to the fixed-complexity K-Best and FSE algorithms is clear, as the amount of branches computed by the LC-SBE is in general considerably smaller than the amount required by the fixed-complexity approaches to achieve a close-to-optimum performance.

This way, it can be shown that several semi-parallel LC-SBE tree-search instances can achieve a higher average throughput than the addressed fixed-complexity algorithms for the same amount of allocated hardware resources. From the data displayed in Figure 7.4, it can be extracted that the LC-SBE traverses 4 tree branches or less with a 99 % probability. Consider a fixed-complexity tree-search technique that evaluates the metrics of L candidate branches, where L is assumed to be a multiple of 4 for ease of exposition. Additionally, consider that $L/4$ LC-SBE tree-search instances are implemented, each one of them processing 4 tree branches in parallel. In this case, the run-time constraint φ_{TH} will determine which computed branches are considered as candidates whenever $\varphi_{\mathrm{TH}} < 4$, and will indicate the occasions in which the 4 branch processing structures should be traversed for a second time, i.e $\varphi_{\mathrm{TH}} > 4$. Hence, despite the parallel processing of the branches, the LC-SBE implementation remains of variable nature. This way, both the fixed-complexity and variable LC-SBE function on a L-branch computation hardware architecture. Nevertheless, the throughput the of the LC-SBE model is on average $L/4$ times higher.

7.4 Simulation Results

A comparative analysis on the error-rate performance and complexity of the proposed SBE and LC-SBE algorithms with respect to the previously addressed tree-search techniques will

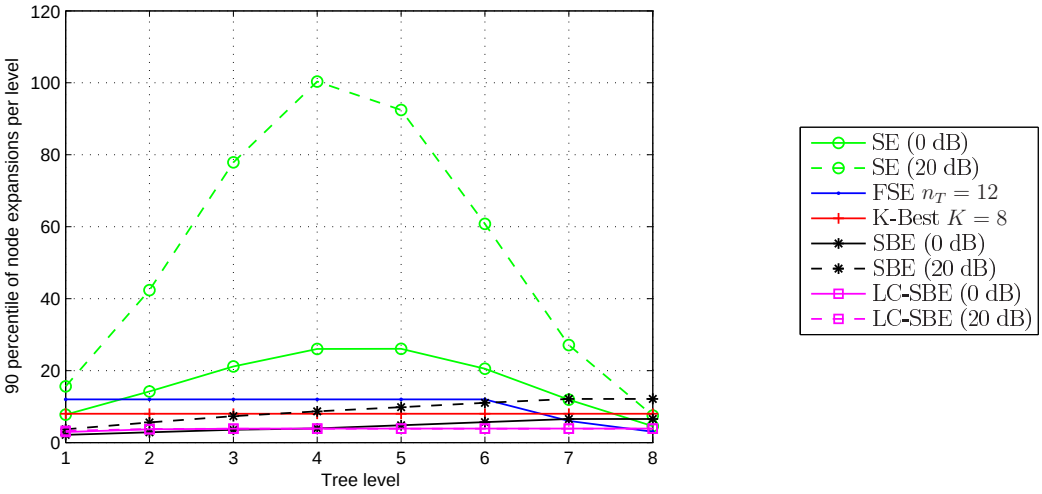

Figure 7.8: 90-percentile of the number of node expansions per level for the SE, FSE, K-Best, SBE and LC-SBE tree-search algorithms in an 8×8 antenna system.

be provided in this section. The flexible threshold values that have been used for the current study are depicted in Figure 7.7.

7.4.1 Computational Complexity

The computational complexity of the proposed SBE and LC-SBE algorithms will be assessed in terms of evaluated nodes and total operation count with the aim of extending the complexity analysis of the SE, K-Best and FSE approaches already presented in Section 4.4.

7.4.1.1 Number of Evaluated Nodes

The 90-percentile of the amount of node expansions per level for the variable complexity schemes, i.e. the SE, SBE and LC-SBE, is shown in Figure 7.8, whereas a more general scope on the SNR-dependency of the iterative algorithms is given in Figure 7.9. Additionally, the amount of visited nodes for the addressed fixed-complexity algorithms, namely the K-Best and FSE, is also depicted for completion.

The notably smaller complexity in terms of evaluated nodes of the SBE-based algorithms when compared to the SE is a consequence of considering only the most favorable nodes in levels $i < N$ of the tree traversal. This implies that the number of expanded nodes at level i is always equal or inferior to the amount of nodes evaluated at level $i + 1$, as is shown in Figure 7.8. Consequently, the great escalation of node expansions experienced by the SE in the central part of the tree search, where the small AEDs derive in a great amount of nodes satisfying the sphere constraint, can be avoided. The single node expansion policy of the SBE-based schemes also reflects positively on the SNR-dependency of their iteration effort. Even if the (ℓ^2-)SBE also performs more node computations in the high SNR range due to

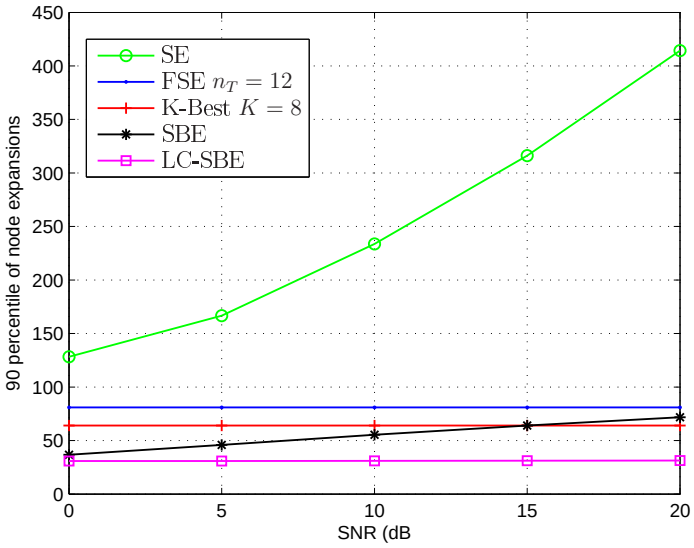

Figure 7.9: 90-percentile of node expansions for the SE, FSE, K-Best, SBE and LC-SBE in an 8×8 antenna system.

its variable nature, its growth rate is considerably smaller than that of the SE due to the additional pruning policy.

What is more, the limitation set by the LC-SBE on the number of expanded root nodes results not only in the elimination of the unfavorable *false positive* events, but additionally in a considerable reduction of the complexity of the algorithm in terms of amount of processed nodes. The data depicted in Figure 7.8 show that the proposed low-complexity scheme iterates through the search tree in a very efficient way, expanding a remarkably smaller amount of nodes than the optimum SE or the addressed fixed-complexity schemes. Additionally, the novel node expansion constraint implemented on the LC-SBE enables an iteration pattern with virtually no SNR dependency, as is shown in Figures 7.8 and 7.9. Actually, the SNR-imposed variation of the LC-SBE iteration pattern is so reduced that it is unnoticeable in the provided figure.

7.4.1.2 Number of Operations

The analysis of the total operation count for the SE, FSE and K-Best approaches performed in Section 4.4.1.2 will be extended to the proposed SBE-based algorithms. The data depicted in Figure 7.10 therefore represent the total amount of arithmetic (addition, substraction, multiplication and division) and logical (comparison, swapping, branching) operations performed during an $N = 8$ tree traversal.

As one can notice, there is a broad complexity gap between the optimum SE and the proposed sequential SBE and LC-SBE schemes. This data reflects the great computational complexity entailed in the management of the candidate lists at each level of the SE tree. When suppressing the possibility of sibling node expansions in levels $i < N$, the complexity

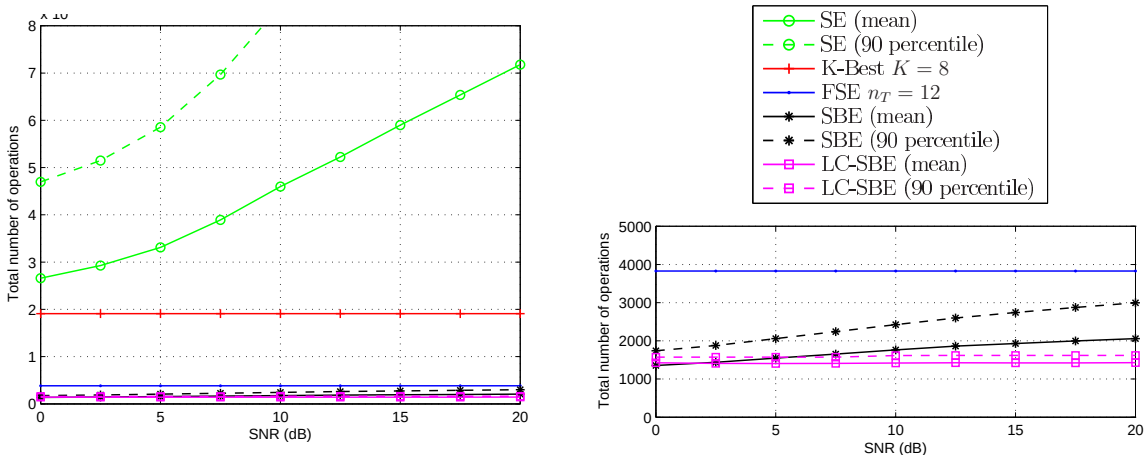

Figure 7.10: Total number of operations for the SE, K-Best, FSE, SBE and LC-SBE tree-search approaches on an $N = 8$ user system. A zoomed image of the lower section of the main figure is also included for a better appreciation of the complexity results of the SBE-based schemes.

of the algorithm in terms of number of operations and node expansions, as already discussed in the previous section, is greatly reduced. Moreover, the elimination of the candidate lists also works in favor of a more constant computational load, as is shown in the zoomed plot of Figure 7.10.

When compared to the addressed fixed-complexity approaches, the data in Figure 7.10 show that an important complexity reduction can be achieved by traversing the tree in both forward and backward directions in such a way that a sphere constraint can be implemented to get rid of unnecessary distance computations. This way, the proposed SBE and LC-SBE approaches yield a total operation count that is at least half of that required by the fixed-complexity FSE, and roughly 10 times smaller than that of the K-Best.

7.4.2 BER Performance

The data displayed in Figure 7.11 show the BER performance of the optimum SE and the proposed tree-search approaches for 4×4 and 8×8 antenna setups. It is noticeable that the SBE yields better error-rate performance results in systems where the number of tree-search levels is more reduced, e.g. $N = 4$. Nevertheless, the performance in the 8 antenna case is still close to the optimum, with just a 0.6 dB degradation in the high-SNR regime. As for its low-complexity counterpart, the performance degradation with respect to the SBE is negligible in both scenarios. This is mainly due to the fact that, by means of the variable threshold approach, the iteration pattern of the LC-SBE is *forced* to resemble that of the SBE.

Given the close-to-optimum error-rate performance and great complexity reduction of the proposed SBE techniques when compared to the optimum SE, it is concluded that there

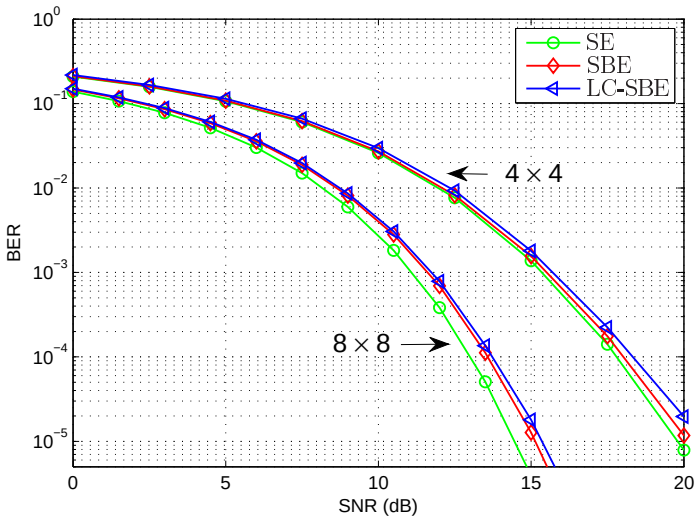

Figure 7.11: BER performance curves for the SE, SBE and LC-SBE tree-search approaches for 4×4 and 8×8 antenna setups.

is little to be gained by considering all sibling nodes for expansion in levels $i < N$ of the tree search.

7.4.2.1 BER Performance under an Overall Run-time Constraint

The iterative tree-search algorithms discussed in this section require a variable amount of cycles to find the perturbation vector. This variable nature usually constitutes a problem in practical systems, where the user data streams are supposed to be processed at a fixed rate. It is therefore important to assess the error-rate performance of the tree-search algorithms under an overall run-time constraint.

The data in Figure 7.12 reflect the error-rate performance of an 8×8 VP system where an overall run-time constraint of $\varpi = 32$ evaluated nodes has been imposed on the analyzed tree-search algorithms. Given the fixed-complexity of the K-Best and FSE approaches, the design parameters that yield a node expansion count of ϖ have been selected. Consequently, values of $K = 4$ and $n_T = 4$ have been considered for the K-Best and FSE tree-search approaches, respectively. Note that the optimum tree configuration vector that yields $n_T = 4$ in an $N = 8$ user system entails the computation of 30 PEDs. This is the closer amount of PED computations to $\varpi = 32$ that can be achieved with the FSE, as dictated by (4.4).

As is shown in Figure 7.12(a), the performance of the run-time-constrained algorithms that use the optimum ℓ^2 norm is similar, being the K-Best the one that performs slightly better in the low and mid-SNR ranges. Nevertheless, the K-Best shows its characteristic error-rate performance degradation in the high-SNR regime. Despite their suboptimum performance [see Figures 4.13 and 7.11], the additional pruning strategy carried out in the fixed-complexity and SBE-based systems favors the expansion of the most promising nodes

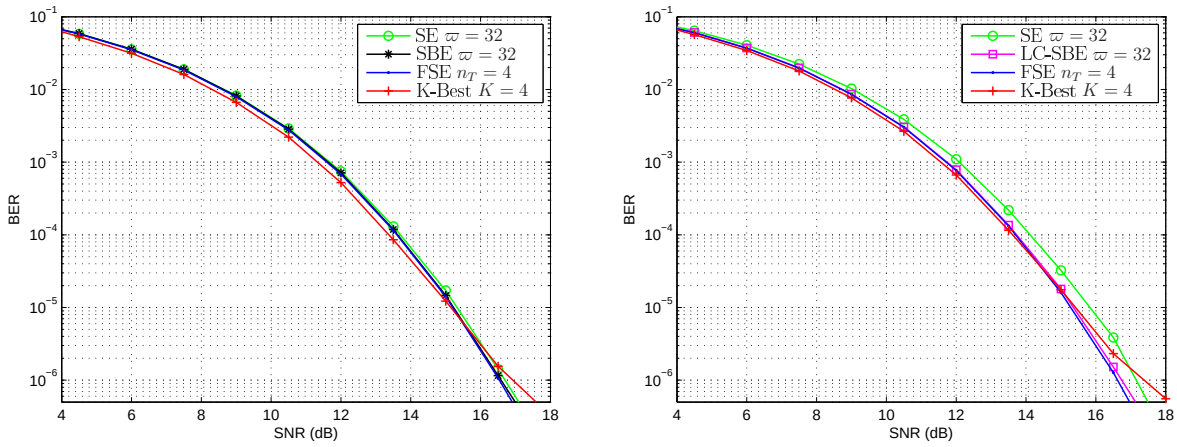

(a) ℓ^2 (b) $\ell^{\tilde{1}}$

Figure 7.12: BER performance of various tree-search approaches with an overall run-time constraint of $\varpi = 32$ evaluated nodes for an 8×8 antenna setup.

as opposed to the SE, which considers all eligible nodes at every level. As a consequence to this, a slightly higher BER performance degradation is introduced by the SE when imposing an overall run-time constraint.

As for the algorithms with an approximate norm, which are depicted in Figure 7.12(b), the same rationale applies. Nevertheless, the performance gap between the optimum SE and the suboptimum techniques is wider in this case. The reason for this is that the use of an approximate norm in the SE derives in an extended set of eligible nodes at every level, which causes a considerable performance degradation when applying an overall run-time constraint. The effect of the $\ell^{\tilde{1}}$-norm on the LC-SBE is more limited due to its variable threshold scheme, and therefore, the imposition of an additional run-time constraint does not significantly affect its performance, as the limitation on the amount of processed nodes is only effective for those channel realizations that yield $\varphi_{\mathrm{TH}} > 4$.

It is worth pointing out that the proposed LC-SBE scheme yields the same BER performance as the (ℓ^2-)SE when an overall run-time constraint is applied, even when the architecture of the former features an approximate distance computation model and a single-node expansion strategy.

7.5 Chapter Summary

The present chapter has dealt with the design of sequential and low-complexity tree-search algorithms for VP. The main motivation for the change of tree traversal strategy is the implementation of a sphere constraint that will allow for additional pruning. This way, a considerable reduction of the unnecessary distance computations performed in the fully-parallel scheme can be achieved.

The proposed SBE approach features, on one hand, a simplified tree-search architecture that only considers the most favorable nodes for expansion and, on the other hand, a specially tailored ordering strategy that provides a good trade-off between performance and computational complexity. Provided BER results have shown that the proposed approach achieves an error-rate performance that is close to the optimum set by the SE.

In addition to this, a low-complexity counterpart of the SBE algorithm has been presented. The proposed LC-SBE combines a simplified norm approach and a variable run-time constraint strategy to further reduce the complexity of the SBE in terms of amount of evaluated nodes and distance computation effort. The LC-SBE presents a negligible performance degradation when compared to the SBE considering the high complexity reduction that it is achieved.

Furthermore, when an overall run-time constraint is imposed to guarantee a fixed throughput, the proposed algorithms show a better error-rate performance than the SE. Moreover, despite featuring an approximate distance computation model and a single-node expansion tree search, the LC-SBE yields an error-rate performance that is similar to that of the SE with the optimum ℓ^2 norm.

Chapter 8
Summary and Conclusions

This chapter summarizes the work developed during this PhD thesis, along with its main contributions and associated publications. The main focus of this thesis has been on the circumvention of the analytical and computational difficulties posed by the perturbation process in VP systems. In this respect, the performance improvement obtained by adding a perturbation signal before the linear precoding stage entails two main drawbacks. On one hand, the computation of the perturbation vector is a non-trivial task of NP-hard complexity and, on the other hand, the analytical assessment of the performance of VP is hindered by the lack of knowledge of the statistical properties of the perturbed signal.

Chapter 3 has provided some insight into the main difficulties in the analytical assessment of the performance of VP systems. From an analytical perspective, the intricacies derived from the incorporation of non-linear algorithms are twofold: first, the AWGN at the receivers is filtered by a modulo operator, which causes the modification of the statistical properties of the effective noise at the detection stage. Second, the computation of the power of the precoded symbols involves the expectation over a closest-point lattice search. These issues have been overcome by introducing several lower and upper bounds on the power of the precoded symbols in the high-SNR range, where the effect of the modulo-filtered noise can be neglected. The user sum rate and three other applications have then been analyzed (weighted sum rate, QoS and user rate balancing) for linear precoding and VP. In the latter, the proposed bounds have been used to solve the optimization problems. The computation of the performance gap between the linear and vector precoding systems has provided an analytical proof of the superior performance of the non-linear scheme. As an additional result, the ergodic performance gap has also enabled the quantization of the gain to be expected from the incorporation of a perturbation process into a linear precoding system.

Chapters 4 and 7 have dealt with the design of tree-search architectures for an efficient computation of the perturbation vector. Initially, non-iterative structures that allow for a high-throughput implementation of vector precoders have been considered. In this respect, a fixed-complexity algorithm entitled the fixed-sphere encoder (FSE) has been proposed for the VP tree search. The lack of sorting stages and the reduced amount of expanded child nodes result in a low complexity of the proposed algorithm when compared to the optimum SE

and the non-recursive K-Best scheme. Furthermore, the optimum structure of the FSE tree search in terms of computational complexity and error-rate performance of the algorithm has been introduced. Despite its simple architecture, the FSE still achieves a close-to-optimum error-rate performance.

The excessive distance computations performed by the pruning criteria of non-iterative algorithms is the main motivation for the study of an alternative tree traversal method for VP systems. The iterative SBE algorithm presented in Chapter 7 features a single-node expansion structure and the implementation of a sphere constraint for additional pruning. The former characteristic of the proposed SBE avoids the computationally expensive task of sibling node management performed during the SE tree traversal, and therefore allows for a more regular data path of the proposed scheme. This particular tree structure, along with the possibility of implementing a sphere constraint, results in an efficient tree traversal approach in terms of amount of expanded nodes, allocated hardware resources and power consumption.

Additionally, a low-complexity counterpart of the SBE algorithm has been presented. The proposed LC-SBE combines a simplified norm calculation approach and a variable run-time constraint strategy to further reduce the complexity of the SBE in terms of amount of evaluated nodes and distance computation effort. The resulting scheme is specially attractive for hardware implementation due to its single node expansion policy and modest resource occupation derived from the lack of squaring operations in the distance computation process.

An interesting insight into the hardware implementation of tree-search algorithms for precoding scenarios has been provided in Chapters 5 and 6. As an initial step, the pivotal process of determining the ordered sequence of complex-valued child nodes has been analyzed. With the aim of minimizing the computational complexity of this task, a novel non-sequential algorithm has been presented. When compared to other state-of-the-art complex enumeration algorithms, the proposed puzzle enumerator has shown a remarkably more reduced hardware resource occupation and latency. Due to these properties, the presented enumeration unit is highly suitable for its incorporation into the hardware architecture of recursive or non-recursive tree-search algorithms.

Furthermore, the proposed enumeration scheme has been included in the fully-pipelined hardware implementation of the non-recursive FSE and K-Best algorithms presented in Chapter 6. Despite the fact that the implemented hardware architectures account for a similar amount of allocated embedded multipliers, the more complex node selection procedure in the K-Best algorithm results in a higher FPGA slice occupation for this tree-search scheme. Due to the good performance, occupation results and simplicity of implementation, it is concluded that the FSE is best suited for a fully-pipelined practical implementation.

8.1 Thesis Contributions

This section outlines the main contributions of the presented research work and the associated scientific publications.

- Analysis and computation of upper and lower bounds for the sum rate in VP systems. The results of this study have been published in [Barrenechea10c] and partially in [Barrenechea10b].
- Optimization problem resolution and delimitation of the performance of VP system with respect to the applications of quality of service, weighted sum rate and user rate balancing [Barrenechea10b].
- Fixed-complexity tree-search algorithm design for the computation of the perturbation vector in VP systems. The results for the Wiener filter variant have been published in [Barrenechea09b] and [Barrenechea09a], while the regularized inversion alternative has been presented in [Barrenechea09c].
- Hardware implementation of fixed-complexity and high-throughput algorithms for precoding scenarios. The results and conclusions of the aforementioned contribution have been published in [Barrenechea10a] and [Barrenechea11b].
- Design and hardware implementation of a non-sequential and low complexity complex-plane enumerator for tree-search-based precoding systems [Barrenechea11d].
- Design and hardware implementation of a distributed sorting and high-throughput K-Best tree-search architecture for the computation of the perturbation vector in VP systems. The results of this research work have been presented in [Barrenechea11a].
- Design of iterative tree-search algorithms for a low-complexity implementation of a vector precoder. The publication associated with the aforementioned research work is pending for review [Barrenechea11c].

8.2 Suggestions for Further Research

The research work presented in this PhD dissertation can serve as the foundation for future work in the area of non-linear precoding. These are some of the directions for a possible extension of the presented results:

- Analysis of the performance gain of VP with respect to linear precoding in the low and mid-SNR ranges. The study of this scenario requires the incorporation of the

modulo-filtered noise parameter Ω (3.13) into the optimization problems presented in Chapter 3.

- Extension of the analysis on the performance of VP-ZF to other precoding structures such as VP-WF. In this case, the effect of the noise variance will also affect the generator matrix of the lattice, and therefore, the lattice parameters used to delimit the performance of VP systems.

- Assessment of the performance of the proposed tree-search algorithms, namely the FSE, SBE and LC-SBE, in a more realistic scenario. The adverse conditions to be considered include imperfections in the channel state information or interference generated from users in an adjacent cell. Moreover, it would be of great interest to analyze the performance of the aforementioned tree-search algorithms with real measured channel coefficients instead of the widely-used Rayleigh fading model.

- Hardware implementation of the proposed iterative algorithms, namely the SBE and LC-SBE. This way, a comparative analysis on the hardware occupation results and average throughput of the non-iterative (FSE and K-Best) and iterative schemes could be provided.

Appendix A
Publications

The research work that has been carried out during the development of this thesis has been published in the refereed conference papers and journal articles listed below.

Journal papers:

- M. Barrenechea, L. Barbero, M. Mendicute and J. Thompson, "Design and implementation of a low-complexity multiuser vector precoder", *International Journal on Embedded and Real-Time Communication Systems*(accepted).

- I. Sobrón, M. Barrenechea, P. Ochandiano, L. Martínez, M. Mendicute and J. Altuna, "Low-complexity detection of space-frequency block codes in LDPC-based OFDM systems", (accepted) *IEEE Transactions on Communications (Transaction Letter).*

- M. Barrenechea, A. Burg and M. Mendicute, "Reduced complexity tree-search-based vector precoding for multiuser systems", (submitted to) *IEEE Transactions on Communications.*

International conference papers:

- M. Barrenechea, M. Mendicute, J. Del Ser, and J. S. Thompson, "Wiener filter-based fixed- complexity vector precoding for the MIMO downlink channel", in *Proceedings of the 10th IEEE Signal Processing Workshop on Signal Processing Advances in Wireless Communications (SPAWC'09)*, pp. 216-220, June 2009.

- M. Barrenechea, J. S. Thompson, M. Mendicute, and J. Del Ser, "Fixed-complexity regularized vector precoding for the multiuser MIMO downlink channel", in *Proceedings of the 17th EURASIP European Signal Processing Conference (EUSIPCO'09)*, pp. 2406-2410, August 2009.

- M. Barrenechea, L. Barbero, M. Mendicute, and J. S. Thompson, "Design and hardware implementation of a low-complexity multiuser vector precoder", in *Proceedings of the Conference on Design and Architectures for Signal and Image Processing (DASIP'10)*, pp. 160-167, October 2010.

- M. Barrenechea, M. Joham, M. Mendicute, and W. Utschick, "Analysis of vector precoding at high SNR: Rate bounds and ergodic results", in *Proceedings of the IEEE Global Communications Conference (GLOBECOM'10)*, pp. 1-5, December 2010.

- M. Barrenechea, L. Barbero, I. Jiménez, E. Arruti, and M. Mendicute, "High-throughput implementation of tree-search algorithms for Vector Precoding", in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'11)*, pp. 1689-1692, May 2011.

- I. Sobrón, M. Barrenechea, P. Ochandiano, L. Martínez, M. Mendicute and J. Altuna, "Low-Complexity detection of golden codes in LDPC-coded OFDM systems", in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP '11)*, pp. 3160-3163, May 2011.

- M. Barrenechea, M. Mendicute, I. Jimenez and E. Arruti, "Implementation of complex enumeration for multiuser MIMO vector precoding",in *Proceedings of the 19th EURASIP European Signal Processing Conference (EUSIPCO'11)*, pp. 739-743, August 2011.

National conference papers:

- M. Barrenechea, M. Mendicute, and E. Arruti, "Vector precoding de complejidad fija para el canal MIMO multiusuario", in *Proceedings of the 24th National Symposium of the International Radio Scientific Union (URSI'09)*, September 2009.

- M. Barrenechea, M. Joham, M. Mendicute, and W. Utschick, "Cotas en la capacidad suma de la precodificacion vectorial", in *Proceedings of the 25th National Symposium of the International Radio Scientific Union (URSI'10)*, September 2010.

References

[Airy06] M. Airy, S. Bhadra, R. W. Heath, and S. Shakkottai, "Transmit precoding for the multiple antenna broadcast channel", in *Proceedings IEEE Vehicular Technology Conference (VTC'06)*, pp. 1396–1400, May 2006.

[Anderson84] J. B. Anderson and S. Mohan, "Sequential coding algorithms: A survey and cost analysis", *IEEE Transactions on Communications*, vol. 32, no. 2, 169–176, February 1984.

[Babai86] L. Babai, "On Lovász's Lattice Reduction and the Nearest Lattice Point Approximation", *Combinatorica*, vol. 6, no. 1, 1–14, March 1986.

[Barbero06a] L. Barbero, *Rapid prototyping of a fixed-complexity sphere decoder and its application to iterative decoding of turbo-MIMO systems*, Ph.D. thesis, University of Edinburgh, 2006.

[Barbero06b] L. Barbero and J. Thompson, "A fixed-complexity MIMO detector based on the complex sphere decoder", in *Proceedings IEEE International Workshop on Signal Processing Advances for Wireless Communications (SPAWC'06)*, pp. 1 – 5, July 2006.

[Barbero06c] L. Barbero and J. Thompson, "FPGA design considerations in the implementation of a fixed-throughput sphere decoder for MIMO systems", in *Proceedings IEEE International Conference on Field Programmable Logic and Applications (FPL '06)*, pp. 1–6, August 2006.

[Barbero06d] L. Barbero and J. Thompson, "Rapid prototyping of a fixed-throughput sphere decoder for MIMO systems", in *Proceedings IEEE International Conference on Communications (ICC'06)*, vol. 7, pp. 3082 – 3087, June 2006.

[Barbero08] L. Barbero and J. S. Thompson, "Extending a fixed-complexity sphere decoder to obtain likelihood information for turbo-MIMO systems", *IEEE*

Transactions on Vehicular Technology, vol. 57, no. 5, 2804–2814, September 2008.

[Barrenechea09a] M. Barrenechea, M. Mendicute, and E. Arruti, "Precodificación vectorial de complejidad fija para el canal MIMO multiusuario", in *Proceedings National Symposium of the International Radio Scientific Union (URSI'09)*, pp. 1–4, Septiembre 2009.

[Barrenechea09b] M. Barrenechea, M. Mendicute, J. Del Ser, and J. Thompson, "Wiener filter-based fixed-complexity vector precoding for the MIMO downlink channel", in *Proceedings IEEE International Workshop on Signal Processing Advances for Wireless Communications (SPAWC'09)*, pp. 216 – 220, June 2009.

[Barrenechea09c] M. Barrenechea, J. S. Thompson, M. Mendicute, and J. Del Ser, "Fixed-complexity regularized vector precoding for the multiuser MIMO downlink channel", in *Proceedings EURASIP European Signal Processing Conference (EUSIPCO'09)*, pp. 2406–2410, August 2009.

[Barrenechea10a] M. Barrenechea, L. Barbero, M. Mendicute, and J. Thompson, "Design and hardware implementation of a low-complexity multiuser vector precoder", in *Proceedings Conference on Design and Architectures for Signal and Image Processing (DASIP'10)*, pp. 160–167, October 2010.

[Barrenechea10b] M. Barrenechea, M. Joham, M. Mendicute, and W. Utschick, "Analysis of vector precoding at high SNR: rate bounds and ergodic results", in *Proceedings IEEE Global Communications Conference (GLOBECOM'10)*, pp. 1–5, December 2010.

[Barrenechea10c] M. Barrenechea, M. Joham, M. Mendicute, and W. Utschick, "Cotas en la capacidad suma de la precodificacion vectorial", in *Proceedings National Symposium of the International Radio Scientific Union (URSI'10)*, pp. 1–4, September 2010.

[Barrenechea11a] M. Barrenechea, L. Barbero, I. Jiménez, E. Arruti, and M. Mendicute, "High-throughput implementation of tree-search algorithms for vector precoding", in *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'11)*, pp. 1689–1692, May 2011.

[Barrenechea11b] M. Barrenechea, L. Barbero, M. Mendicute, and J. Thompson, "Design and implementation of a low-complexity multiuser vector precoder", *International Journal on Embedded and Real-Time Communication Systems*, (accepted) 2011.

[Barrenechea11c] M. Barrenechea, A. Burg, and M. Mendicute, "Reduced complexity tree-search-based vector precoding for multiuser systems", *IEEE Transactions on Communications*, (submitted) 2011.

[Barrenechea11d] M. Barrenechea, M. Mendicute, I. Jiménez, and E. Arruti, "Implementation of complex enumeration for multiuser MIMO vector precoding", in *Proceedings EURASIP European Signal Processing Conference (EUSIPCO'11)*, pp. 739–743, August 2011.

[Bhagawat08] P. Bhagawat, P. Wang, M. Uppal, G. Choi, Z. Xiong, M. Yeary, and A. Harris, "An FPGA implementation of dirty paper precoder", in *Proceedings IEEE International Conference on Communications (ICC'08)*, pp. 2761 – 2766, June 2008.

[Boccardi06] F. Boccardi and G. Caire, "The p-sphere encoder: peak-power reduction by lattice precoding for the MIMO Gaussian broadcast channel", *IEEE Transactions on Communications*, vol. 54, no. 11, 2085–2091, November 2006.

[Burg04] A. Burg, M. Wenk, M. Zellweger, M. Wegmueller, N. Felber, and W. Fichtner, "VLSI implementation of the sphere decoding algorithm", in *Proceedings European Solid-State Circuits Conference (ESSCIRC'04)*, pp. 303–306, September 2004.

[Burg05] A. Burg, M. Borgmann, M. Wenk, M. Zellweger, W. Fichtner, and H. Bölcskey, "VLSI implementation of MIMO detection using the sphere decoding algorithm", *IEEE Journal of Solid-State Circuits*, vol. 40, no. 7, 1566–1577, July 2005.

[Burg07] A. Burg, D. Seethaler, and G. Matz, "VLSI implementation of a lattice-reduction algorithm for multi-antenna broadcast precoding", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'07)*, pp. 673 – 676, May 2007.

[Caire00] G. Caire and S. Shamai, "On achievable rates in a multi-antenna broadcast downlink", in *Proceedings Annual Allerton Conference on Communications and Control and Computing (ALLERTON'00)*, pp. 1–10, October 2000.

[Callard06] A. Callard, A. Khandani, and A. Saleh, "Vector precoding with MMSE for the fast fading and quasi-static multi-user broadcast channel", in

Proceedings Annual Conference on Information Sciences and Systems (CISS'06), pp. 1002–1007, March 2006.

[Chen07] S. Chen, T. Zhang, and Y. Xin, "Relaxed K-Best MIMO signal detector design and VLSI implementation", *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 3, 328–337, March 2007.

[Christensen07] S. S. Christensen and E. de Carvalho, "Achievable sum-rates in MIMO broadcast channels with vector precoding techniques using coded modulation", in *Proceedings IEEE Vehicular Technology Conference (VTC'07)*, pp. 2261–2265, April 2007.

[Conway82] J. Conway and N. Sloane, "Voronoi regions of lattices, second moments of polytopes, and quantization", *IEEE Transactions on Information Theory*, vol. 28, no. 2, 211 – 26, March 1982.

[Conway84] J. H. Conway and N. J. A. Sloane, "On the Voronoi regions of certain lattices", *Journal on Algebraic Discrete Methods*, vol. 5, no. 3, 294–305, September 1984.

[Conway85] J. Conway and N. Sloane, "A lower bound on the average error of vector quantizers", *IEEE Transactions on Information Theory*, vol. 9, no. 31, 106–110, January 1985.

[Conway98] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, Springer, 3 edn., 1998.

[Costa83] M. Costa, "Writing on dirty paper", *IEEE Transactions on Information Theory*, vol. 29, no. 3, 439–441, May 1983.

[Damen03] M. Damen, H. El Gamal, and G. Caire, "On maximum-likelihood detection and the search for the closest lattice point", *IEEE Transactions on Information Theory*, vol. 49, no. 10, 2389 –2402, October 2003.

[Dietrich08] F. A. Dietrich, *Robust Signal Processing for Wireless Communications*, Springer, 2008.

[Dirichlet50] G. L. Dirichlet, "Über die Reduktion der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen", *Journal für die Reine und Angewandte Mathematik*, vol. 40, 209–227, 1850.

[Fincke85] U. Fincke and M. Pohst, "Improved methods for calculating vectors of short length in a lattice, including a complexity analysis", *Mathematics of Computation*, vol. 44, no. 170, 463 – 71, April 1985.

[Fischer02a] R. F. Fischer, C. Windpassinger, A. Lampe, and J. B. Huber, "MIMO precoding for decentralized receivers", in *Proceedings IEEE International Symposium on Information Theory (ISIT'02)*, p. 496, July 2002.

[Fischer02b] R. F. H. Fischer, *Precoding and Signal Shaping for Digital Transmission.*, John Wiley & Sons, Inc., 2002.

[Fischer02c] R. F. H. Fischer, C. Windpassinger, A. Lampe, and J. B. Huber, "Space-time transmission using Tomlinson-Harashima precoding", in *Proceedings International ITG Conference on Source and Channel Coding (SCC'02)*, pp. 139 – 147, January 2002.

[Fischer03] R. F. Fischer, C. Stierstorfer, and C. Windpassinger, "Precoding and signal shaping for transmission over MIMO channels", in *Proceedings Canadian Workshop on Information Theory (CWIT'03)*, pp. 83–87, May 2003.

[Foschini98] G. Foschini and M. Gans, "On limits of wireless communications in a fading environment when using multiple antennas", *Wireless Personal Communications*, vol. 6, no. 3, 311 – 335, January 1998.

[Gesbert07] D. Gesbert, M. Kountouris, R. Heath, and C. Chae, "From single user to multi user communications: shifting the MIMO paradigm", *IEEE Signal Processing Magazine*, vol. 24, no. 5, 36–46, October 2007.

[Goldsmith03] A. Goldsmith, S. A. Jafar, N. Jindal, and S. Vishwanath, "Capacity limits of MIMO channels", *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 5, 684–702, June 2003.

[Grötschel93] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Springer, 1993.

[Guo06] Z. Guo and P. Nilsson, "Algorithm and implementation of the K-Best sphere decoding for MIMO detection", *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, 491–503, March 2006.

[Gupta99] A. Gupta and D. Nagar, *Matrix Variate Distributions*, Chapman & Hall, 1999.

[Habendorf06] R. Habendorf, I. Riedel, and G. Fettweis, "Reduced complexity vector precoding for the multiuser downlink", in *Proceedings IEEE Global Telecommunications Conference (GLOBECOM'06)*, pp. 1–5, November 2006.

[Habendorf07] R. Habendorf and G. Fettweis, "Vector precoding with bounded complexity", in *Proceedings IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC'07)*, pp. 1–5, June 2007.

[Hadamard93] J. Hadamard, "Résolution dŠune question relative aux dťeterminants", *Bulletin des Sciences Mathématiques*, vol. 2, no. 17, 240Ű248, 1893.

[Harashima72] H. Harashima and H. Miyakawa, "Matched-transmission technique for channels with intersymbol interference", *IEEE Transactions on Communications*, vol. CM-20, no. 4, 774 – 780, August 1972.

[Hess07] C. Hess, M. Wenk, A. Burg, P. Luethi, C. Studer, N. Felber, and W. Fichtner, "Reduced-complexity MIMO detector with close-to ML error rate performance", in *Proceedings ACM Great Lakes Symposium on VLSI (GLSVLSI'07)*, pp. 200–203, March 2007.

[Hochwald03] B. M. Hochwald and S. Ten Brink, "Achieving near-capacity on a multiple-antenna channel", *IEEE Transactions on Communications*, vol. 51, no. 3, 389 – 399, March 2003.

[Hochwald05] B. Hochwald, C. Peel, and A. Swindlehurst, "A vector-perturbation technique for near-capacity multiantenna multiuser communication. Part II: perturbation", *IEEE Transactions on Communications*, vol. 53, no. 3, 537–544, March 2005.

[Hur07] S. Hur, N. Kim, H. Park, and J. Kang, "Enhanced lattice-reduction-based precoder with list quantizer in broadcast channel", in *Proceedings IEEE Vehicular Technology Conference (VTC'07)*, pp. 611–615, October 2007.

[Jaldén08] J. Jaldén, J. Maurer, and G. Matz, "On the diversity order of vector perturbation precoding with imperfect channel state information", in *Proceedings IEEE Signal Processing Advances in Wireless Communications (SPAWC'08)*, pp. 211–215, July 2008.

[Jindal01] N. Jindal, S. Vishwanath, and A. Goldsmith, "On the Duality of Multiple-Access and Broadcast Channels", in *Proceedings Annual Allerton Conference on Commununications and Control and Computing (ALLERTON'01)*, pp. 3–5, September 2001.

[Joham04a] M. Joham, *Optimization of linear and nonlinear transmit signal processing*, Ph.D. thesis, Munich Technical University, 2004.

[Joham04b] M. Joham, J. Brehmer, and W. Utschick, "MMSE approaches to multiuser spatio-temporal Tomlinson-Harashima precoding", in *Proceedings International ITG Conference on Source and Channel Coding (SCC'04)*, pp. 387 – 394, January 2004.

[Joham05a] M. Joham and W. Utschick, *Smart Antennas - State-of-the-Art*, chap. Ordered spatial Tomlinson Harashima precoding, pp. 401–422, Hindawi Publishing Corporation, 2005.

[Joham05b] M. Joham, W. Utschick, and J. A. Nossek, "Linear transmit processing in MIMO communications systems", *IEEE Transactions on Signal Processing*, vol. 53, no. 8, 2700–2712, August 2005.

[Krithivasan07] D. Krithivasan and S. S. Pradhan, "A proof of the existence of good nested lattices.", Tech. rep., University of Michigan, 2007.

[Kusume05] K. Kusume, M. Joham, W. Utschick, and G. Bauch, "Efficient Tomlinson-Harashima precoding for spatial multiplexing on flat MIMO channel", in *Proceedings IEEE International Conference on Communications (ICC'05)*, vol. 3, pp. 2021 –2025, May 2005.

[Kusume07] K. Kusume, M. Joham, W. Utschick, and G. Bauch, "Cholesky factorization with symmetric permutation applied to detecting and precoding spatially multiplexed data streams", *IEEE Transactions on Signal Processing*, vol. 55, no. 6 II, 3089 – 3103, June 2007.

[Lenstra82] A. Lenstra, H. Lenstra, and L. Lovász, "Factoring polynomials with rational coefficients", *Mathematische Annalen*, vol. 216, 513–534, 1982.

[Li06] Q. Li and Z. Wang, "Improved K-Best Sphere Decoding Algorithms for MIMO Systems", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'06)*, pp. 1159–1162, May 2006.

[Lin06] K. Lin, H. Lin, R. Chang, and C. Wu, "Hardware architecture of improved Tomlinson-Harashima precoding for downlink MC-CDMA", in *Proceedings IEEE Asia Pacific Conference on Circuits and Systems (APCCAS'06)*, pp. 1200–1203, December 2006.

[Liu07] F. Liu, L. Jiang, and C. He, "Low complexity MMSE vector precoding using lattice reduction for MIMO systems", in *Proceedings IEEE International Conference on Communications (ICC'07)*, pp. 2598–2603, June 2007.

[Mahdavi10] M. Mahdavi, M. Shabany, and V. Vahdat, "A modified complex K-best scheme for high-speed hard-output MIMO detectors", in *Proceedings IEEE International Midwest Symposium on Circuits and Systems (MWSCAS'10)*, pp. 845–848, August 2010.

[Mennenga09] B. Mennenga and G. Fettweis, "Search Sequence Determination for Tree Search Based Detection Algorithms", in *Proceedings IEEE Sarnoff Symposium (SARNOFF'09)*, pp. 1–6, March 2009.

[Messenger09] C. Messenger, R. Prix, and M. Papa, "Random template banks and relaxed lattice coverings", *Physical Review D*, vol. 79, no. 10, 104–117, May 2009.

[Mohaisen11] M. Mohaisen and K. Chang, "Fixed-complexity sphere encoder for multi-user MIMO systems", *Journal of Communications and Networks (JCN)*, vol. 13, no. 1, 63–39, February 2011.

[Mondal08] S. Mondal, W. Ali, and K. Salama, "A novel approach for K-Best MIMO detection and its VLSI implementation", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'08)*, pp. 936–939, May 2008.

[Müller08] R. Müller, D. Guo, and A. Moustakas, "Vector precoding for wireless MIMO systems and its replica analysis", *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 3, 530–540, April 2008.

[Park08] K. Park, J. Cha, and J. Kant, "A computationally efficient stack-based iterative precoding for multiuser MIMO broadcast channel", in *Proceedings IEEE Vehicular Technology Conference (VTC'08)*, pp. 1–5, September 2008.

[Peel05] C. B. Peel, B. M. Hochwald, and A. L. Swindlehurst, "A vector-perturbation technique for near-capacity multiantenna multiuser communication. Part I: channel inversion and regularization", *IEEE Transactions on Communications*, vol. 53, no. 3, 195–202, March 2005.

[Razi09] A. Razi, D. J. Ryan, I. B. Collings, and J. Yuan, "Sum rates and user scheduling for multi-user MIMO vector perturbation precoding", in *Proceedings IEEE International Conference on Communications (ICC'09)*, pp. 1–5, June 2009.

[Razi11] A. Razi, D. Ryand, J. Yuan, and I. Collings, "Sum rates for multi-user MIMO regularized vector perturbation precoding", in *Proceedings International Symposium on Wireless Communication Systems (ISWCS'11)*, November 2011.

[Rogers58a] C. Rogers, "Lattice coverings of space: The Minkowski-Hlawka theorem", *Proceedings of the London Mathematical Society*, vol. 8, no. 3, 447–465, 1958.

[Rogers58b] C. Rogers, "Lattice coverings of space with convex bodies", *Journal of London Mathematic Society*, vol. 33, no. 2, 208–212, April 1958.

[Ryan08] D. Ryan, I. B. Collings, I. V. Clarkson, and R. Heath, "A lattice-theoretic analysis of vector perturbation for multi-user MIMO systems", in *Proceedings IEEE International Conference on Communications (ICC'08)*, pp. 3340–3344, May 2008.

[Ryan09] D. J. Ryan, I. B. Collings, I. V. L. Clarkson, and R. W. Heath Jr., "Performance of vector perturbation multiuser MIMO systems with limited feedback", *IEEE Transactions on Communications*, vol. 57, no. 9, 2633 – 2644, September 2009.

[Sato78] H. Sato, "An outer bound to the capacity region of broadcast channel", *IEEE Transactions on Information Theory*, vol. 24, no. 3, 374–377, May 1978.

[Schmidt05] D. Schmidt, M. Joham, and W. Utschick, "Minimum mean square error vector precoding", in *Proceedings IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'05)*, vol. 1, pp. 107–111, September 2005.

[Schmidt08] D. Schmidt, M. Joham, and W. Utschick, "Minimum mean square error vector precoding", *European Transactions on Telecommunications*, vol. 19, 219–231, September 2008.

[Schnorr91] C. P. Schnorr and M. Euchner, "Lattice basis reduction: improved practical algorithms and solving subset sum problems", in *Proceedings International Symposium on Fundamentals of Computation Theory (FCT'91)*, vol. 529, pp. 68–85, September 1991.

[Seethaler06] D. Seethaler and G. Matz, "Efficient vector perturbation in multi-antenna multi-user systems based on approximate integer relations", in *Proceed-*

ings EURASIP European Signal Processing Conference (EUSIPCO'06), pp. 1–5, September 2006.

[Shabany08a] M. Shabany and P. Gulak, "Scalable VLSI architecture for K-best lattice decoders", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'08)*, pp. 940 – 943, May 2008.

[Shabany08b] M. Shabany, K. Su, and P. Gulak, "A pipelined scalable high-throughput implementation of a near-ML K-Best complex lattice decoder", in *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'08)*, pp. 3173 – 3176, March 2008.

[Shannon48] C. E. Shannon, "A Mathematical Theory of Communications", *Bell Systems Technical Journal*, vol. 27, 623–656, July 1948.

[Shen10] C. Shen and A. Eltawil, "A radius adaptive K-best decoder with early termination: algorithm and VLSI architecture", *IEEE Transaction on Circuits and Systems*, vol. 57, no. 9, 2476–2486, September 2010.

[Taherzadeh05] M. Taherzadeh, A. Mobasher, and A. Khandani, "LLL lattice-basis reduction achieves the maximum diversity in MIMO systems", in *Proceedings IEEE International Symposium on Information Theory (ISIT'05)*, pp. 1300 –1304, September 2005.

[Taherzadeh07] M. Taherzadeh, A. Mobasher, and A. K. Khandani, "Communication over MIMO broadcast channels using lattice-basis reduction", *IEEE Transactions on Information Theory*, vol. 53, no. 12, 4567–4582, December 2007.

[Telatar95] I. Telatar, "Capacity of multi-antenna Gaussian channels", Tech. rep., AT&T Bell Laboratories, 1995.

[Telatar99] I. Telatar, "Capacity of multi-antenna Gaussian channels", *European Transactions on Telecommunications*, vol. 10, no. 6, 585–595, November 1999.

[Tomlinson71] M. Tomlinson, "New automatic equaliser employing modulo arithmetic", *Electronics Letters*, vol. 7, no. 5-6, 138 –139, March 1971.

[Tsai10] P. Tsai, W. Chen, and M. Huang, "A 4x4 64-QAM Reduced-Complexity K-Best MIMO Detector up to 1.5 Gbps", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'10)*, pp. 3953–3956, May 2010.

[Tse05] P. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, Cambridge University Press, 2005.

[Tulino04] A. Tulino and S. Verdú, *Random Matrix Theory*, Now Publishers Inc, 2004.

[vanEmde Boas81] P. van Emde Boas, "Another NP-complete partition problem and the complexity of computing short vectors in lattices", Tech. rep., University of Amnsterdam, 1981.

[Verdú98] S. Verdú, *Multiuser Detection*, Cambridge University Press, 1998.

[Vishwanath02] S. Vishwanath, N. Jindal, and A. Goldsmith, "On the capacity of multiple-input multiple-output broadcast channels", in *Proceedings of the International Conference on Communications (ICC'02)*, vol. 3, pp. 1444–1450, April 2002.

[Viterbo99] E. Viterbo and J. Boutros, "A universal lattice code decoder for fading channels", *IEEE Transactions on Information Theory*, vol. 45, no. 5, 1639–1642, July 1999.

[Voronoi07] G. Voronoi, "Nouvelles applications des paramètres continus à la théorie des formes quadratiques", *Journal für die Reine und Angewandte Mathematik*, vol. 133, 97–178, January 1907.

[Wübben01] D. Wübben, R. Böhnke, J. Rinas, V. Kühn, and K. Kammeyer, "Efficient algorithm for decoding layered space-time codes", *IEEE Electronic Letters*, vol. 37, no. 22, 1348–1350, October 2001.

[Wenk06] M. Wenk, M. Zellweger, A. Burg, N. Felber, and W. Fichtner, "K-Best MIMO detection VLSI architectures achieving up to 424 Mbps", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'06)*, pp. 1151–1154, September 2006.

[Wenk10a] M. Wenk, *MIMO-OFDM Testbed: Challenges, Implementations, and Measurement Results*, Ph.D. thesis, Eidgennössische Technische Hochschule (ETH) Zürich, 2010.

[Wenk10b] M. Wenk, L. Bruderer, A. Burg, and C. Studer, "Area- and throughput-optimized VLSI architecture of sphere decoding", in *Proceedings VLSI and System-on-Chip (VLSI-SoC'10)*, pp. 189–194, September 2010.

[Windpassinger04] C. Windpassinger, R. F. H. Fischer, and J. B. Huber, "Lattice-reduction-aided broadcast precoding", *IEEE Transactions on Communications*, vol. 52, no. 12, 2057–2060, December 2004.

[Wong02] K. Wong, C. Tsui, R. Cheng, and W. Mow, "A VLSI architecture of a K-best lattice decoding algorithm for MIMO channels", in *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'02)*, vol. 3, pp. 273–276, May 2002.

[Yu01] W. Yu and J. M. Cioffi, "Trellis precoding for the broadcast channel", in *Proceedings IEEE Global Communications Conference (GLOBECOM'01)*, vol. 2, pp. 1344–1348, November 2001.

[Zador63] P. Zador, *Development and Evaluation of Procedures for Quantizing Multivariate Distributions*, Ph.D. thesis, Stanford University, 1963.

[Zaidel08] B. M. Zaidel, R. R. Müller, R. de Miguel, and A. L. Moustakas, "On spectral efficiency of vector precoding for Gaussian MIMO broadcast channels", in *Proceedings IEEE International Symposium on Spread Spectrum Techniques and Applications (ISSSTA'08)*, pp. 232–236, August 2008.

[Zamir96] R. Zamir and M. Feder, "On lattice quantization noise", *IEEE Transactions on Information Theory*, vol. 42, no. 4, 1152 – 1159, July 1996.

[Zhang05] J. Zhang and K. J. Kim, "Near-capacity MIMO multiuser precoding with QRD-M algorithm", in *Proceedings Asilomar Conference on Signals, Systems, and Computers (ACSSC'05)*, vol. 1, pp. 1498–1502, October 2005.