

Compressed Sensing of Memoryless Sources: A Deterministic Hadamard Construction

THÈSE N° 6356 (2014)

PRÉSENTÉE LE 12 DÉCEMBRE 2014

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

LABORATOIRE DE COMMUNICATIONS MOBILES

PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Saeid HAGHIGHATSHOAR

acceptée sur proposition du jury:

Prof. J.-P. Hubaux, président du jury

Prof. B. Rimoldi, directeur de thèse

Prof. H. Bölcskei, rapporteur

Prof. A. Montanari, rapporteur

Prof. E. Telatar, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2014

Thesis No. 6356 (September 2014)

The thesis is presented to the faculty of Computer and Communication Sciences for obtaining the degree of Docteurs Sciences.

Accepted by the jury members:

Bixio Rimoldi

Thesis director

Emre Telatar

Expert

Helmut Bölskei

Expert

Andrea Montanari

Expert

Jean-Pierre Hubaux

President of the jury

To my father, the man of honor.

and

To my mother, the ocean of love

Acknowledgements

My PhD journey at EPFL started in September 2010. Now that after four years I have finally defended my thesis and finished my PhD, I owe my gratitude to all the people who made this journey nothing short of amazing.

First of all, I would like to thank my advisor Prof. Bixio Rimoldi. Definitely, without his constant support and encouragement, this thesis would have not been possible. He always kept his office doors open for me, supported me financially for essentially any conferences that I wanted to attend, and gave me the courage and self-confidence to be an independent researcher. Over the past four years, he has taught me more than one can imagine. He helped me a lot to improve my writing skills and taught me to keep a high standard for my publications. It is really difficult to imagine an advisor who gives full freedom to his student and does not even accept his name to be on a paper because he has not actively worked on it. No words can express my gratitude for his trust and fatherly support during these four years.

I would like to thank Prof. Emre Telatar, Prof. Helmut Bölcskei, Prof. Andrea Montanari and Prof. Jean-Pierre Hubaux who accepted to be in my defense committee.

I am really grateful to Prof. Emre Telatar and Prof. Ruediger Urbanke who along with Bixio have made the Information Processing Group (IPG), in my opinion, the best research lab at EPFL. I am specially thankful to Emre for all his care and friendly advices. He is a great teacher and a kind friend whom one can always trust in many aspects. His astonishingly relevant comments about the research problems have always been my main motivation to approach and solve them. I should also thank Ruediger for all his private advices, usually with his office door closed, specially when I encountered really frustrating problems in my research.

I am also grateful to two great mentors in IPG, Dr. Olivier Lévêque and Dr. Nicolas Macris. I learned probability from Olivier the first year I came to EPFL. I attended his “Advanced Probability” course, which motivated me to study this amazingly beautiful mathematical topic. I am also very grateful to Nicolas for all the intuitive discussions that we had about Statistical Physics and Quantum Information Theory while I was a Teaching Assistant (TA) for his course.

The life would be so difficult in IPG without the help of our secretaries Mrs. Francoise Behn and Mrs. Muriel Bardet. All their selfless efforts have made the IPG the most pleasant lab without any need to take care of the administrative stuff such as reserving flights, registering conferences, etc. I am deeply grateful for all their dedication, help, warm welcomes and friendly support during these four years. I am sure that one day I will regret missing the opportunity that I had to learn French

from them.

I am thankful to our system administrator Mr. Damir Laurenzi for helping me with so many naive mistakes that I made on my computer system.

My deepest gratitude goes to all my PhD colleagues at IPG who made it the best work environment possible. Their cheerful, friendly and warm characters have always been a blessing for me. Among them I am specially grateful to my Office-mate Mr. Ayan Sengupta. I will really miss all the discussions that I had with him on many diverse topics such as culture, politics, psychology, etc. I learned a lot about India and Indians from him. He has always been a huge source of encouragement and friendly support for me, which I would always appreciate.

Four years of research, with its ups and downs, would have not been possible without the help of three special friends at IPG whom I deeply acknowledge. First, I would like to thank Dr. Emmanuel Abbe who was a postdoctoral researcher at IPG when I started my PhD (He is a tenure-track professor at Princeton University at the moment). All the scientific discussions and fruitful meetings with him finally motivated me to follow the research that has shaped my PhD thesis. I am also grateful for all his funny discussions and for his pleasant sense of humor. I wish him success and happiness in all his life and scientific career. Second, I am grateful to Dr. Amin Karbasi who was also a PhD student at IPG (He defended in 2012 and he is a tenure-track professor at Yale University at the moment). There are many things that I learned from him, too many to mention here. We took many courses at EPFL together and had many discussions on research and life. I always recall his advice while deciding to do or learn something: “divide it into so many small pieces that you can not evade doing them!”. Third, I would like to thank my friend Dr. Hamed Hassani who was also a PhD student at IPG (He defended in 2013). He has always been a source of motivation for me in many aspects, specially in research. He made the research so smooth for me in so many tough occasions. I am really grateful for all his friendly support.

I am grateful to all my Iranian friends in Lausanne who made the life in Switzerland so amazing and memorable. In particular, I am so thankful to my friends known among iranians as Tir-Federal crew: Mohammad Parhizkar, Reza Parhizkar, Hamed Hassani, Farid Movahedi Naini and Masih Nilchian. We shared an apartment for more than a year and we had so many memories together. All the time that we spent together (literally wasted on watching movies) and all the friendship and dedication that we had for one another is beyond imagination. I hope the best wishes for them all. I am sure I will miss forever these amazing friends and the unforgettable moments that I spent with them.

I would have not succeeded either in life or in my research without the unconditional love, support and patience of my family. My deepest gratitude goes to my father for being an amazing father, a true friend and a great support all throughout my life. He taught me how to follow the path of honor, how to devote myself to others and how to always live away from pride and ignorance. I am thankful to my mother for her unconditional love and all the sacrifices that she made in her life for her children and specially for me. There are no words that can express my deepest gratitude to her. My special gratitudes go to my father-in-law and mother-in-law who always treated me like a son and trusted me in crucial occasions of my married life. I also deeply thank my two brother-in-laws who have always been like true brothers to me. My deepest gratitude goes to my three lovely sisters. Being the only

brother, I have always had the great privilege to be loved and to be supported by them all the time. I am also deeply grateful to my three little nephews who have always cheered me up as an uncle and have been giving me the happiest moments in my life.

Special thanks and recognitions go to my beloved wife for her deep love and heart-filling affection. I truly thank her for sticking by my side and believing in me even in those moments that I did not believe in myself. There are no words that can truly express my gratitude and love for her.

Abstract

Compressed sensing is a new trend in signal processing for efficient data sampling and signal acquisition. The idea is that most real-world signals have a sparse representation in an appropriate basis and this can be exploited to capture the sparse signal by taking a few linear projections. The recovery is possible by running appropriate low-complexity algorithms that exploit the sparsity (prior information) to reconstruct the signal from the linear projections (posterior information). The main benefit is that the required number of measurements is much smaller than the dimension of the signal. This results in a huge gain in sensor cost (in measurement devices) or a dramatic saving in data acquisition time. The drawback is that one needs complicated devices in order to implement the linear projections.

Some difficulties naturally arise in applying the compressed sensing to real-world applications such as robustness issues in taking the linear projections and computational complexity of the recovery algorithm. Robustness issue arises because even if the devices are precisely calibrated, there is still a mismatch with the intended linear projection. Obtaining stable and low-complexity recovery algorithms has also been a challenge in compressed sensing. Although there are numerically stable convex optimization algorithms for recovery, their complexity usually scales like $O(n^3)$ in signal dimension n , which prohibits their use in high dimensional applications that are encountered more and more nowadays. Consequently, there have been different attempts to reduce this complexity as much as possible.

In this thesis, we design structured matrices for compressed sensing. In particular, we claim that some of the practical difficulties can be reasonably solved by imposing some structure on the measurement matrices. Almost all the thesis evolves around the Hadamard matrices, which are $\{+1, -1\}$ -valued matrices with many applications in signal processing, coding theory, optics and theoretical mathematics. As the title of the thesis implies, there are two main ingredients to this thesis. First, we use a memoryless assumption for the source, i.e., we assume the nonzero components of the sparse signal are independently generated by a given probability distribution and their position is completely random. This is not a major restriction because most of the results obtained are not sensitive to the shape of the distribution. The advantage is that the probabilistic model of the signal allows us to use tools from probability, information theory and coding theory to rigorously assess the achievable performance. Second, using the mathematical properties of the Hadamard matrices, we design deterministic matrices for compressed sensing of memoryless sources by selecting specific rows of a Hadamard matrix according to a deterministic criterion. We call the resulting matrices *partial Hadamard matrices*.

We design partial Hadamard matrices for three distinct signal models: memoryless discrete sources and sparse signals with linear or sub-linear sparsity. A signal has linear sparsity if the number of its nonzero components k is proportional to n , the dimension of signal, whereas it has a sub-linear sparsity if the k scales like $O(n^\alpha)$ for some $\alpha \in (0, 1)$. We develop tools to rigorously analyze the performance of the proposed Hadamard constructions by borrowing ideas from information theory and coding theory.

In the last part of the thesis, we extend our construction to distributed (multi-terminal) signals. Distributed compressed sensing is a very interesting and ubiquitous problem in distributed data acquisition systems such as ad-hoc sensor networks. From both a theoretical and an engineering point of view, it is important to know how many measurements per dimension are necessary from different terminals in order to have a reliable estimate of the distributed data. We analyze this problem for a very simple setup, where the components of the distributed signal are generated by a joint probability distribution which, in some sense, captures the spatial correlation among different terminals. We give an information-theoretic characterization of the measurement-rate region that results in a negligible recovery distortion. We also propose a low-complexity distributed message passing algorithm to achieve the theoretical limits.

Keywords: Compressed sensing, Hadamard matrices, Deterministic matrix construction, Sparse Fast Hadamard Transform (SFHT), Distributed compressed sensing.

Résumé

L'acquisition comprimée est une nouvelle tendance dans le domaine du traitement de signal pour l'échantillonnage de données et l'acquisition de signal efficace. L'idée est que dans la plupart des applications réelles, les signaux ont une représentation parcimonieuse dans une base appropriée et celle-ci peut être exploitée pour capter un signal parcimonieux avec un petit nombre de projections linéaires. La reconstruction est possible par l'exécution d'algorithmes de faible complexité appropriés qui exploitent la parcimonie (information préalable) afin de reconstruire le signal à partir de projections linéaires (information postérieure). L'avantage est que le nombre de mesures nécessaires est beaucoup plus moins que la dimension du signal. Cela implique un gain important dans le coût du capteur (dans les dispositifs de mesures) ou une économie considérable en temps d'acquisition des données. L'inconvénient est que l'on a besoin d'appareils compliqués pour réaliser les projections linéaires.

Certaines difficultés surgissent naturellement dans l'application de la théorie de l'acquisition comprimée aux applications concrètes, notamment les questions de la robustesse du processus de prise de projections linéaires et de la complexité de l'algorithme de reconstruction. Les problèmes de robustesse se posent car même si les appareils sont précisément calibrés, un mésappariement avec la projection linéaire prévue existe tout de même. Obtenir des algorithmes de reconstruction stables et de faible complexité est aussi un défi pour l'acquisition comprimée. Bien qu'il existe des algorithmes d'optimisation convexe numériquement stable pour la reconstruction, leurs complexités sont généralement d'ordre $\Theta(n^3)$ dans la dimension n du signal, empêchant leur utilisation dans le cas des applications de grande dimension que l'on rencontre de plus en plus de nos jours. Par conséquent, différentes tentatives ont été faites pour réduire cette complexité au minimum.

Dans cette thèse, nous concevons des matrices structurées pour l'acquisition comprimée. En particulier, nous affirmons que certaines des difficultés pratique peuvent être raisonnablement résolues en imposant une structure aux matrices de mesure. Presque toute la thèse évolue autour des matrices de Hadamard qui sont des matrices dont les coefficients sont $\{+1, -1\}$ avec de nombreuses applications dans les domaines du traitement de signal, la théorie du codage, l'optique, et les mathématiques théoriques. Comme le titre de la thèse l'indique, il existe deux ingrédients principaux à cette thèse. Premièrement, nous supposons que la source est sans mémoire, c'est-à-dire que nous supposons que les composantes non nulles du signal parcimonieux sont sélectionnées indépendamment selon une distribution de probabilité donnée et leur position est complètement aléatoire. Cette hypothèse n'impose pas une restriction majeure car la plupart des résultats obtenus ne sont

pas sensibles à la forme de la distribution. L'avantage est que le modèle probabiliste du signal nous permet d'utiliser des outils de probabilité, la théorie de l'information, et la théorie du codage pour évaluer rigoureusement les performances réalisables. Deuxièmement, en utilisant les propriétés mathématiques des matrices de Hadamard, nous développons des matrices déterministes pour l'acquisition comprimée des sources sans mémoire en sélectionnant des lignes spécifiques d'une matrice de Hadamard selon un critère déterministe. Nous appelons les matrices résultantes les *matrices de Hadamard partielles*.

Nous concevons des matrices de Hadamard partielles pour trois modèles distincts: les sources discrètes sans mémoire et les signaux parcimonieux avec parcimonie linéaire ou sous-linéaire. Un signal a une parcimonie linéaire si le nombre k de ses composantes non nulles est proportionnelle à n , la dimension du signal, alors qu'il a une parcimonie sous-linéaire si k est d'ordre $\Theta(n^\alpha)$ pour un certain $\alpha \in (0, 1)$. Nous développons des outils pour analyser rigoureusement la performance des constructions de Hadamard proposées en empruntant des idées de la théorie de l'information et de la théorie du codage.

Dans la dernière partie de cette thèse, nous étendons notre construction aux signaux distribués (multi-terminaux). L'acquisition comprimée distribuée est un problème très intéressant et omniprésent dans les systèmes d'acquisition de données distribuées tels que les réseaux ad-hoc de capteurs. Tant d'un point de vue théorique et de l'ingénierie, il est important de savoir combien de mesures par dimension sont nécessaires des différents terminaux afin d'avoir une estimation fiable des données distribuées. Nous analysons ce problème pour une configuration très simple, où les composantes du signal distribué sont générées par une distribution de probabilité conjointe qui, dans un certain sens, capte la corrélation spatiale entre les différents terminaux. Nous donnons une caractérisation dans le sens de la théorie de l'information de la région de mesure-débit qui entraîne une distorsion de reconstruction négligeable. Nous proposons aussi un algorithme de propagation de message distribué de faible complexité pour atteindre les limites théoriques.

Mots clés: Acquisition comprimée, Matrices de Hadamard, Construction de matrice déterministe, Transformée parcimonieuse de Hadamard rapide (SFHT), Acquisition comprimée distribué

Contents

Acknowledgements	v
Abstract	ix
Résumé	xi
Contents	xiii
List of Figures	xvi
List of Tables	xviii
1 Introduction	1
1.1 Linear Inverse Problems	2
1.1.1 Main Essence of a Linear Inverse Problem	3
1.1.2 Challenges in Application	5
1.2 Contribution of this Thesis	7
1.2.1 Hadamard Matrices	7
1.2.2 Probabilistic Model for the Signal	9
1.2.3 Distributed Compressed Sensing	9
1.3 Applications beyond Compressed Sensing	11
1.4 Outline of the Thesis	13
2 Hadamard Construction for Discrete Memoryless Sources	17
2.1 Introduction and Related Work	17
2.2 Polarization Theory over Finite Field \mathbb{F}_q	18
2.3 What about Integer-valued Sources?	19
2.4 Hadamard Matrices and the Entropy Process	20
2.5 Deterministic Partial-Hadamard Matrix Construction	22
2.6 Lower Bound on the Rate of Absorption	23
2.7 Simulation Results	24
2.7.1 Absorption Phenomenon	24
2.7.2 Nested Property	25
2.7.3 Robustness to Measurement Noise	25
3 Fast Hadamard Transform: Construction for Signals with Sub-linear Sparsity	27

3.1	Walsh-Hadamard Transform: Overview and Related Work	28
3.2	Main Results	30
3.3	Walsh-Hadamard Transform and its Properties	31
3.3.1	Basic Properties	32
3.4	Hadamard Hashing Algorithm	34
3.4.1	Properties of the Hadamard Hashing	35
3.5	Sparse Fast Hadamard Transform	36
3.5.1	Explanation of the Algorithm	36
3.5.2	Complexity Analysis	38
3.6	Performance Analysis of the very Sparse Regime	40
3.6.1	Hash Construction	41
3.6.2	Random Bipartite Graph Construction	42
3.6.3	Performance Analysis of the Peeling Decoder	46
3.7	Performance Analysis of the Less Sparse Regime	49
3.7.1	Hash Construction	50
3.7.2	Bipartite Graph Representation	51
3.7.3	Performance Analysis of the Peeling Decoder	51
3.7.4	Generalized Hash Construction	54
3.8	Simulation Results	55
3.9	Conclusion	58
3.10	Proof of the Auxiliary Results	59
3.10.1	Proof of the Properties of the WHT	59
3.10.2	Proof of Proposition 3.2	60
3.10.3	Proof of Proposition 3.3	61
3.10.4	Proof of Proposition 3.9	62
4	Rényi Polarization: Hadamard Construction for Signals with Linear Sparsity	65
4.1	Related Work	66
4.2	Rényi information dimension	68
4.3	Main results	71
4.3.1	Polarization of the Rényi information dimension	71
4.3.2	A2A compression	72
4.4	Proof Techniques	75
4.4.1	Rényi Information Dimension	75
4.4.2	Polarization of the RID	77
4.4.3	A2A Compression	79
4.5	Operational vs. Informational Characterization	80
4.6	Simulation Results	81
4.6.1	Signal Model and the Recovery Algorithm	81
4.6.2	Sensitivity to Signal Distribution	82
4.6.3	Comparison of the Performance of ℓ_1 -minimization and AMP	82
4.6.4	Comparison with Random Gaussian Matrices	82
4.7	Conclusion and Further Discussion	85
5	Multi-Terminal Compressed Sensing	87
5.1	Introduction and Related Work	88
5.2	Hadamard Construction for Multi-terminal A2A Compression	90

5.3	Gaussian Measurement Matrices	91
5.4	Spatially Coupled Gaussian Measurement Matrices	94
5.5	Simulation Results	98
5.5.1	Signal Model	98
5.5.2	Performance without Spatial Coupling	98
5.5.3	Performance with Spatial Coupling	101
5.6	Appendix	104
5.6.1	Proofs of the Hadamard Construction	104
5.6.2	Heuristic Derivation of the Multi-Terminal AMP	108
5.6.3	Heuristic Derivation of the State Evolution	110
5.6.4	MMSE Estimator Linearly Correlated Bernoulli-Gaussian Signals	111
A	Entropy Power Inequality for Integer-valued Random Variables	113
A.1	History and Introduction	114
A.2	Statement of the Results	115
A.3	Proof Techniques	116
A.3.1	EPI for i.i.d. Random Variables	117
A.3.2	EPI for non-i.i.d. random variables	118
A.3.3	Conditional EPI	120
A.4	Open problems	123
A.4.1	Closure convexity of the entropy set \mathcal{H}	123
A.5	Proof of Auxiliary Lemmas	124
A.5.1	EPI for i.i.d. random variables	124
A.5.2	EPI for non-i.i.d. random variables	126
A.5.3	Conditional EPI	128
	Bibliography	131

List of Figures

1.1	Compressed Sensing problem.	2
1.2	Linear Inverse Problem.	3
2.1	Representation of the Entropy Process on a Binary Rooted Tree	21
2.2	Absorption phenomenon for a binary source with $p(1) = 0.05$ and $N = 512$	24
2.3	Nested property for the absorption scheme of a binary source for different values of p	25
2.4	Stability of the Maximum Likelihood (ML) Decoder to additive Gaussian measurement noise.	26
3.1	Comparison of the Hadamard Construction for CS and WHT	29
3.2	Illustration of the downsampling property on a hypercube for $N = 2^3$. .	34
3.3	Illustration of the Hadamard Hashing.	35
3.4	Hadamard Hashing sparsifies the underlying graph.	37
3.5	A block diagram of the SFHT algorithm in the time domain.	37
3.6	Tree-like neighborhood an an edge $e = (v, c)$	48
3.7	Density Evolution equation for $C = 3$ and different values of $\beta = \frac{K}{B}$. . .	49
3.8	Bipartite graph representation for the less sparse case $\alpha = \frac{2}{3}$, $C = 3$. .	51
3.9	Probability of success of the algorithm as a function of α and C for deterministic hash construction.	55
3.10	Probability of success of the algorithm as a function of α and C for random hash construction.	56
3.11	Probability of success of the algorithm in the less sparse regime as a function of $\beta = K/B$	56
3.12	Comparison of the Median runtime in ms of the SFHT and conventional WHT.	57
3.13	Marginal value of α for better performance of SFHT compared with WHT	58
4.1	Polarization of the RID for $N = 512$ and $d(X) = 0.5$	72
4.2	Boundary of the Low-Distortion Region for ℓ_1 -minimization for Different Signal Distributions	83
4.3	Boundary of the Low-Distortion Region for AMP and ℓ_1 -minimization .	83
4.4	Rate-Distortion Region for Hadamard Construction and ℓ_1 -minimization	84
4.5	Rate-Distortion Region for Random Gaussian Matrices and ℓ_1 -minimization	84
5.1	Graphical Model Representation for Two-Terminal Compressed Sensing.	92

5.2	The Structure of a band-diagonal Gaussian matrix with non-homogenous entry variances.	95
5.3	Empirical and State Evolution result for T_X for $\rho_x = 0.5, \rho_y = 0.6$. . .	99
5.4	Empirical and State Evolution result for T_Y for $\rho_x = 0.5, \rho_y = 0.6$. . .	100
5.5	Empirical and State Evolution result for T_X for $\rho_x = 0.5, \rho_y = 0.7$. . .	100
5.6	Empirical and State Evolution result for T_Y for $\rho_x = 0.5, \rho_y = 0.7$. . .	101
5.7	Rate-Distortion region for a linearly correlated Bernoulli-Gaussian source with $d(X) = d(Y) = 0.44$ and $d(X Y) = d(Y X) = 0.248$	102
5.8	Rate-Distortion Region for the Linearly Correlated Bernoulli-Gaussian Source with $d(X) = d(Y) = 0.44$ and $d(X Y) = d(Y X) = 0.1802$	102
5.9	Rate-Distortion region for a linearly correlated Bernoulli-Gaussian source with $d(X) = d(Y) = 0.44$ and $d(X Y) = d(Y X) = 0.0916$	103
5.10	Effect of Correlation on the Measurement Rate Region.	103
5.11	Spatial Coupling Wave for A Linearly Correlated Source with $\rho_x = 1.1d(X)$ and $\rho_y = 1.1d(Y X)$	104
5.12	Spatial Coupling Wave for A Linearly Correlated Source with $\rho_x = 1.1d(X)$ and $\rho_y < d(Y X)$	105
5.13	Phase Transition Boundary for MAMP and Comparison with SE Prediction.	105
A.1	EPI gap (A.4) for i.i.d. integer-valued random variables	119

List of Tables

2.1	Summary of Notations	17
3.1	Summary of Notations	30
4.1	Summary of Notations	67
4.2	Duality between H and d	71

1

Introduction

In most of the applications in signal processing, one deals with the acquisition or processing of a usually high-dimensional signal x that can be modeled as a vector in an abstract linear space. For example, for discrete-time signals, this linear space can be \mathbb{R}^n for finite-dimensional signals or $\ell^2(\mathbb{Z})$ for infinite-dimensional finite-energy sequences, where $\ell^2(\mathbb{Z})$ is the space of all real-valued sequences $x : \mathbb{Z} \rightarrow \mathbb{R}$ with finite energy, i.e., $\sum_{n \in \mathbb{Z}} x_n^2 < \infty$. For continuous-time signals, the appropriate linear space can be $L^2(\mathbb{R})$, the space of all real-valued functions over \mathbb{R} , i.e., $x : \mathbb{R} \rightarrow \mathbb{R}$ with a finite energy $\int x(t)^2 dt < \infty$. Albeit being very high dimensional, most of the time x is a very structured signal and has an inherently low “*information content*”. More precisely, it has a very low-dimensional representation in some appropriate basis depending on the application. For example, a collection of images taken by a high-resolution photograph might have only a few dominant components in the frequency domain.

The first step of processing the signal is to appropriately sample or measure the signal by using suitable sensors or measurement devices. The suitable sampling procedure highly depends on the application. For example, an air-conditioning system can use a thermal sensor and sample the output of this sensor every second or every minute by an ADC (analog-to-digital converter), in order to capture the room temperature. In a camera the sampling is done by using an array of photo sensitive sensors to capture a scene. In an optical device such as a telescope, the suitable sampling might need applying an array of mirrors and lenses to obtain a high resolution image of the sky. A trivial but costly approach to capture the signal is to measure or sample all of its components separately. I.e., the number of measurements is equal to the dimension of the signal. For example, in a camera this necessitates using one sensor per an image pixel. Therefore, the acquisition task can be very expensive specially for higher dimensions. In particular, because of the underlying structure of the signal, the resulting measurements are highly redundant which implies the inefficiency of the sampling process.

Compressed sensing (CS) is a new paradigm in signal processing for efficient data acquisition and feature extraction [1–4]. In contrary to what the name suggests, compressed sensing is not involved with the compression but it mostly deals with




Figure 1.1 – Compressed sensing problem. One takes measurements from a high-dimensional signal x via an $m \times n$ measurement matrix A . The possible additive measurement noise is denoted by z .

efficient sampling or acquisition of the signal. The main idea is to optimally sample the signal up to its information content by taking suitably designed linear measurements and faithfully reconstruct it by running low-complexity recovery algorithms. This results in a tremendous saving in the sensor cost because one can use less number of sensors and at the same time keep the desired performance. For example, one can decrease the number of sensors in a camera without sacrificing the quality of the resulting image. The drawback is that the sampling is more sensitive to noise and mismatch. Moreover, one needs more complex sampling devices and numerically stable algorithms in order to recover the initial signal from the measurements. An illustration of the compressed sensing problem is given in Figure 1.1. The goal is to capture an n dimensional signal x by taking a vector y of $m \ll n$ possibly noisy measurements. The signal x is not completely arbitrary but it has some underlying structure, e.g., it has only a few nonzero components as depicted in the figure.

1.1 Linear Inverse Problems

A more general but closely related problem to compressed sensing is the “*Linear Inverse Problem*” (LIP). A linear inverse problem is a general framework that is used to infer some information about a physical object or system by taking linear measurements [5]. For example, a famous equation in physics connecting the mass density and the gravitational field intensity is given by:

$$g(r) = \int_{r'} \frac{\rho(r')(r - r')dr'}{|r - r'|^3}, \quad (1.1)$$

where $g(r)$ is the gravitational field at the observation point r and ρ is the mass density all over the space. One can consider this equation as a linear operator with mass density ρ as the input and the field intensity g as the output. What the name ‘*linear*’ implies in this example is that the equation connecting the observations $g(r)$ and the parameters ρ is linear. If one has measurements of the Earth’s gravity field, then one might ask the question: “Given the available data (measurements), what can one say about the mass density distribution of the Earth in a specific area?” The solution to this problem, i.e., the density distribution that best matches the data, is

useful because it generally tells us something about a physical parameter that we cannot directly observe. The term ‘*inverse*’ is more dominant because the direct problem is frequently straightforward and easy to solve, for example by knowing the mass density ρ one can easily find the gravitational intensity everywhere by a simple integration. Solving the inverse part is however more challenging and one needs to suitably combine the information obtained from the measurements by the structure or the underlying model for data generation in order to recover the initial signal or to find the associated parameters. What distinguishes the compressed sensing problem from a general LIP is that, in the compressed sensing one has more flexibility in designing measurement matrices or in general measurement operators in order to take more informative measurements from the signal. This is not generally possible in a typical linear inverse problem. For example, in Equation 1.1 the measurement process is done by a fixed integral operator connecting the input data ρ to the observation g , whereas in Figure 1.1, the measurement matrix A is at our hands to design.

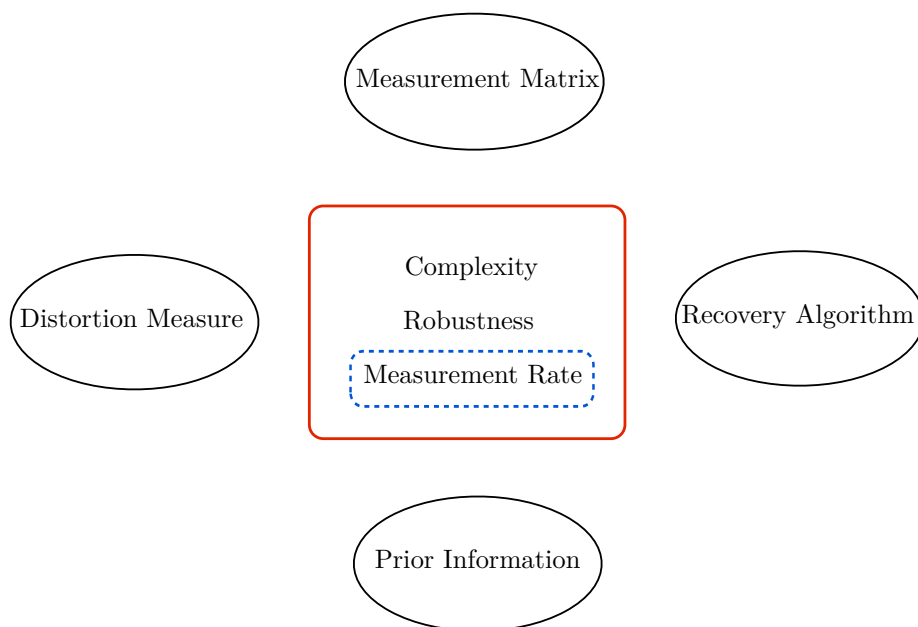


Figure 1.2 – Different aspects of a Linear Inverse Problem

1.1.1 Main Essence of a Linear Inverse Problem

A simple literature review reveals the vast applications and extensive variety of LIPs in general and the compressed sensing in particular that has been studied recently. Apparently this makes it really difficult to give a comprehensive picture of the different aspects of the problem. However, we have tried to, at least, summarize the main essence in Figure 1.2:

- **Prior Information:** Usually, in a linear inverse problem, one has a suitable

model or prior information about the signal. For example, most of natural signals such as images, and sounds have a sparse representation in the frequency or in the wavelet domain. Signals received in an airport radar are sparse with a few spikes corresponding to arriving or departing airplanes; and images from the night sky or deep-space photos from galaxies are sparse images with a dominant dark background and a few sparsely illuminated points showing the position of stars or planets.

Although sparsity is the dominant and, of course, one of the most useful models in signal processing, the LIP is not necessarily restricted to sparse signals. An interesting and recently well-studied model is the graphical model for signal [6–8]. It allows to further extend the usual sparsity model to more comprehensive hierarchical or group sparsity models represented over a graph with many applications and interesting theoretical results [9, 10]. Another interesting application of the graphical models is in probabilistic graphical models, a marriage of probability theory with graphical models, where it is assumed that the samples of the signals are generated by a probability distribution structurally consistent with a given graph [11, 12].

- **Distortion Measure:** In a general LIP, one might be interested in recovering a specific feature of the signal via the measurements, instead of the whole signal. For example, consider the sparse signal received in an airport traffic radar. An air traffic agent only needs to know the signal support (positions of the nonzero components) because it roughly specifies the position of the traveling airplanes. In this case, probably knowing the amplitude of the signal received from each airplane is not so important because it might drastically change due to the shape, orientation or distance of the airplanes from the receiving radar. On the contrary, for a military radar the received amplitude of the signal is as important as its support because it allows to identify the target by estimating its shape or structure. Mathematically speaking, one can assume that there is a distortion measure $d(x, \hat{x})$ between the true signal x and the recovered signal \hat{x} and one adjusts the number of measurements and the recovery algorithm to achieve an acceptable distortion.
- **Measurement Matrix and Recovery Algorithm:** Depending on the signal model and the distortion measure, one needs suitable matrices to encode relevant features in measurements taken from the signal so that the recovery algorithm can successfully estimate the signal. However, there are also other factors that should be taken into account such as the ease of implementation, and physical constraints on measurement devices, which might restrict the choice of the measurement matrix. For example, assuming a sparsity model for the signal, it has been shown that random measurement matrices, such as random Gaussian matrices with i.i.d. entries, provide a close-to-optimal measurement rate under robust and low-complexity recovery algorithms (implemented as a convex optimization problem) [3, 4]. Although random matrices are interesting from an analytic point of view and usually require a minimum number of measurements, implementing them in real-world applications might not be easy and one needs to use more structured matrices. There are similarities with coding theory, where a random code achieves information-theoretically

optimal rates but one has to look for more structured codes to facilitate the implementation and control the decoding complexity. Moreover, there are other issues that might prevent the use of random matrices, such as storage limitation for the measurement matrix, robustness problem in implementing the matrix resulting in measurement mismatches, computational complexity of the recovery algorithm, etc. There have been several attempts to build structured matrices with a performance close to optimal. In [13], $\{0, 1\}$ matrices based on the random expander graph construction were proposed for compressed sensing. It was shown that the iterative peeling decoding can be used to recover sparse signals with a very high probability and with an optimal measurement rate. There has been other constructions based on the *restricted isometry property* (RIP) criterion that also use ideas from coding theory; see [14–17] and the references therein.

1.1.2 Challenges in Application

As depicted in Figure 1.2, at the heart of an LIP, specifically when applied to real-world applications is robustness, complexity and measurement rate that we explain briefly.

- **Robustness:** Assume that $a \in \mathbb{R}^n$ and that one needs to take the measurement $\langle a, x \rangle$ by projecting the signal x on the measurement vector a . For example, this can be done by adjusting the position and the angle, or by masking specific mirrors in an optical measurement device according to the components of a . In practice, one essentially gets $\langle \hat{a}, x \rangle$, where \hat{a} is only an approximation of a . The mismatch results due to imperfect adjustment of the components or use of uncalibrated devices. In general, after the acquisition step, instead of the desired measurements $y = Ax$, one obtains $\hat{y} = \hat{A}x + z$, where \hat{A} is a mismatched measurement matrix and z is the unavoidable additive measurement noise. It is not difficult to see that the measurement imperfections become more problematic for high-dimensional settings because they generally scale like $\|A - \hat{A}\|_1 \approx O(nm)$, quadratically in the dimension n , in contrast to the measurement noise effect which essentially scales linearly in n . In particular, contrary to the noise effect, the mismatch effect can not be moderated by taking extra measurements. The mismatch problem usually precludes the use of random Gaussian matrices because very high precision measurement devices are required for their implementation. Therefore, one needs to use more structured measurement matrices for example $\{0, 1\}$ -valued matrices that can be implemented as on-off pattern of the measurement devices which can be implemented with a very high precision on optical devices. In some applications, $\{+1, -1\}$ -valued matrices are more suitable because they can be mapped to a $\{0, \pi\}$ phase shift pattern in electronic devices.
- **Complexity:** After taking the measurements, in order to recover the initial signal, one needs to combine the information provided with the measurements with a good regularization provided by the prior information or the model. The solution is usually an optimization problem that can be implemented in software. For example, for compressed sensing of sparse signals, one can use

the convex ℓ_1 -minimization algorithm [1–4]

$$\hat{x}(y) = \arg \min \|w\|_1 \quad \text{subject to} \quad \|y - Aw\|_2^2 \leq \epsilon, \quad (1.2)$$

to recover the signal, where ϵ is an estimate of the noise power. Being a convex optimization problem, this can be solved very efficiently by well-developed convex optimization tools. There are also other recovery algorithms that have been proposed with different performance guaranties. A close look to all these algorithms shows that their run time grows rapidly with the dimension, prohibiting their use for high-dimensional signals. Recently, this has motivated many researchers to look for better and faster recovery algorithms in order to extend the compressed sensing to high-dimensional setups. There have been different attempts such as using stochastic optimization algorithms [18, 19] or developing low-complexity message-passing algorithms [20, 21] to speed up the recovery, but essentially it seems that to get any further improvement in computational complexity, one inevitably needs to impose some kind of ‘*structure*’ on the measurement matrices. For example, in [13] the $\{0, 1\}$ -valued adjacency matrix of a bipartite random expander graph was used as the measurement matrix. It was shown that the expander structure of the graph can be used to recover the nonzero components of the underlying sparse signal via a fast iterative peeling algorithm which was much faster than the traditional convex optimization.

Another important factor is the space complexity or the amount of memory required to implement the recovery algorithm. For example, in order to save an unstructured measurement matrix such as a random Gaussian matrix, one needs a memory of size $O(nm)$. Even, in the $\{0, 1\}$ -valued matrix example that we mentioned, although the recovery algorithm runs very fast but the requirement is to implement the adjacency matrix as a graph data structure on the software in order to efficiently recover the nonzero components. Hence, the space complexity might still be a problem for high-dimensional applications.

- **Measurement Rate:** Usually the main *figure of merit* for assessing the performance of signal acquisition and reconstruction is the required measurement rate, i.e., the number of measurements needed per signal dimension to achieve the desired performance. There are a few reasons for this. For example, in an MRI imaging the time that patient should spend inside the imaging instrument is proportional to the number of measurements to be taken, thus for medical convenience, it is reasonable to reduce this time as far as possible. Another reason reducing the measurements could be important is the unstationary behavior of the signal. More precisely, in most of the applications, the measurement process is a sequential and time-consuming process and during this process, the desired signal can change significantly. To cope with this problem, one needs to take all the required measurements as soon as possible. This will be more facilitated if one can reduce the measurement rate. There are also other theoretical interests to characterize the achieved performance (such as recovery distortion) in terms of measurement rate because it provides invaluable engineering insights and some rule of thumb measures for design purposes.

There is no doubt that the picture we provided is a very rough and imperfect view of what happens in reality, but it might be sufficient to point out some of the main difficulties that arise in applying the pure theoretical results to more realistic setups in applications. In particular, another glance at all the parts we mentioned, in particular robustness and complexity issues, reveals the importance of the *structured matrices* in order to obtain robust measurements and low-complexity recovery algorithms.

1.2 Contribution of this Thesis

In previous sections, we explained the main essence of the compressed sensing problem and all the benefits that it provides compared with the traditional sampling and data acquisition. We also outlined some of the issues that arise in applying the compressed sensing to real-world problems. In particular, we emphasized the inherent role that the structured measurements play in order to obtain more robust implementations along with low-complexity recovery algorithms. These are the main requirements in order to extend the compressed sensing to high-dimensional signals that become ubiquitous recently.

This thesis includes three main parts that can be summarized as follows:

- Constructing deterministic and structured partial-Hadamard matrices for compressed sensing and designing low-complexity algorithms for recovery
- Using probabilistic model for the signal and applying tools from probability, coding theory and information theory to rigorously analyze the performance of the proposed algorithms for compressed sensing
- Extending the traditional (single-terminal) compressed sensing to a distributed (multi-terminal) setting

In this section, we will explain these parts in more detail. In particular, we try to emphasize that by the approach taken in this thesis, one can expect to solve some of the issues that naturally arise in high-dimensional setups. Moreover, under some mild assumptions on the signal model, which is not difficult to meet in applications, we use tools from other areas such as information theory and coding theory to assess the performance of the resulting constructions and the proposed recovery algorithms.

1.2.1 Hadamard Matrices

In this thesis, we emphasize the importance of the Hadamard matrices and their efficiency for compressed sensing. In particular, we show that under some mild assumptions on the signal that can be met in applications, using Hadamard matrices can provide close to optimal performance (e.g., measurement rate or recovery distortion). We propose a new matrix construction for compressed sensing using the Hadamard matrices. Specifically, we choose some of the rows of a given Hadamard matrix deterministically according to an information-theoretic metric, and use the resulting matrix as a measurement matrix for compressed sensing. We call the resulting matrices partial Hadamard matrices. Interestingly, the resulting construction can be applied for the compressed sensing of sparse signals which is the dominant

model used in signal processing. The requirement is that the position of the nonzero components in the signal be uniformly random. We use different techniques from information theory and coding theory to analyze the performance of the resulting constructions.

We design deterministic partial Hadamard matrices for two sparsity regimes: sub-linear and linear. In sub-linear sparsity regime, we assume that the number of nonzero components of the signal $K = O(N^\alpha)$ scales sub-linearly with the dimension N for some $\alpha \in (0, 1)$ whereas in the linear regime $K \approx N\delta$ is proportional to N with some $\delta \in (0, 1)$. A summary of the results proved for these two regimes is as follows:

Sub-linear Regime: We prove that by using the underlying properties of the Hadamard matrices, it is possible to deterministically choose some of the rows of a Hadamard matrix for compressed sensing. Furthermore, assuming that the position of the nonzero components is uniformly random, we prove that the recovery of the nonzero components of the signal can be formulated as a belief-propagation algorithm over a sparse graph. We use tools from probability and coding theory to rigorously analyze the performance of the resulting algorithm. In particular, we show that the proposed construction needs to select only $O(K \log_2(\frac{N}{K}))$ rows of a Hadamard matrix, and the recovery algorithm decodes all the nonzero components with a computational complexity $O(K \log_2(K) \log_2(\frac{N}{K}))$ and with a very high probability.

Linear Regime: We formulate the compressed sensing problem from an information theoretic point of view. We prove that in order to capture the information of a memoryless source, one needs to select some specific rows of a Hadamard matrix according to some information-theoretic metric (Rényi information dimension). We use results from coding theory and more recently developed polarization theory to analyze the proposed construction asymptotically for large block-lengths. Using the resulting partial Hadamard matrices is beneficial from different aspects. We now briefly overview some of these advantages.

1. It is not necessary to save the resulting partial Hadamard matrices because they can be easily generated by a closed form formula. One only needs to keep the index of the selected rows which is much easier. To explain more, consider a Hadamard matrix H of order $N = 2^n$ and for any number $i \in \{0, 1, \dots, N - 1\}$, let i_1, i_2, \dots, i_n denote its binary expansion. Then it's possible to show that the i - j -th component of a Hadamard matrix is given by $H_{ij} = (-1)^{\langle i, j \rangle}$, where $\langle i, j \rangle = \sum_{k=1}^n i_k j_k$. This implies that by knowing the index of the selected rows, one can simply generate the resulting partial Hadamard matrix. Compared with a random matrix, such as random Gaussian matrix or even a random binary-valued matrix, where one needs to save all the elements, this results in a huge saving in the memory of the computer software specially for high-dimensional settings.
2. Let x be the signal that we are interested to capture and let $y = Ax$ be the measurements taken from x via the measurement matrix A . In order to obtain the signal from the measurements, it is necessary to run a recovery algorithm. In almost all the recovery algorithms in compressed sensing (e.g., ℓ_1 minimization, Lasso, AMP, MP and so on), one has the *matched-filter* phase of computation, where it is necessary to compute A^*y , i.e., to compute the correlation of the

columns of the measurement matrix with the measurements y . The recursive structure of the Hadamard matrices allows us to compute A^*y faster than the traditional matrix multiplication by at least a factor of $O(\frac{N}{\log_2(N)})$, where N denotes the dimension of the signal x . For a typical dimension $N = 10^3$, this is around 100 times faster. As we mentioned in Section 1.1.2, nowadays, one of the main difficulties in using compressed sensing in real-world applications is the dimensionality scaling problem. In other words, for high dimensional signals, the computational complexity of most of the algorithms prohibits their use for recovery. By replacing the unstructured measurement matrices by structured partial Hadamard matrices, one can get a huge gain in computational complexity and much better dimensionality scaling.

3. Hadamard matrices are $\{+1, -1\}$ -valued matrices. Generally speaking, they can be more easily implemented on sensors or measurement devices. In particular, compared with the unstructured matrices such as random Gaussian matrices, their implementation can be more resilient to mismatch. This makes them a favorable option specially for high-dimensional setups.
4. Our simulation results show that Hadamard matrices give a close-to-optimal performance like random Gaussian matrices. For example, a simple comparison shows that for a similar measurement rate, they have equal and sometimes slightly better performance than the Gaussian matrices in terms of the recovery distortion.

1.2.2 Probabilistic Model for the Signal

Traditionally, the compressed sensing theory was developed assuming a sparsity model for the signal. Specifically, it is assumed that the signal has only a few number of nonzero components. This is a reasonable model for most of applications in signal processing. However, in order to give an information-theoretic flavor to the compressed sensing problem, we use a probabilistic prior for the signal. More precisely, we assume that the components of the signal are generated by a given probability distribution. This model is rich enough to cover most of the applications in signal processing. In particular, it encompasses a probabilistic relaxation of the traditional sparsity model. The requirement is that the support (the position of the nonzero components) of the signal be uniformly random. This is not so restrictive because one can randomly shuffle the columns of any measurement matrix before taking the measurements. This results in a uniformly random support for the signal. The benefit of the resulting probabilistic model is that it allows for an exact information-theoretic analysis of the compressed sensing problem. Specifically, this is used to design partial Hadamard matrices that achieve the information-theoretic limits. Moreover, the probabilistic model of the signal allows to extend the traditional compressed sensing problem to a distributed case that we will explain in the next section.

1.2.3 Distributed Compressed Sensing

One of the advantages of the probabilistic model for the signal is that it can be naturally extended to a multi-terminal or distributed setting, where one has a distributed signal (e.g., temperature, humidity, etc.) that is observed via several

terminals. For example, in a sensor network, a collection of sensors are spread across a field in order to measure a distributed signal. It is always important to acquire the signal by taking as few measurements as possible. One can consider each sensor as a terminal or an observation point. Naturally, there is an inherent correlation among the signals in different terminals. The main challenge in this case is to build a good distributed model that can capture the underlying correlation between the terminals. There have been several attempts to analyze this problem by extending the traditional sparsity model to the multi-terminal setting (specially [22, 23]). Specifically, there was an attempt to make a connection between multi-terminal compressed sensing and the distributed source coding (Slepian-Wolf) counterpart in information theory; refer to [24] for extra references. To explain more, let (X_1, X_2, \dots, X_t) be a multi-terminal memoryless source with a probability distribution $p(x_1, x_2, \dots, x_t)$, where all the X_i have a common finite alphabet \mathcal{X} . Suppose each terminal i encodes its message with rate R_i and sends the encoded message to a data fusion center. One is interested to recover the multi-terminal signal from the received encoded messages. Distributed source coding theorem [25] states that, for large block-lengths, the distributed signal is recoverable from the encoded messages with a negligible distortion, if and only if for any subset of terminals $S \subset \{1, 2, \dots, t\}$,

$$\sum_{i \in S} R_i \geq H(X_S | X_{S^c}),$$

where H denotes the discrete entropy and S^c is the complement of S . Intuitively, this implies that the amount of the information received from each subset of the terminals must be larger than the innovation of those subset given all the rest. In particular, for a simplified three-terminal case, this implies the following constraints on the encoding rates:

$$R_1 \geq H(X_1 | X_2), \quad R_2 \geq H(X_2 | X_1), \quad R_1 + R_2 \geq H(X_1, X_2).$$

We study the counterpart of this problem in compressed sensing by assuming a probabilistic model for the signal to imitate the correlation among the different terminals. As a result, we obtain an information-theoretic characterization of the required measurement rate very similar to the distributed source coding problem. For simplicity, consider a two terminal memoryless source (X_1, X_2) with a joint probability distribution $p(x_1, x_2)$. We show that, under a mild assumption on the distribution p , the two terminal source is recoverable from a set of measurements from both terminals if and only if the following constraints are satisfied:

$$\rho_1 \geq d(X_1 | X_2), \quad \rho_2 \geq d(X_2 | X_1), \quad \rho_1 + \rho_2 \geq d(X_1, X_2),$$

where ρ_1 and ρ_2 are the measurement rates, the number of measurements per dimension of the signal, in the first and the second terminal and d denotes the joint or the conditional Rényi information dimension of the signal. Hence, the probabilistic model of the signal allows to information-theoretically characterize the required measurement rate in terms of signal parameters. In particular, in the compressed sensing setting, the Rényi information dimension plays a role similar to entropy in the distributed source coding problem. Definitely, it is difficult to realistically model or efficiently approximate the complicated correlation among the different terminals in a real-world scenario. However, the probabilistic modeling of the multi-terminal signal

is rich enough to reasonably approximate the spatial correlation among different terminals. The insight gained by this approximation can be helpful for designing efficient data acquisition in such complex settings.

1.3 Applications beyond Compressed Sensing

As we explained in Section 1.2, our main goal in this thesis is to build deterministic and structured measurement matrices to solve some of the main issues that arise while applying the compressed sensing to real-world problems. To do so, we essentially select a specific set of the rows of Hadamard matrices according to a deterministic criterion. Very briefly, we emphasize that essentially in any compressed sensing application, as far as the measurement device allows the implementation, one can replace the unstructured matrices by our proposed partial Hadamard matrices. This solves some of the practical issues such as robustness, and reduces the computational complexity dramatically while keeping a close to optimal performance in terms of the required measurement rate.

Since Hadamard matrices have many applications in signal processing, coding theory and wireless communication, one might wonder if the tools developed in this thesis have applications in these other areas. In this section, we deviate from the main theme of the thesis to explain a potential application of the results developed in the thesis to wireless and mobile communication.

We start from a very brief introduction to mobile communication, specially the Code Division Multiple Access (CDMA) technique vastly applied in cellular networks. In a mobile network, the geographical area decided for communication is typically divided into cells and there is a base station per cell. When a mobile device physically moves inside the communication area, depending on the strength of the signal it receives from different base stations, it is automatically allocated to a specific base station. The set of all base stations are typically connected to a big central hub which is responsible for all necessary computations that make the communication possible such as packetizing data, controlling possible channel errors, identifying and managing the mobility and handover of the mobile devices in the network, etc. More importantly, the cellular structure allows frequency reuse in nonadjacent cells, provided that the resulting interference is kept below some threshold by a suitable design of the network.

To make the discussion as simple of possible, we restrict ourselves to a specific base station which is responsible for signal communication inside its cell. Typically, there are many users inside this cell and some kind of Multiple Accessing (MA) technique is required to allocate the finite communication capacity of the base station among the mobile users. In CDMA, a distinct code is assigned to each user in a cell by means of which the user can send its information in a shared media. Very briefly, assume that we have K users inside the cell and let us denote the code for each user by c_1, \dots, c_K . In the discrete time model for the communication channel, without loss of essential generality, one can assume that these codes are vectors in \mathbb{R}^N , where the code length N is proportional to the time required to transmit one information symbol. The base station transmits the signal s_1, \dots, s_K for users 1 up to K by simply multiplying the signal for each user by its associated code. More precisely, the received signal for user i is given by $R_i = \sum_{j=1}^K s_j c_j + Z_i$, where Z_i is

the additive white Gaussian noise for user i .

To decode its desired signal, the user i uses a matched filter receiver that takes the inner-product or correlation of the received signal with the code \mathbf{c}_i . More precisely, it computes $r_i = \langle \mathbf{c}_i, R_i \rangle$ and makes decision about s_i based on r_i , where

$$r_i = \sum_{j=1}^K s_j \langle \mathbf{c}_i, \mathbf{c}_j \rangle + \langle \mathbf{c}_i, Z_i \rangle = s_i \|\mathbf{c}_i\|^2 + \sum_{j=1, j \neq i}^K s_j \langle \mathbf{c}_i, \mathbf{c}_j \rangle + \langle \mathbf{c}_i, Z_i \rangle. \quad (1.3)$$

A glance at Equation (1.3) shows three contributing terms: the desired signal, the interference, and the additive channel noise. In a traditional design, the codes for different users are selected to be orthogonal, which implies that the interference term vanishes. Moreover, the orthogonality constraint essentially implies that there can not be more than N active users in a cell (recall that N is the code length). There have been different classes of codes that have been designed and frequently used in implementations and in standards such as Gold codes, Kasami codes, and Walsh-Hadamard codes to name a few [26].

To make a connection with the compressed sensing problem, let us denote by $\{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N\}$ the full set of codes that the base station can allocate to different users and let us denote by $A = [\mathbf{c}_1 \mathbf{c}_2 \dots \mathbf{c}_N]$ the matrix obtained by concatenating these codes. We denote by $\theta_j \in \{0, 1\}$ the code occupancy inside a cell, i.e., $\theta_j = 1$ if the code j is allocated to some active user and $\theta_j = 0$ otherwise. We define the whole signal of the cell by an N dimensional column vector x whose j -th component is given by $x_j = \theta_j s_j$. In particular, if the code j is not assigned to a user, we have $x_j = 0$. Typically, the number of active users in a cell is less than N , which implies that the vector x is usually a sparse vector. Therefore, finding the signal s_i is equivalent to solving the sparse equation $R_i = Ax + Z_i$, where R_i denotes the noisy received signal of the user i .

Under the orthogonal code assumption, matrix A is an orthogonal matrix, thus it is invertible. However, we can take advantage of the compressed sensing technique, in the sense that instead of keeping the whole matrix A , we only need to keep a specific set of rows of A to recover all the nonzero components since x is sparse. This is equivalent to shortening the codes allocated to different users which equivalently increases the transmission rate. If we use Walsh-Hadamard codes for multiple accessing inside the cell, the problem is completely transformed to the setting that we have in this thesis, namely, we select a specific set of rows of a Hadamard matrix according to a deterministic criterion, which directly provides a collection of shorter codes (shorter than N) for signal transmission. More importantly, the required length of the code (equivalent to the number of measurements in the compressed sensing setup) is deeply related to the number of active users (equivalent to the sparsity $\delta = \frac{K}{N}$ of the vector x with K being the number of active users), which is typically known to the base station and can be made known to all the users.

There are also other nice properties of our proposed construction that can be used to build an adaptive transmission scenario. More specifically, in Chapter 4, we propose a deterministic Hadamard construction that depends on the source sparsity. As we emphasize there, the set of constructed matrices has an embedding property with respect to the source sparsity. To be more precise, let $0 < \delta_1 < \delta_2 \leq 1$ and let H_1 and H_2 be partial Hadamard matrices constructed for sparsity δ_1 and δ_2 respectively. The embedding property implies that all the rows of H_1 are included

among the rows of H_2 . This essentially means that one can order all the rows of a Hadamard matrix based on their importance. Looking from a communication point of view, this implies that one can build hierarchical codes, namely, the base station starts the communication by sending more important sections of the code and refines them further and further by sending the less important parts until all the users get enough information to resolve the interference caused by the other active users. Interestingly, this also provides some adaptive power allocation scheme in the sense that the more the refinements are transmitted to the users, the more power is delivered to them.

There are also some issues that should be mentioned. For example, in this scenario, it is implicitly assumed that each user knows all the codes (the codebook) exploited in a cell, which might not be the case in reality. Furthermore, every user jointly decodes its signal along with the signals from all the other users. This might create some privacy issues that should be appropriately managed.

In this brief section, our main goal was to give a very simple example to show some of the potential applications of the Hadamard constructions proposed in this thesis beyond the compressed sensing setting.

1.4 Outline of the Thesis

In this section, we briefly overview the structure of the thesis by explaining each chapter separately.

Chapter 2

In this chapter, we present a Hadamard construction for the compressed sensing of memoryless integer-valued sources, which is a good model for those applications in signal processing that deal with quantized signals. Compressed sensing of integer-valued signals allows to extend the linear compression problem from finite alphabet sources to infinite alphabet ones, which is also of independent theoretical interest. To analyze the problem and characterize the ultimate compression bounds for this class of signals, we develop information-theoretic tools such as entropy power inequality.

Chapter 3

In this chapter, we propose a new Hadamard construction for a more practical signal model. We assume that the number of nonzero components of the signal K scales like $O(N^\alpha)$, sub-linearly with the dimension N for some $\alpha \in (0, 1)$. Moreover, we suppose the position of these nonzero components is uniformly random. Using the underlying properties of the Hadamard matrices, we build a measurement matrix by selecting specific rows of the Hadamard matrices. We also develop a low-complexity algorithm in order to recover the nonzero components of the signal from the measurements.

The proposed construction can be equivalently seen as a collection of linear hash functions and a peeling decoder that iteratively reconstructs the nonzero components from the hash outputs. In particular, using the properties of the Hadamard matrices and assuming a random support model for the signal, we show that one can randomly partition the nonzero components in a collection of bins and each hash function computes the weighted sum of all the signal components that are mapped to the

same bin. We show that the proposed recovery algorithm is equivalent to a belief-propagation algorithm over a specific sparse graph. Fortunately, this allows us to fully analyze the performance of the algorithm for large dimensions by borrowing results from belief propagation analysis over the sparse graphs. We give a full characterization of the required number of measurements, the hash construction and the decoding procedure.

Fortunately, the proposed construction can be applied for computing the Walsh-Hadamard transform of a signal provided that its Walsh-Hadamard transform is sparse, which would be of great interest in many applications in signal processing. In particular, compared with the traditional recursive FFT-like method, the new algorithm computes the Walsh-Hadamard transform much faster and requires much less number of signal samples. The only requirement is that the signal be sparse in the transform domain with a uniformly random support.

Chapter 4

In this chapter, we extend the Hadamard construction to memoryless signals with the number of nonzero components proportional to the dimension of the signal. We call this regime the linear sparsity regime in contrast to the sub-linear sparsity regime studied in Chapter 3. The construction is based on the polarization theory, recently developed in coding theory [27, 28]. We show that the importance of any row among the rows of a Hadamard matrix can be characterized by a parameter called the Rényi information dimension (RID). In particular, we show that for large block lengths this parameter polarizes to the two extremal values 0 and 1, with 1 denoting the most important and with 0 denoting the less important rows. To obtain this result, we develop the properties of the RID in vector setting and related information measures. It is then shown that the RID polarization is obtained with an analytical pattern. In other words, there is no need to rely on algorithms to compute the set of components that tend to 0 or 1, as this is given by a known pattern equivalent to the binary erasure channel (BEC) polarization [27]. This is then used to construct explicit partial Hadamard matrices for compressed sensing. Numerical simulations provide evidence that off-the-shelf recovery algorithms such as ℓ_1 -minimization or approximate message passing (AMP) for compressed sensing can be used in conjunction with the constructed matrices. As discussed in this chapter, using our deterministically-constructed matrices, we achieve the same performance as the traditional Gaussian matrices while reducing the time complexity of the recovery algorithm by approximately $O(\frac{N}{\log_2(N)})$ in block length N . This allows to extend the compressed sensing to higher dimensional signals. Moreover, the simple $\{+1, -1\}$ structure of the resulting measurement matrices facilitates the implementation and reduces the mismatch effect.

Chapter 5

This is the final chapter of the thesis. In this chapter, we extend the traditional (single-terminal) compressed sensing to a distributed (multi-terminal) scenario. In particular, we discuss about different correlations that exist in a multi-terminal signal: the temporal correlation of the signal in each terminal, and the spatial correlation among signals in different terminals. For simplicity, we restrict ourselves

to a memoryless case (independent samples across time), where the components of the distributed signal at each time are generated by a multi-dimensional probability distribution that models the spatial correlation among the different terminals. In particular, we neglect the time correlation among the samples in each terminal. This is a good approximation because the temporal correlation can be moderated by suitable sampling and processing in each terminal and it seems that the spatial correlation plays a more important role. Moreover, the memoryless assumption on the signal allows to give a full information-theoretic characterization of the required measurement rate of every terminal.

Assuming a memoryless model for the distributed signal also allows us to extend the AMP algorithm, developed for the recovery of single-terminal sources, to the distributed setting (MAMP). We prove that for large block lengths the performance of the MAMP algorithm can be fully characterized by a state evolution equation. By analyzing the behavior of this equation, we rigorously specify the rate-distortion region of a multi-terminal compressed sensing problem. In particular, we show that by spatially coupling the measurement matrices and running the MAMP algorithm, we can asymptotically obtain all the measurement rate region predicted information-theoretically.

Hadamard Construction for Discrete Memoryless Sources

2

In this chapter¹, we study the compressed sensing of discrete memoryless sources using Hadamard matrices. To do this, we select a specific set of the rows of a Hadamard matrix to take linear measurements from the signal and use the recursive structure of the Hadamard matrices to recover the signal efficiently (with a very low complexity). The proposed matrix construction is deterministic and only depends on the distribution of the source. We are mainly interested to know how many number of measurements is necessary to recover the source with a negligible distortion. We introduce the problem in Section 2.1 and review some of the related works. We review the polarization theory for source coding in Section 2.2 and make a link between the polarization theory and the Hadamard matrix construction in Section 2.3. Section 2.4 formulates the problem. Section 2.5 explains the proposed deterministic Hadamard matrix construction to capture the information of the source. Section 2.7 includes the simulation results and further intuitions about the problem. Table 2.1 summarizes the notations used in this chapter.

Table 2.1 – Summary of Notations

\mathbb{Z}	the set of integers	$[m]$	$\{1, 2, \dots, m\}$
\mathbb{Z}_+	the set of positive integers	X_i^j	$\{X_i, X_{i+1}, \dots, X_j\}$
\mathbb{N}	the set of strictly positive integers	H	discrete entropy
\mathbb{R}	the set of reals	H_n	Hadamard matrix of order 2^n
\mathbb{F}_q	Galois field of order q	\tilde{H}_n	partial Hadamard matrix

2.1 Introduction and Related Work

One of the most important problems in information and communication theory is source coding or source compression. Different variants of source coding problem have been studied in the literature. The main idea is that any information source

¹This chapter is the result of collaboration with Emmanuel Abbe and Emre Telatar.

has some inherent redundancy which can be used systematically in order to compress the source. Moreover, the initial data is recoverable from the compressed version with a very high probability. Shannon in his seminal paper [29] studied this problem and showed that for discrete memoryless sources over finite alphabets, the ultimate compression bound is given by the entropy of the source.

Arikan, in his break-through work [27, 28], introduced the polarization phenomenon which can be exploited to build optimal linear source and channel codes achieving the information theoretic limits. In particular, in the source coding setting, he showed that based on the polarization theory one can *deterministically* build linear block codes to encode any memoryless binary-valued source up to its entropy by taking linear measurement from the source with arithmetic over the Galois field \mathbb{F}_2 . The result was later extended to memoryless sources over finite field \mathbb{F}_q , where q is a prime number [30].

Although for most of the applications in signal processing, at least those dealing with discrete signals, modeling the information sources over finite fields seems to be sufficient, there are also cases in which the alphabet of the sources is inherently unbounded, where for simplicity one can assume that the source takes its values in the set of integers \mathbb{Z} . One can still use linear polar encoders in order to compress the source with the difference that for the integer-valued source all the arithmetics is done over \mathbb{Z} . One can naturally ask whether a phenomenon similar to the polarization occurs over the integers and if the answer is positive, how one can use this phenomenon to build deterministic linear matrices to capture the information of memoryless integer-valued sources. In this chapter, our goal is to answer these questions.

2.2 Polarization Theory over Finite Field \mathbb{F}_q

Let $G_n = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}^{\otimes n}$, where \otimes denotes the Kronecker product. Assume that X_1^N is a memoryless source with alphabet \mathbb{F}_q , where q is a prime number and $N = 2^n$. Let $Y_1^N = G_n X_1^N$ be the resulting linear measurements with arithmetic over \mathbb{F}_q . For $i \in [N]$, let us define $I_i = H_q(Y_i | Y_1^{i-1})$, where H_q denotes the discrete entropy in base q . The polarization phenomenon studied in [27, 30] states that for any $\delta > 0$, as n goes to infinity

$$\frac{\#\{i \in [N] : I_i \in (\delta, 1 - \delta)\}}{N} \rightarrow 0.$$

As $I_i \in [0, 1]$, this implies that for large n the values I_i , $i \in [N]$, polarize to either 0 or 1. In other words, all the components of the measurement Y_1^N are either highly informative ($I_i \approx 1$) or highly predictable ($I_i \approx 0$). Furthermore, for every $\delta \in (0, 1)$

$$\frac{\#\{i \in [N] : H_i \in (1 - \delta, 1]\}}{N} \rightarrow H_q(X), \quad (2.1)$$

which implies that the number of informative components is approximately $NH_q(X)$. Notice that every measurement Y_i is associated with one of the rows of the matrix G_n and (2.1) indicates that the “measurement rate” (number of measurements per source symbol) required to extract the informative components is close to the entropy of the

source $H_q(X)$ for large N . In a nutshell, this provides a linear source compression scheme, i.e., for $N = 2^n$, assuming that \tilde{G}_n is a $m_N \times N$ submatrix of G_n obtained by dropping the rows corresponding to the highly predictable Y 's

$$E_N : \{0, 1, \dots, q-1\}^N \rightarrow \{0, 1, \dots, q-1\}^{m_N},$$

$$E_N(x_1^N) = \tilde{G}_n x_1^N,$$

is an ensemble of fixed-to-fixed linear source encoders capturing the information of the source with a measurement rate very close to $H_q(X)$. From classical information theory, it is known that asymptotically this measurement rate is the best one can hope for and the polarization theory provides a mathematical tool to design asymptotically optimal linear source encoders.

2.3 What about Integer-valued Sources?

Suppose that we have a source with alphabet size q' , which we assume to be prime for simplicity. To design a linear code for this source one can treat it like a source with a larger alphabet \mathbb{F}_q for sum prime number $q > q'$ with the arithmetic operations induced by \mathbb{F}_q . By doing so and by constructing a polar code for the source over \mathbb{F}_q , one can decrease the required measurement rate from $H_{q'}(X)$ to $H_q(X)$, i.e., by a factor of $\frac{\log(q)}{\log(q')}$. This implies that by working over larger fields, it is possible to pack more information in a specific measurement and decrease the measurement rate. The drawback is that the polarization will happen slowly and one needs to take larger block-lengths N . Going beyond finite alphabets, one can naturally ask what happens if one works over infinite alphabets such as the integers \mathbb{Z} ? Is it possible to get a *zero* measurement rate asymptotically?

In this chapter, we provide a positive answer to this question. More precisely, we prove that, asymptotically, one can fully capture the information of an integer-valued memoryless source by taking linear measurements with a vanishing measurement rate (measurement per signal dimension). However, there are also some caveats that one should be aware of in order to get a reasonable result:

- We only consider linear encoders. Otherwise, since for any $m \in \mathbb{N}$ the cardinality of \mathbb{R}^m is the same as \mathbb{R} , any number of random variables such as $X_1^m \in \mathbb{R}^m$ can be non-linearly but reversibly imbedded in \mathbb{R} , thus, only one *nonlinear* measurement is sufficient for their recovery. Consequently, the asymptotic measurement rate is 0. However, the resulting measurement scheme is highly nonlinear and extremely susceptible to measurement noise.
- If the source alphabet is finite, e.g., \mathbb{F}_q , there is a *magic linear measurement* $(1, q, q^2, q^3, \dots)$ which packs all of the information of the source in only one measurement. Therefore, it seems that the answer to the posed question is trivial. However, similar to the previous nonlinear case, the resulting linear measurement is highly susceptible to measurement noise. Notice that what we are interested to know is the compression power of the Hadamard matrices H_n consisting of rows with $\{+1, -1\}$ values. In particular, the numerical simulations at the end of this chapter show that the resulting matrices seem to be stable to additive measurement noise. For these family of matrices, a simple

argument shows that each measurement can carry at most $O(\log_2(N))$ bits of information, which upon considering the total information $NH(X)$ of the source, gives a measurement rate of at least $O(\frac{1}{\log_2(N)})$ vanishing asymptotically. Thus, it seems that the measurement rate must be 0.

- The problem becomes more challenging if we assume that the source has an arbitrary distribution over the integers, which does not necessarily have a finite support. In this case, it is not even obvious that such a magic linear measurement exists and the problem becomes really interesting.

In the rest of this chapter, we further study this problem and develop mathematical tools in order to analyze it. In particular, we introduce the family of Hadamard matrices and use their underlying properties to design appropriate measurement matrices to capture the information of the signal. We also make a connection with the polarization theory in order to assess the performance of the constructed measurement matrices in terms of their measurement rate.

2.4 Hadamard Matrices and the Entropy Process

Let X be an integer-valued memoryless source. For $n \in \mathbb{N}$, let H_n denote the Hadamard matrix of order $N = 2^n$ defined by $H_n = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}^{\otimes n}$ and let $Y_1^N = H_n X_1^N$ be the measurements taken from the i.i.d. sequence X_1^N . Notice that the Hadamard matrices $\{H_n\}$ and the polar matrices $\{G_n\}$ introduced in Section 2.2 have the same recursive structure and the construction that we proposed for $\{I_n\}$ is the same as source polar codes with the difference that the arithmetic operations are done over the reals. Similar to the source polar codes, we define the process I_n , where for $i \in [N]$, $I_n(i) = H(Y_i | Y_1^{i-1})$ is the conditional entropy of the i -th measurement given the previous ones. One can convert $\{I_n\}$ to a branching process by the following procedure (see Figure 2.1):

1. I_0 is the root of the tree defined by $I_0(1) = H(X)$, where $H(X)$ denotes the discrete entropy of the source.
2. From $I_n(i)$ the process branches to $I_{n+1}(2i-1)$ and $I_{n+1}(2i)$ with equal probability (as shown in Figure 2.1).

With this construction, it is easy to see that for the stochastic process $\{I_n\}$, any random variable such as I_n takes all of its 2^n possible values $\{I_n(i)\}_{i=1}^{2^n}$ equiprobably.

Using the results of source polarization, it is possible to prove the following proposition.

Proposition 2.1. *Let X be a memoryless source with a probability distribution p_X over \mathbb{Z} and let $\{I_n\}$ be its corresponding entropy process (branching process). I_n is a martingale converging to the random variable I_∞ .*

Proof. The martingale property follows from an argument similar to the source polar codes. Moreover, from the positivity of the discrete entropy, it results that I_n is a positive martingale, thus, from the *martingale convergence theorem*, it must converge




Figure 2.1 – Representation of the Entropy Process on a Binary Rooted Tree

to a random variable I_∞ almost surely [31]. Applying Fatou's lemma [31] to I_n , it results that

$$\mathbb{E}(I_\infty) = \mathbb{E}(\liminf_{n \rightarrow \infty} I_n) \leq \liminf_{n \rightarrow \infty} \mathbb{E}(I_n) = E(I_0) = H(X),$$

where we used the martingale property that $\mathbb{E}(I_n) = \mathbb{E}(I_0) = H(X)$. Thus, I_∞ has a finite expected value. In particular, $\mathbb{P}(I_\infty = \infty) = 0$, which implies that I_∞ is finite-valued almost surely. \square

Theorem 2.1 (“Absorption phenomenon”). *Let X be a memoryless source with the entropy martingale I_n converging to I_∞ . Then $\mathbb{P}(I_\infty = 0) = 1$.*

Remark 2.1. *For polar codes over \mathbb{F}_q , the resulting entropy process I_n is bounded in the interval $[0, 1]$. The polarization theorem states that, for large values of n , almost all of the branches of the process are polarized to the boundary points $\{0, 1\}$. However, the process I_n for integer-valued sources has only one extreme point 0 and that is the reason why we call this phenomenon absorption rather than polarization.*

Proof. Let $n \in \mathbb{N}$ and let X_1^N be an i.i.d. sequence from the memoryless source X and let $Y_1^N = H_n X_1^N$ be the resulting measurements taken by the Hadamard matrix H_n . By the definition of the entropy process, one has $I_n(i) = H(Y_i | Y_1^{i-1})$. Let $\{I_{n+1}(2i-1), I_{n+1}(2i)\}$ be the entropy values for time $n+1$ resulted by branching $I_n(i)$ and let \tilde{Y}_1^N be an independent copy of Y_1^N . Using the recursive structure of the Hadamard matrices and using the result for the source polar codes as in [28, 30], one can obtain the following expressions for $\{I_{n+1}(2i-1), I_{n+1}(2i)\}$:

$$I_{n+1}(2i-1) = H(Y_i + \tilde{Y}_i | Y_1^{i-1}, \tilde{Y}_1^{i-1}), \quad I_{n+1}(2i) = H(Y_i - \tilde{Y}_i | Y_1^{i-1}, \tilde{Y}_1^{i-1}, Y_i + \tilde{Y}_i).$$

Applying the chain rule for discrete entropy, one can simply check that

$$\frac{I_{n+1}(2i-1) + I_{n+1}(2i)}{2} = I_n(i),$$

confirming the martingale property of $\{I_n\}$. As I_n converges almost surely, there is an event E with $\mathbb{P}(E) = 1$ that consists of all convergent branches ω , for which $\lim_{n \rightarrow \infty} I_n^\omega$ exists and is a well-defined real number. Let $\omega \in E$ be a fixed convergent branch in E . Then for any $\delta > 0$, there exist $n_0(\delta, \omega)$ such that for all $n > n_0$, $|I_{n+1}^\omega - I_n^\omega| < \delta$. Let i be the branch number of ω at time n , i.e., $I_n^\omega = I_n(i)$ and notice that at time $n + 1$ the branch number of ω can be either $2i - 1$ or $2i$. From the martingale property of $\{I_n\}$, it results that

$$H(Y_i + \tilde{Y}_i | Y_1^{i-1}, \tilde{Y}_1^{i-1}) - H(Y_i | Y_1^{i-1}) < \delta.$$

Using the *Conditional Entropy Power Inequality* derived in Appendix A, one has

$$g(I_n^\omega) = g(H(Y_i | Y_1^{i-1})) \leq H(Y_i + \tilde{Y}_i | Y_1^{i-1}, \tilde{Y}_1^{i-1}) - H(Y_i | Y_1^{i-1}) < \delta, \quad (2.2)$$

which implies that $I_n^\omega < g^{-1}(\delta)$, where g is a universal continuous and strictly increasing function lower bounding the gap between the sum conditional entropy and individual conditional entropy with the property that $g(0) = 0$. As $\delta > 0$ is arbitrary, it results that $I_\infty^\omega = \lim_{n \rightarrow \infty} I_n^\omega = 0$ over all branches $\omega \in E$, which implies that $\mathbb{P}(I_\infty = 0) = \mathbb{P}(E) = 1$. \square

As we explained in Section 2.2, the polarization phenomenon allows to construct linear measurement matrices for source compression. More precisely, if one keeps all the rows of the polar matrix G_n corresponding to the measurements with significant conditional entropy and drops the remaining rows, the resulting matrix is capable of preserving approximately all the information of the source with an asymptotically optimal measurement rate (the entropy of the source). Following a similar procedure as in source polar codes and using the absorption phenomenon proved in Theorem 2.1, in the next section, we will construct appropriate measurement matrices for an integer-valued source and will evaluate their asymptotic measurement rate.

2.5 Deterministic Partial-Hadamard Matrix Construction

Let X be a memoryless source with corresponding entropy process $\{I_n\}$. Notice that for a given source distribution, all the possible values of $\{I_n(i) : i \in [2^n]\}$ can be exactly computed. In order to capture the information of the source, we construct the ensemble of measurement matrices $\{\tilde{H}_n\}$, where \tilde{H}_n consists of those rows of the Hadamard matrix H_n with indices in $S_n = \{i \in [2^n] : I_n(i) \geq \epsilon H(X)\}$, where $\epsilon \in (0, 1)$ is a fixed parameter and $H(X)$ denotes the discrete entropy of the source X . In other words, we keep those rows of the Hadamard matrix which capture a significant amount of the information of the source assuming the access to all the previous measurements. We denote by $m_n = |S_n|$ the number of the measurements or the number of the rows of \tilde{H}_n and define the asymptotic measurement rate of the ensemble by $\rho = \limsup_{n \rightarrow \infty} \frac{m_n}{2^n}$.

Theorem 2.2. *Let X be a memoryless source with a probability distribution p_X over \mathbb{Z} and let $\{\tilde{H}_n\}$ be the ensemble of partial Hadamard matrices designed for some $\epsilon \in (0, 1)$. Then for any n ,*

$$\frac{H(X_1^N | \tilde{H}_n X_1^N)}{H(X_1^N)} \leq \epsilon, \rho = 0.$$

Theorem 2.2 implies that for any $\epsilon \in (0, 1)$ the constructed family can capture more than $1 - \epsilon$ ratio of the information content of the source and has asymptotically zero measurement rate.

Proof. Let $n \in \mathbb{N}$ and let S_n denotes the set of rows selected from H_n to construct \tilde{H}_n . For simplicity, we drop the dependence on n . We have

$$\begin{aligned} \frac{H(X_1^N | \tilde{H} X_1^N)}{H(X_1^N)} &= \frac{H(Y_1^N | Y_S)}{NH(X)} = \frac{H(Y_S, Y_{S^c} | Y_S)}{NH(X)} = \frac{H(Y_{S^c} | Y_S)}{NH(X)} \\ &\leq \frac{\sum_{i \in S^c} H(Y_i | Y_1^{i-1})}{NH(X)} \leq \frac{|S^c| \epsilon H(X)}{NH(X)} \leq \epsilon, \end{aligned}$$

where $Y_S = \{Y_i : i \in S\}$ and $S^c = [N] \setminus S$ denotes the complement of S . To prove the other part, notice that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{m_n}{2^n} &= \limsup_{n \rightarrow \infty} \frac{|S_n|}{2^n} \stackrel{(a)}{=} \limsup_{n \rightarrow \infty} \mathbb{P}(I_n \geq \epsilon H(X)) \\ &\stackrel{(b)}{\leq} \mathbb{P}(\limsup_{n \rightarrow \infty} I_n \geq \epsilon H(X)) = \mathbb{P}(I_\infty \geq \epsilon H(X)) = 0, \end{aligned}$$

where (a) follows from the uniform probability assumption on the branches $\{I_n(i) : i \in [2^n]\}$ and (b) follows from reverse Fatou's lemma. In particular, this shows that the asymptotic measurement rate of the ensemble is 0. \square

2.6 Lower Bound on the Rate of Absorption

Let $\{I_n\}$ be the entropy process of a memoryless source X . Recall that $I_n(i) = H(Y_i | Y_1^{i-1})$ for $i \in [N]$. From the chain rule for the discrete entropy, one can check that $\sum_{i \in [N]} I_n(i) = NH(X)$. Moreover, $I_n(i) = H(Y_i | Y_1^{i-1}) \leq H(Y_i)$.

Proposition 2.2. *Let Y_1^N be the measurement sequence as defined before. Then for every $i \in [N]$,*

$$H(Y_i) \leq \frac{1}{2} \log_2 \left\{ (2\pi e) \left(N\sigma_X^2 + \frac{1}{12} \right) \right\},$$

where σ_X^2 denotes the variance of the source X and the discrete entropy is in bits.

Proof. The proof follows from the result of Problem 8.7 of Cover and Thomas [32], which states that for a discrete random variable Z taking values in the set $\mathcal{Z} = \{a_1, a_2, \dots\}$ with $\mathbb{P}(Z = a_j) = p_j$,

$$H(Z) = H(p_1, p_2, \dots) \leq \frac{1}{2} \log_2 \left\{ (2\pi e) \left(\sum_{j=1}^{\infty} p_j j^2 - \left(\sum_{j=1}^{\infty} j p_j \right)^2 + \frac{1}{12} \right) \right\}.$$

In particular, for the integer-valued random variable Y_i , one obtains that

$$H(Y_i) \leq \frac{1}{2} \log_2 \left\{ (2\pi e) (\text{Var}(Y_i) + \frac{1}{12}) \right\} = \frac{1}{2} \log_2 \left\{ (2\pi e) (N\sigma_X^2 + \frac{1}{12}) \right\},$$

where we used the fact that $Y_i = \sum_{j \in [N]} h_{ij} X_j$ with $h_{ij} \in \{+1, -1\}$ being the components of the i -th row of the Hadamard matrix H_n , thus from the i.i.d. assumption, $\text{Var}(Y_i) = \sum_{j \in [N]} h_{ij}^2 \text{Var}(X_j) = N\sigma_X^2$. \square

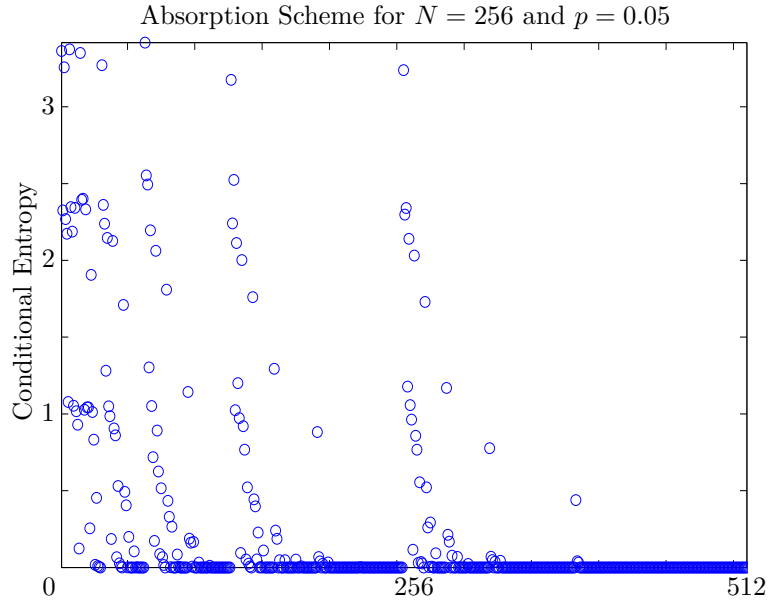


Figure 2.2 – Absorption phenomenon for a binary source with $p(1) = 0.05$ and $N = 512$. It is seen that almost all the entropy values are absorbed to 0.

Proposition 2.3. *Let $\{I_n\}$ be the entropy process of a memoryless source X . Suppose $\epsilon \in (0, 1)$ and let $S_n = \{i \in [N] : I_n(i) \geq \epsilon H(X)\}$ and $m_n = |S_n|$. Then*

$$\liminf_{N \rightarrow \infty} \frac{m_n}{\frac{2NH(X)}{\log_2(N)}} \geq 1 - \epsilon.$$

Proof. It is easy to check that $\sum_{i=1}^N I_n(i) = NH(X)$ which implies that

$$\begin{aligned} NH(X) &= \sum_{i \in S_n} I_n(i) + \sum_{i \in [N] \setminus S_n} I_n(i) \leq \frac{m_n}{2} \log_2 \left\{ (2\pi e) \left(N\sigma_X^2 + \frac{1}{12} \right) \right\} + (N - S_n)\epsilon H(X) \\ &\leq \frac{m_n \log_2(N)}{2} \left(1 + O\left(\frac{1}{\log_2(N)}\right) \right) + N\epsilon H(X). \end{aligned}$$

Therefore, one obtains that $\liminf_{N \rightarrow \infty} \frac{m_n}{\frac{2NH(X)}{\log_2(N)}} \geq 1 - \epsilon$. \square

2.7 Simulation Results

In this section, we evaluate the performance of the proposed deterministic partial Hadamard matrices via numerical simulation. For simulation, we use a binary random variable X with a probability distribution $p_X(0) = 1 - p$ for some $0 < p \leq \frac{1}{2}$.

2.7.1 Absorption Phenomenon

Figure 2.2 shows the absorption phenomenon for the binary source with $p = 0.05$ and for block-length $N = 512$. As seen from the figure, almost all the entropy values are absorbed to 0.

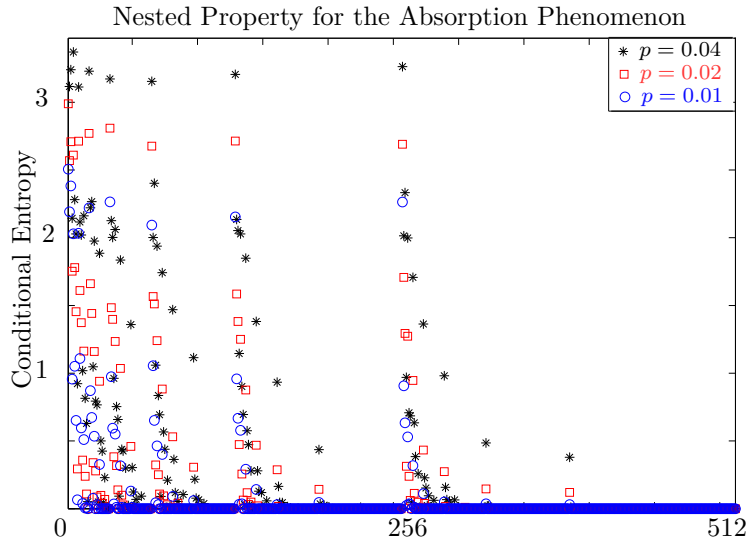


Figure 2.3 – Nested property for the absorption scheme of a binary source for different values of p . It is seen that the entropy process for the larger p dominates the entropy process for the smaller ones.

2.7.2 Nested Property

Absorption phenomenon for $N = 512$ and different values of p is shown in Figure 2.3. It is seen that the set of high-entropy indices for lower p are *included* among the the set of high-entropy indices of higher p . We call this the “*nested*” property. The nested property can be useful in applications because it allows to take measurements adaptively if the distribution of the binary source is a priori unknown. In other words, one takes some measurements corresponding to high-entropy indices and if the recovery is not successful, refines them by adding extra measurements corresponding to low-entropy ones to improve the quality of the recovery.

2.7.3 Robustness to Measurement Noise

Having the measurement taken from the source, one can use different algorithms to reconstruct it. We use the optimal Maximum Likelihood (ML) algorithm for recovery which uses the recursive structure of the Hadamard matrices to speed up the reconstruction. For simulation, we used $N = 512$, $p = 0.05$ and took all of the indices with entropy greater than 0.01. Let us denote the input random variables by X_1^N and assume that we keep all of the rows of the matrix H_n with indices in the set S . Suppose that Y_S is the set of all measurements. We denote by Z_S the resulting measurements after adding noise, where for $i \in S$, $Z_i = Y_i + W_i$ and W_i are i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables. We define the *signal to noise ratio* (SNR) at the input of the decoder by $\text{SNR}_{\text{in}} = \frac{\sum_{i \in S} \mathbb{E}(Y_i^2)}{|S| \sigma^2}$, where σ^2 is the noise variance. The SNR at the output of the decoder is defined by $\text{SNR}_{\text{out}} = \frac{N \mathbb{E}(X^2)}{\sum_{i=1}^N \mathbb{E}(|X_i - \hat{X}_i|^2)}$, where \hat{X}_i is the the estimate of X_i from the measurements at the output of the ML decoder. Figure 2.4 shows the stability analysis of the ML algorithm to the Gaussian measurement noise. The result shows approximately 4 dB loss in SNR for high SNR regime. Notice that

part of this loss is because of the distortion created by removing the measurements corresponding to low-entropy indices.

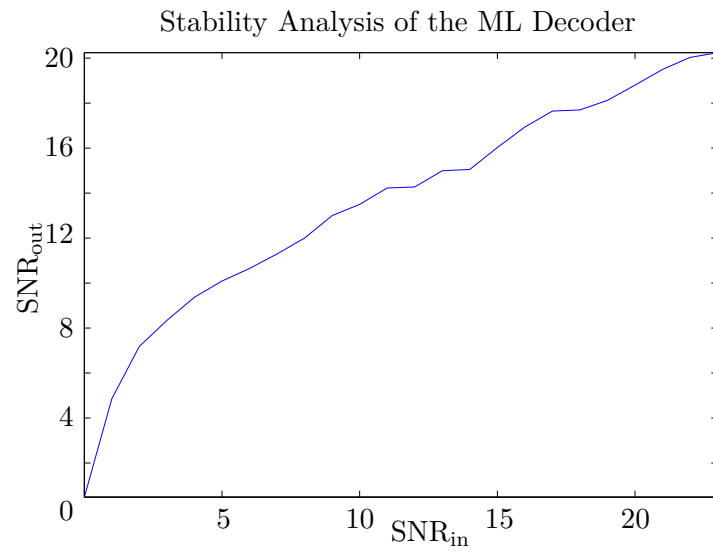


Figure 2.4 – Stability of the Maximum Likelihood (ML) Decoder to additive Gaussian measurement noise.

Fast Hadamard Transform: Construction for Signals with Sub-linear Sparsity

3

In Chapter 2, we explained the Hadamard construction to capture the information of an integer-valued memoryless source¹. In this chapter, we give a new Hadamard construction for compressed sensing of signals with sub-linear sparsity. Specifically, we assume that for an N -dimensional signal, the number of nonzero components K scales like $K = O(N^\alpha)$, sub-linearly in N for some $0 < \alpha < 1$. Depending on the value of α , we select a specific collection of rows of a Hadamard matrix to build the measurement matrix. Although in the thesis we mainly focus on matrix construction for the compressed sensing, the proposed construction in this chapter can be considered from two completely different points of view:

1. From a compressed sensing point of view, the proposed construction allows to build ensemble of measurement matrices to capture the information of sparse signals. By assuming a sub-linear sparsity model for the source and by using the underlying properties of the Hadamard matrices, we design a fast iterative and low-complexity algorithm to recover the nonzero components with a very high probability. The only requirement for the algorithm to work successfully is that the position of the nonzero components (support) of the signal be uniformly random. This is not a major problem because one can shuffle the position of the nonzero components by simply permuting the columns of the measurement matrix uniformly randomly before taking the measurements. Moreover, as we will explain, it is possible to know whether the recovery algorithm succeeds to recover all the nonzero components or not. In particular, if the recovery is not successful and if it is possible to repeat the measurement process, one can take a new set of measurements by using a new random permutation of the columns of the matrix. We show that the failure probability of the algorithm decays exponentially fast in the number of repetitions.
2. From another point of view, the proposed matrix construction and low-complexity recovery algorithm can be exploited to speed up the traditional Walsh-Hadamard Transform (WHT). To explain more, it is known that the

¹This chapter is the result of collaboration with Robin Scheibler and Martin Vetterli.

recursive structure of the Hadamard matrices allows to compute the WHT of an N -dimensional signal, which we call the *time domain* signal, in $O(N \log_2(N))$ operations assuming that one has access to all the samples of the signal. We call the output of the WHT the *transform domain* signal. Exploiting the construction proposed in this chapter, we improve this algorithm in two directions. First, the time complexity of the algorithm is reduced significantly. Second, the sample complexity is much less, i.e., it is not necessary to have all the N samples of the signal since a few number of time domain samples is sufficient to compute the transform domain signal. The requirement is that the support of the signal be random. Specially, in contrast to the compressed sensing setup already mentioned, where the random support can be created by permuting the columns of the measurement matrix, in this case the randomness of the support is crucial for the success of the recovery because in the first place, one has access only to the time domain signal and one can not shuffle the components of the transform domain signal before taking the measurements. This has been pictorially shown in Figure 3.1.

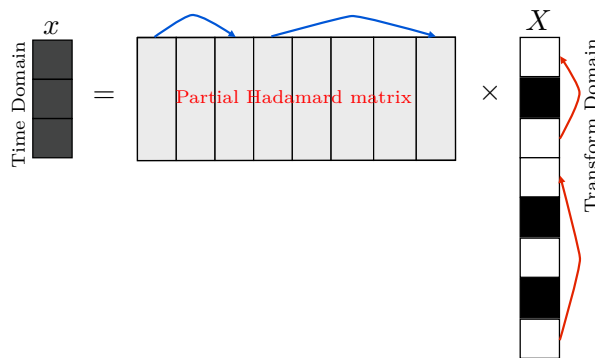
In this chapter, we explain the construction with an inclination towards WHT. However, the results can be immediately applied to the compressed sensing setup.

The structure of this chapter is as follows. In Section 3.1, we give an overview of the history and recent results concerning the WHT and its twin *Discrete Fourier Transform* (DFT) and their application in signal processing. Section 3.2 gives the main results. In Section 3.3, we prove some of the properties of the WHT that are crucial for our construction. Using these properties, we develop a *Hashing Algorithm* that is explained in Section 3.4. In Section 3.5, we introduce the Sparse Fast Hadamard Transform (SFHT) whose performance is analyzed in Section 3.6 and 3.7 for two different sparsity regimes: very sparse and less sparse. Simulation results are given in Section 3.8. Finally, Section 3.9 concludes the chapter and provides some suggestions for further extension.


3.1 Walsh-Hadamard Transform: Overview and Related Work

The Walsh-Hadamard transform (WHT) is a well-known signal processing tool with application in areas as varied as image compression and coding [33], spreading sequence for multi-user transmission in cellular networks (CDMA) [34], spectroscopy [35] as well as compressed sensing [36]. It has interesting properties studied in different areas of mathematics [37]. It also shares many underlying properties with the Discrete Fourier Transform (DFT). For example, both have a nice recursive structure which allows a fast computation with a time complexity $O(N \log_2(N))$ in the dimension of the signal N [38, 39].

A number of recent publications have addressed the particular problem of computing the DFT of an N -dimensional signal under the assumption of K -sparsity of the signal in the frequency domain [40–44]. In particular, it has been shown that the well-known computational complexity $O(N \log_2(N))$ of the FFT algorithm can be strictly improved. Such algorithms are generally known as *sparse FFT* (sFFT) algorithms. The authors in [45] by extending the results of [44], gave a very low-



(a) Hadamard Construction for CS: One starts from the signal X and takes linear measurements to recover X . The random support for X can be created by permuting the columns of the matrix before taking the measurements.



(b) Hadamard Construction for WHT: One starts from the signal x (rather than X) and needs to reconstruct X . Therefore, the measurement matrix is fixed and it is not possible to shuffle its columns in order to build a random support for the signal X .

Figure 3.1 – Comparison of the Hadamard Construction for CS with WHT

complexity algorithm for computing the 2D-DFT of a $\sqrt{N} \times \sqrt{N}$ signal. In a similar line of work, based on the subsampling property of the DFT in the time domain resulting in aliasing in the frequency domain, the authors in [46, 47] developed a novel low-complexity iterative algorithm to recover the non-zero frequency elements using ideas from sparse-graph codes [48].

Since the Hadamard matrices share the same recursive structure as the DFT matrices, one might naturally ask if the results developed for the sparse FFT can be exploited to develop a fast Walsh-Hadamard Transform. In this chapter, we develop such an algorithm to speed up the traditional Walsh-Hadamard Transform. Although the construction seems similar to the sparse FFT (for example [46, 47]), there are some unique features of the Hadamard matrices that make the construction distinct. For example, in a DFT matrix of order N , different rows of the matrix correspond to different harmonics but a Hadamard matrix consists of only ± 1 -valued components

which, roughly speaking, correspond to two different harmonics.

To give a brief overview of the construction of fast Walsh-Hadamard Transform, we first develop some useful properties of the WHT, specially the subsampling and the modulation property, which play a key role in the development the algorithm. In particular, we show that by suitable subsampling a signal in the time domain, one can induce a well-designed aliasing pattern over the transform domain components. In other words, it is possible to obtain a linear combination of a controlled collection of transform domain components (aliasing). This creates interference between the non-zero components if more than one of them are involved in the induced linear combination. Similar to [47] and borrowing ideas from sparse-graph codes, we construct a bipartite graph by treating the non-zero values in the transform domain as variable nodes and interpreting any induced aliasing pattern as a parity check constraint over those variable nodes. We analyze the structure of the resulting graph assuming a random support model for the non-zero coefficients in the transform domain. Moreover, we give an iterative peeling decoder to recover those non-zero components. Very briefly, our proposed sparse fast Hadamard transform (SFHT) consists of a set of deterministic linear hash functions (explicitly constructed) and an iterative peeling decoder that uses the hash outputs to recover the non-zero transform domain variables. It recovers the K -sparse WHT of the signal in sample complexity (number of time domain samples used) $O(K \log_2(\frac{N}{K}))$, total computational complexity $O(K \log_2(K) \log_2(\frac{N}{K}))$ and with a high probability approaching 1 as N tends to infinity. Table 3.1 summarizes the notations used in this chapter.

Table 3.1 – Summary of Notations

\mathbb{Z}	the set of integers	$[m]$	$\{1, 2, \dots, m\}$
\mathbb{Z}_+	the set of positive integers	X_i^j	$\{X_i, X_{i+1}, \dots, X_j\}$
\mathbb{N}	the set of strictly positive integers	v_i	the i -th components of v
\mathbb{R}	the set of reals	H_N	Hadamard matrix of order N
$a \wedge b$	minimum of a and b	$a \vee b$	maximum of a and b
\mathbb{F}_2	Binary field $\{0, 1\}$	\mathbb{F}_2^n	n -dimensional binary vectors
\mathcal{N}	null space of a matrix	\mathcal{N}	neighborhood a node in a graph
$\text{supp}(v)$	support of the vector v , $\{i : v_i \neq 0\}$		
$v_{i_0, i_1, \dots, i_{n-1}}$	the i -th component of v with binary expansion $\langle i_0, i_1, \dots, i_{n-1} \rangle$		
x and X	time domain and frequency domain signals		
k -sparse signal	signal having only k nonzero components		
WHT	Walsh-Hadamard Transform		
SFHT	Sparse Fast Hadamard Transform		

3.2 Main Results

We summarize the main result of this chapter in the following theorem.

Theorem 3.1. *Let $0 < \alpha < 1$, $N = 2^n$ a power of two and $K = N^\alpha$. Let $x \in \mathbb{R}^N$ be a time domain signal with a WHT signal $X \in \mathbb{R}^N$. Assume that X is a K -sparse signal whose support is selected uniformly at random among all possible $\binom{N}{K}$ subsets*

of $[N]$ of size K . There is an algorithm that can compute X and has the following properties:

1. **Sample complexity:** The algorithm uses $CK \log_2(\frac{N}{K})$ time domain samples of the signal x . C is a function of α and $C \leq (\frac{1}{\alpha} \vee \frac{1}{1-\alpha}) + 1$, where for $a, b \in \mathbb{R}_+$, $a \vee b$ denotes the maximum of a and b .
2. **Computational complexity:** The total number of operations in order to either successfully recover the position and the value of all the non-zero spectral components or to announce a decoding failure is $O(CK \log_2(K) \log_2(\frac{N}{K}))$.
3. **Success probability:** The algorithm correctly computes the K -sparse WHT of the signal X with very high probability asymptotically approaching 1 as N tends to infinity, where the probability is taken over all random selections of the support of X .

To prove Theorem 3.1, we distinguish between the very sparse case ($0 < \alpha \leq \frac{1}{3}$) and the less sparse one ($\frac{1}{3} < \alpha < 1$). Also, we implicitly assume that the algorithm knows the value of α , which might not be possible in reality. As we will see later, if we know some range to which α belongs, it will be possible to design an algorithm that works for all those values of α . However, the sample and computational complexity of the algorithm might increase compared with the optimal one that knows the value of α . For example, if we know that the signal is very sparse, $\alpha \in (0, \alpha^*]$ for some $\alpha^* \leq \frac{1}{3}$, it is sufficient to design the algorithm for α^* and it will work for all signals with sparsity index less than α^* . Similarly, if the signal is less sparse with a sparsity index $\alpha \in (\frac{1}{3}, \alpha^*)$ for some $\frac{1}{3} < \alpha^* < 1$, then again it is sufficient to design the algorithm for α^* and it will automatically work for all $\alpha \in (\frac{1}{3}, \alpha^*)$.

Remark 3.1. In the very sparse regime ($0 < \alpha \leq \frac{1}{3}$), we prove that for any value of α , the success probability of the optimally designed algorithm is at least $1 - O(1/K^{3(C/2-1)})$, with $C = \lceil \frac{1}{\alpha} \rceil$ where for $u \in \mathbb{R}_+$, $\lceil u \rceil = \max\{n \in \mathbb{Z} : n \leq u\}$. It is easy to show that for every value of $\alpha \in (0, \frac{1}{3})$ the success probability can be lower bounded by $1 - O(N^{-\frac{3}{8}})$.

3.3 Walsh-Hadamard Transform and its Properties

Let x be an $N = 2^n$ dimensional signal indexed with elements $m \in \mathbb{F}_2^n$. The N -dimensional WHT of the signal x is defined by

$$X_k = \frac{1}{\sqrt{N}} \sum_{m \in \mathbb{F}_2^n} (-1)^{\langle k, m \rangle} x_m, \quad (3.1)$$

where $k \in \mathbb{F}_2^n$ denotes the corresponding binary index of the transform domain component. Moreover, $\langle k, m \rangle = \sum_{i=0}^{n-1} k_i m_i$, where the arithmetic is done over the binary field \mathbb{F}_2 . We also use the convention that $(-1)^0 = 1$ and $(-1)^1 = -1$. Borrowing some terminology from the DFT, we call the transform domain samples $X_k, k \in \mathbb{F}_2^n$, frequency or spectral domain components of the time domain signal x .

Notice that with our notation both the time-domain signal $x : \mathbb{F}_2^n \rightarrow \mathbb{R}$ and the frequency-domain signal $X : \mathbb{F}_2^n \rightarrow \mathbb{R}$ are functions from the index set \mathbb{F}_2^n to reals.

32 Fast Hadamard Transform: Construction for Signals with Sub-linear Sparsity

Therefore, the WHT given by the Equation 3.1 is a mapping between the function x (time domain signal) and the function X (transform domain signal). For simplicity of the notation, we will use x_m and X_k for the time and frequency domain function with the convention that both m and k belong to the index set \mathbb{F}_2^n .

3.3.1 Basic Properties

This subsection is devoted to reviewing some of the basic properties of the WHT. Some of the properties are not directly used and we have included them for the sake of completeness. They can be of independent interest. The proofs of all the properties are provided in Section 3.10.1.

Property 1 (Shift/Modulation). *Let X_k be the WHT of the signal x_m and let $p \in \mathbb{F}_2^n$. Assume that $y_m = x_{m+p}$ denotes the time domain signal $y : \mathbb{F}_2^n \rightarrow \mathbb{R}$, where $y = x \circ i_p$ is the composition of the function x with the index transformation function $i_p : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ given by $i_p(m) = m + p$, where $m + p$ denotes the component-wise module two addition of binary vectors m and p . Then*

$$y_m = x_{m+p} \xleftrightarrow{\text{WHT}} X_k(-1)^{\langle p, k \rangle}.$$

The next property explains one of the interesting properties of the WHT. Suppose that one has access only to the time domain signal and for some reason one is interested to shuffle the transform domain values by some simple manipulation of the time domain samples. This property allows to partially permute the Hadamard spectrum in a specific way by applying a corresponding permutation in the time domain. However, the collection of all such possible permutations is limited. Technically, this property is equivalent to finding permutations $\pi_1, \pi_2 : [N] \rightarrow [N]$ with corresponding permutation matrices Π_1, Π_2 such that

$$\Pi_2 H_N = H_N \Pi_1, \tag{3.2}$$

where H_N is the Hadamard matrix of order N and where the permutation matrix corresponding to a permutation π is defined by $(\Pi)_{i,j} = 1$ if and only if $\pi(i) = j$, and zero otherwise. The identity (3.2) is equivalent to finding a row permutation of H_N that can be equivalently obtained by a column permutation of H_N .

Property 2. *all the permutations satisfying (3.2) are described by the elements of*

$$\text{GL}(n, \mathbb{F}_2) = \{A \in \mathbb{F}_2^{n \times n} \mid A^{-1} \text{ exists}\},$$

where $\text{GL}(n, \mathbb{F}_2)$ denotes the set of $n \times n$ non-singular matrices with entries in \mathbb{F}_2 .

Remark 3.2. *The total number of possible permutations in Property 2 is given by $\prod_{i=0}^{n-1} (N - 2^i)$, which is a negligible fraction of all $N!$ permutations over $[N]$.*

Property 3 (Permutation). *Let $\Sigma \in \text{GL}(n, \mathbb{F}_2)$. Assume that X_k is the WHT of the time domain signal x_m . Then*

$$x_{\Sigma m} \xleftrightarrow{\text{WHT}} X_{\Sigma^{-T} k}.$$

Notice that $y_m = x_{\Sigma m}$ is the function given by the composition of the function x and the index transformation i_Σ , $y = x \circ i_\Sigma$, where for $m \in \mathbb{F}_2^n$, $i_\Sigma(m) = \Sigma m$ is the multiplication of the matrix Σ with the index vector m . Moreover, any $\Sigma \in \text{GL}(n, \mathbb{F}_2)$ is a bijection from \mathbb{F}_2^n to \mathbb{F}_2^n , thus $x_{\Sigma m}$ is simply a signal obtained by permuting the components of the signal x_m .

The last property is that of downsampling/aliasing. Notice that for a signal x consisting of $N = 2^n$ components, we index every component by a binary vector of length n , namely, $x_{m_0, m_1, \dots, m_{n-1}}$. To subsample this signal “along dimension i ”, we freeze the i -th component of the index to either 0 or 1. For example, $x_{0, m_1, \dots, m_{n-1}}$ is a 2^{n-1} dimensional signal obtained by subsampling the vector x_m along the first index. More precisely, the subsampled signal is simply the restriction of the function $x : \mathbb{F}_2^n \rightarrow \mathbb{R}$ to the subset $\{0\} \times \mathbb{F}_2^{n-1}$, i.e., $x_{0, m_1, \dots, m_{n-1}} = x|_{\{0\} \times \mathbb{F}_2^{n-1}}$.

Property 4 (Downsampling/Aliasing). *Suppose that x is a vector of dimension $N = 2^n$ indexed by the elements of \mathbb{F}_2^n and assume that $B = 2^b$, where $b \in \mathbb{N}$ and $b < n$. Let*

$$\Psi_b = \begin{bmatrix} \mathbf{0}_{b \times (n-b)} & I_b \end{bmatrix}^T, \quad (3.3)$$

be the subsampling matrix freezing the first $n - b$ components in the index to 0. If X_k is the WHT of x , then

$$x_{\Psi_b m} \xleftrightarrow{\text{WHT}} \sqrt{\frac{B}{N}} \sum_{j \in \mathcal{N}(\Psi_b^T)} X_{\Psi_b k + j},$$

where $m, k \in \mathbb{F}_2^b$ denote the corresponding binary indices of the time and frequency components and $x_{\Psi_b m}$ is a $B = 2^b$ dimensional signal labelled with $m \in \mathbb{F}_2^b$.

Recall that by our notation, $y_m = x_{\Psi_b m}$ is a function $y : \mathbb{F}_2^b \rightarrow \mathbb{R}$ given by $y = x \circ i_{\Psi_b}$, where $i_{\Psi_b} : \mathbb{F}_2^b \rightarrow \mathbb{F}_2^n$ is the index transformation given by $i_{\Psi_b}(m) = \Psi_b m$. One can simply check that $\Psi_b m$, which is the multiplication of $m \in \mathbb{F}_2^b$ with the matrix Ψ_b of dimension $n \times b$, gives an index in \mathbb{F}_2^n which is the right argument for the function x . One can also check that the index $\Psi_b k + j$ with $j \in \mathcal{N}(\Psi_b^T)$ gives the suitable index for the function X . Notice that Property 4 can be simply applied for any matrix Ψ_b that subsamples any set of indices of length b not necessarily the b last ones.

To give further intuition about the downsampling property, notice that the elements of \mathbb{F}_2^n can be visualized as the vertices of an n -dimensional hypercube. The downsampling property just explained implies that downsampling along some of the dimensions in the time domain is equivalent to summing up all the spectral components along *the same dimensions* in the spectral domain. This is illustrated visually in Figure 3.2 for dimension $n = 3$.

In a general downsampling procedure, one can replace the frozen indices by an arbitrary but fixed binary pattern. The only difference is that instead of summing the aliased spectral components, one should also take into account the suitable $\{+, -\}$ sign patterns, i.e., one has

$$x_{\Psi_b m + p} \xleftrightarrow{\text{WHT}} \sqrt{\frac{B}{N}} \sum_{j \in \mathcal{N}(\Psi_b^T)} (-1)^{\langle p, j \rangle} X_{\Psi_b k + j}, \quad (3.4)$$

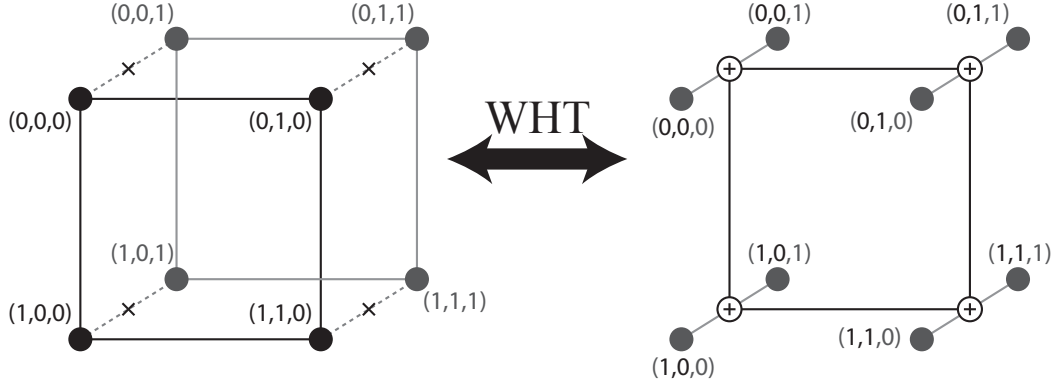


Figure 3.2 – Illustration of the downsampling property on a hypercube for $N = 2^3$. The two cubes represent the time-domain (left) and the Hadamard-domain (right) signal. We decide to drop all the nodes whose third coordinate is ‘1’. We illustrate this by adding a ‘x’ on the edges connecting those vertices through the third coordinate. This is equivalent to summing up vertices along the corresponding edges in the Hadamard domain.

where p is a binary vector of length n with b zeros at the end. To visualize this property, consider Figure 3.2, where we have a signal over a 3-dimensional cube and we subsample it along the third dimension, i.e., we keep only 4 variables with the third index equal to 0. Notice that these variables lie on a 2-dimensional (square) face of the cube that corresponds to a subsampling with $p = 000$. Instead we can use $p = 001$ for subsampling and this value of p will select all 4 variables on the other face of the cube corresponding to those variables with the third index equal to 1. This face of the cube is a square parallel to the square corresponding to $p = 000$.

3.4 Hadamard Hashing Algorithm

By applying the basic properties of the WHT, one can design suitable hash functions in the spectral domain. The main idea is that one does not need the spectral values in order to compute the hash outputs because they can be simply computed by low-complexity operations on the time domain samples of the signal.

Proposition 3.1 (Hashing). *Assume that $\Sigma \in \text{GL}(n, \mathbb{F}_2)$ and $p \in \mathbb{F}_2^n$. Let $N = 2^n$, $b \in \mathbb{N}$, $B = 2^b$ with $b < n$, and let $m, k \in \mathbb{F}_2^b$ denote the time and frequency indices of a B -dimensional signal and its WHT respectively, where the signal is defined by $u_{\Sigma,p}(m) = \sqrt{\frac{N}{B}} x_{\Sigma\Psi_b m+p}$. The length B Hadamard transform of $u_{\Sigma,p}$ is given by*

$$U_{\Sigma,p}(k) = \sum_{j \in \mathbb{F}_2^n \mid \mathcal{H}j=k} X_j (-1)^{\langle p, j \rangle}, \quad (3.5)$$

where \mathcal{H} is the index hashing operator defined by

$$\mathcal{H} = \Psi_b^T \Sigma^T, \quad (3.6)$$

where Ψ_b is as in (3.3).

Remark 3.3. Note that the index of components in the sum (3.5) can be explicitly written as function of the bin index k , i.e., $j = \Sigma^{-T} \Psi_b k + q$, $q \in \mathcal{N}(\mathcal{H})$.

The proof simply follows from the properties 1, 3, and 4. Based on Proposition 3.1, we give Algorithm 1 which computes the hashed Hadamard spectrum. Given an FFT-like fast Hadamard transform (FHT) algorithm, and picking B bins for hashing the spectrum, Algorithm 1 requires $O(B \log B)$ operations. Figure 3.3 shows an illustration of the resulting hashing.

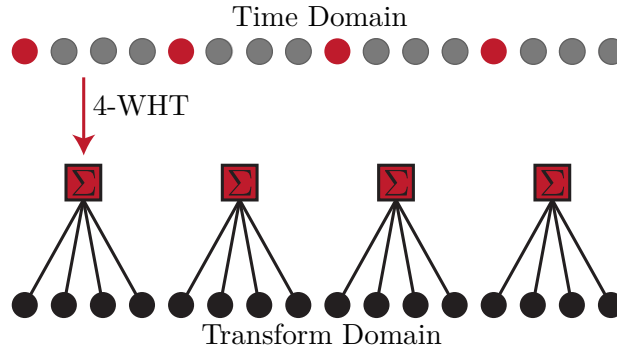


Figure 3.3 – Illustration of the Hadamard Hashing. Time domain signal x has dimension 16 and is subsampled by 4. The red circles show the resulting subsampled signal. Computing the WHT of the 4-sample signal is equivalent to hashing the transform domain signal in 4 bins and taking the summation.

Algorithm 1 FastHadamardHashing(x, N, Σ, p, B)

Require: Signal x of dimension $N = 2^n$, subsampling parameters Σ and p and the number of output bins $B = 2^b$ in a hash.

Ensure: U_k is the hashed Hadamard spectrum of x .

$$u_m = x_{\Sigma \Psi_b m + p}, \text{ for } m \in \mathbb{F}_2^b.$$

$$U_k = \sqrt{\frac{N}{B}} \text{WHT}(u_m) \text{ is the resulting } B\text{-dimensional WHT.}$$

3.4.1 Properties of the Hadamard Hashing

In this part, we review some of the nice properties of the hashing algorithm that are crucial for developing an iterative peeling decoding algorithm to recover the non-zero spectral values. We explain how it is possible to identify collisions between the non-zero spectral coefficients that are hashed to the same bin and also to estimate the support of non-colliding components.

Let us consider $U_{\Sigma, p}(k)$ for two cases: $p = 0$ and some $p \neq 0$. It is easy to see that in the former $U_{\Sigma, p}(k)$ is obtained by summing all the spectral variables hashed to bin k , i.e., those whose index j satisfies $\mathcal{H}j = k$, whereas in the latter the same variables are added together after being weighted by $(-1)^{\langle p, j \rangle}$. Let us define the following ratio test

$$r_{\Sigma, p}(k) = \frac{U_{\Sigma, p}(k)}{U_{\Sigma, 0}(k)}.$$

When the sum in $U_{\Sigma,p}(k)$ contains only one non-zero component, it is easy to see that $|r_{\Sigma,p}(k)| = 1$ for ‘any value’ of p . However, if there is more than one component in the sum, under a very mild assumption on the non-zero coefficients of the spectrum (e.g., if they are jointly sampled from a continuous distribution), one can show that $|r_{\Sigma,p}(k)| \neq 1$ for at least some values of p . In fact, $n - b$ well-chosen values of p allow to detect whether there is only one, or more than one non-zero components in the sum.

When there is only one $X_{j'} \neq 0$ hashed to the bin k ($h_{\Sigma}(j') = k$), the result of the ratio test is precisely 1 or -1 , depending on the value of the inner product between j' and p . In particular, we have

$$\langle p, j' \rangle = \mathbb{1}_{\{r_{\Sigma,p}(k) < 0\}}, \quad (3.7)$$

where $\mathbb{1}_{\{t < 0\}} = 1$ if $t < 0$, and zero otherwise. Hence, if for $n - b$ well-chosen values of p , the ratio test results in 1 or -1 , implying that there is only one non-zero spectral coefficient in the corresponding hash bin. By some extra effort, it is even possible to identify the position of this non-zero component. We formalize this result in the following proposition proved in Section 3.10.2.

Proposition 3.2 (Collision detection/Support estimation). *Let $\Sigma \in \text{GL}(n, \mathbb{F}_2)$ and let $\sigma_i, i \in [n]$ denote the columns of Σ .*

1. *If for all $d \in [n - b]$, $|r_{\Sigma, \sigma_d}(k)| = 1$ then almost surely there is only one non-zero spectral value in the bin indexed by k . Moreover, if we define*

$$\hat{v}_d = \begin{cases} \mathbb{1}_{\{r_{\Sigma, \sigma_d}(k) < 0\}} & d \in [n - b], \\ 0 & \text{otherwise,} \end{cases}$$

the index of the unique non-zero coefficient is given by

$$j = \Sigma^{-T}(\Psi_b k + \hat{v}). \quad (3.8)$$

2. *If there exists a value $d \in [n - b]$ such that $|r_{\Sigma, \sigma_d}(k)| \neq 1$ then the bin k contains more than one non-zero coefficient.*

3.5 Sparse Fast Hadamard Transform

In this section, we give a brief overview of the main idea of Sparse Fast Hadamard Transform (SFHT). In particular, we explain the peeling decoder, which recovers the non-zero spectral components and analyze its computational complexity.

3.5.1 Explanation of the Algorithm

Assume that x is an $N = 2^n$ dimensional signal with a K -sparse WHT X (having only K non-zero values). Suppose that $K = O(N^\alpha)$ scales sub-linearly with N with index $\alpha \in (0, 1)$. As $H_N^{-1} = H_N$, taking the inner product of the vector X with the i -th row of the Hadamard matrix H_N gives the time domain sample x_i . Using the terminology of coding theory, it is possible to consider the spectral components X




Figure 3.4 – On the left, bipartite graph representation of the WHT for $N = 8$ and $K = 3$. On the right, the underlying bipartite graph after applying $C = 2$ different hashing produced by plugging Σ_1, Σ_2 in (3.6) with $B = 4$. The variable nodes (\bullet) are the non-zero spectral values to be recovered. The white check nodes (\square) are the original time-domain samples. The colored squares are new check nodes after applying Algorithm 1.




Figure 3.5 – A block diagram of the SFHT algorithm in the time domain. The downsampling plus small-size low-complexity FHT blocks compute different hash outputs. Delay blocks denote an index shift by σ_i before hashing. The S/P and P/S are serial-parallel and parallel-serial blocks to emphasize that the FHT operates on the whole signal at once. The collision detection/support estimation block implements Proposition 3.2 to identify if there is a collision and if not to find the index of the only non-zero value. The recovered index i is not valid when there is a collision.

as variables nodes (information bits in coding theory), where the inner product of the i -th row of H_N is like a parity check constraint on X . For example, the first row of the Hadamard matrix is the all-one vector which implies that the sum of all the components of X must be equal to the first time domain sample. A similar interpretation holds for the other rows. Thus, the WHT can be imagined as a code defined by a bipartite graph. With this picture in mind, one can consider the recovery of the non-zero spectral values as a decoding problem over this bipartite graph. If we consider the WHT, it is easy to see that the induced bipartite graph on the non-zero spectral values is a complete (dense) bipartite graph because any variable node is connected to all the check nodes as has been depicted in the left part of Figure 3.4, where $\{X_1, X_8, X_{11}\}$ are the only non-zero variables in the spectral domain and each check constraint corresponds to the value of a specific time domain sample. It is also implicitly assumed that one knows the support of X , e.g., $\{1, 8, 11\}$ in our case. At

the moment, it is not clear how one can obtain the position of the non-zero variables. As we will explain later, in the final version of the algorithm this can be done by applying Proposition 3.2.

For codes on bipartite graphs, there is a collection of low-complexity belief propagation algorithms to recover the variable nodes given the value of check nodes. Most of these algorithms perform well if the underlying graph is sparse. Unfortunately, the graph corresponding to WHT is dense, and probably not suitable for any of these belief propagation algorithms.

As explained in Section 3.4, by subsampling the time domain components of the signal, it is possible to hash the spectral components in different bins as depicted for the same signal X in the right part of Figure 3.4. The advantage of the hashing operation must be clear from this picture. The idea is that hashing ‘*sparsifies*’ the underlying bipartite graph. It is also important to notice that in the bipartite graph induced by hashing, one can obtain all the values of parity checks (hash outputs) by using low-complexity operations on a small set of time domain samples as explained in Proposition 3.1.

We propose the following iterative algorithm to recover the non-zero spectral variables over the bipartite graph induced by hashing. The algorithm first tries to find a degree one check node. Using the terminology of [47], we call such a check node a *singleton*. Using Proposition 3.2, the algorithm is able to find the position and the value of the corresponding non-zero component, thus, the algorithm can subtract (peel off) this variable from all the other check nodes that are connected to it. In particular, after this operation the corresponding singleton check node gets value zero, i.e., it is satisfied. Equivalently, we can update the underlying graph by removing the mentioned variable node from the graph along with all the edges connected to it. This creates an isolated (degree zero) check node which we call a *zeroton*. Also notice that by removing some of the edges from the graph, the degree of the associated check nodes decreases by one, thus there is a chance that another singleton be found.

The algorithm proceeds to peel off one singleton at a time until all the check nodes are *zeroton* (decoding succeeds) or all the remaining check nodes have degree greater than one (we call them *multiton*) and the algorithm fails to identify all the non-zero spectral values.

A more detailed pseudo-code of the proposed iterative algorithm is given in Algorithm 2.

3.5.2 Complexity Analysis

Figure 3.5 shows a full block diagram of the SFHT algorithm. Using this block diagram, in this section, we prove part 1 and 2 of Theorem 3.1 regarding the sample and the computational complexity of the SFHT algorithm. The last part of the theorem concerns the success probability of the algorithm which will be separately analyzed in Sections 3.6 and 3.7 for the very and less sparse regimes, respectively.

Computational Complexity: As it will be further clarified in Sections 3.6 and 3.7, depending on the sparsity index of the signal α , we will use C different hash functions, where $C \leq (\frac{1}{\alpha} \vee \frac{1}{1-\alpha}) + 1$ each with $B = 2^b$ different output bins. We always select $B = K$ to keep the average number of non-zero components per bin $\beta = \frac{K}{B}$ equal to 1. This implies that computing the hash outputs via an

Algorithm 2 SFHT(x, N, K, C, L, Σ)

Require: Input signal x of length $N = 2^n$. Sparsity K . Hash count C . Number of iterations of decoder L . Array Σ of C matrices in $\text{GL}(n, \mathbb{F}_2)$, $\Sigma_c = [\sigma_{c,1} \mid \cdots \mid \sigma_{c,n}]$, $\sigma_{c,i} \in \mathbb{F}_2^n$.

Ensure: X contains the sparse Hadamard spectrum of x .

$$B = O(K)$$

$$D = n - b + 1$$

for $c = 1, \dots, C$ **do**

$$U_{c,0} = \text{FastHadamardHashing}(x, N, \Sigma_c, 0, B)$$

for $d = 1, \dots, D$ **do**

$$U_{c,d} = \text{FastHadamardHashing}(x, N, \Sigma_c, \sigma_{c,d}, B)$$

end for

end for

for $l = 1, \dots, L$ **do**

for $c = 1, \dots, C$ **do**

for $k = 0, \dots, B - 1$ **do**

if $U_{c,0,k} = 0$ **then**

continue to next k

end if

$$\hat{v} \leftarrow 0$$

for $d = 1, \dots, D$ **do**

if $U_{c,d,k}/U_{c,0,k} = -1$ **then**

$$\hat{v}_{d-1} \leftarrow 1$$

else if $U_{c,d,k}/U_{c,0,k} \neq 1$ **then**

continue to next k

end if

end for

$$i \leftarrow \Sigma_c^{-T}(\Psi_b k + \hat{v})$$

$$X_i \leftarrow U_{c,0,k}$$

for $c' = 1, \dots, C$ **do**

$$j \leftarrow \Psi_b^T \Sigma_{c'}^T i$$

$$U_{c',0,j} \leftarrow U_{c',0,j} - X_i$$

for $d' = 1, \dots, D$ **do**

$$U_{c',d',j} \leftarrow U_{c',d',j} - X_i (-1)^{\langle \sigma_{c',d'}, i \rangle}$$

end for

end for

end for

end for

end for

FHT block of size B needs $B \log_2(B)$ operations which assuming $K = B$, gives a computational complexity $K \log_2(K)$. Moreover, we need to compute any hash output with $n - b = \log_2(\frac{N}{B})$ different shifts in order to do collision detection/support estimation, thus, the computational cost for each hash is $K \log_2(K) \log_2(\frac{N}{K})$. Since we need to compute C different hash blocks, the total computational complexity of each iteration will be $CK \log_2(K) \log_2(\frac{N}{K})$. We will explain later that the algorithm terminates in a fixed number of iterations independent of the value of α and the dimension of the signal N . Therefore, the total computational complexity of the algorithm will be $O(CK \log_2(K) \log_2(\frac{N}{K}))$.

Sample Complexity: Assuming $K = B$, computing each hash with $n - b$ different shifts needs $K \log_2(\frac{N}{K})$ time domain samples. Hence, the total sample complexity, i.e., the required number of time domain samples, will be $CK \log_2(\frac{N}{K})$.

3.6 Performance Analysis of the very Sparse Regime

In this section, we consider the very sparse regime, where $0 < \alpha \leq \frac{1}{3}$. In this regime, we show that assuming a random support model for non-zero spectral components and a careful design of hash functions, it is possible to obtain a random bipartite graph with variable nodes corresponding to non-zero spectral components and with check nodes corresponding to outputs of hash functions. We explicitly prove that asymptotically this random graph behaves similar to the ensemble of LDPC bipartite graphs. Running the peeling decoder to recover the spectral components is also equivalent to the belief propagation (BP) decoding for a binary erasure channel (BEC). Fortunately, there is a rich literature in coding theory about asymptotic performance of the BP decoder. Specially, it is possible to show that the error (decoding failure) probability can be asymptotically characterized by a ‘*Density Evolution*’ (DE) equation which allows a perfect analysis of the peeling decoder.

We use the following steps to rigorously analyze the performance of the decoder in the very sparse regime:

1. We explain how to construct suitable hash functions depending on the value of $\alpha \in (0, \frac{1}{3}]$. Although, in this chapter we only deal with WHT but as the hash construction is deterministic, it automatically provides a deterministic partial Hadamard matrix that can be used for compressed sensing.
2. We rigorously analyze the structure of the induced bipartite graph obtained by treating the non-zero spectral components as variable nodes and the output of hash functions as check nodes. In particular, we prove that the resulting graph is a fully random left regular bipartite graph similar to the regular LDPC ensemble. We also obtain variable and check degree distribution polynomials for this graph.
3. At every stage, the peeling decoder recovers some of the variable nodes, removing all the edges incident to those variable nodes. We use Wormald’s method given in [49] to prove the concentration of the number of unpeeled edges around its expected value which we characterize as well. Wormald’s method as exploited in [50], uses the differential equation approach to track the evolution of the number of edges in the underlying bipartite graph. Specifically,

it shows that the number of edges at every step of the algorithm is very well concentrated around the solution of the associated differential equations.

4. Wormald's method gives a concentration bound to the number of remaining edges as far as their count is a fixed ratio $\gamma \in (0, 1)$ of the initial edges in the graph. Another expander argument as in [50] is necessary to show that if the decoder peels a $1 - \gamma$ fraction of the edges successfully for a small enough value of γ , it can continue to peel off all the remaining edges with very high probability.

3.6.1 Hash Construction

For the very sparse regime, $0 < \alpha \leq \frac{1}{3}$, consider those values of α equal to $\frac{1}{C}$ for some positive integer $C \geq 3$. We will explain later how to cover the intermediate values. For $\alpha = \frac{1}{C}$, we will consider C different hash functions as follows. Let x be an N -dimensional time domain signal with a WHT X , where $N = 2^n$ and let $b = \frac{n}{C}$. As we did before, we label the components of the vector X by an n -dimensional binary vector from \mathbb{F}_2^n . We design C different subsampling operators, where the i -th operator takes a binary index-vector from \mathbb{F}_2^n , keeps indices ib up to $(i+1)b - 1$ of this vector intact and sets the other binary-indices to zero. Using the terminology of Proposition 3.1, we define Σ_i to be the identity matrix with columns circularly shifted by $(i+1)b$ to the left. Then, the hash operator given by (3.6) is

$$\mathcal{H}_i = \Psi_b^T \Sigma_i^T = [\mathbf{0}_{b \times ib} \ I_b \ \mathbf{0}_{b \times (n - (i+1)b)}],$$

where I_b is the identity matrix of order b and Ψ_b is as defined in (3.3).

To give further intuition about the hash construction, consider an $N = 2^n$ dimensional signal x . We can label the components of x by a length n binary vector. With some abuse of notation, suppose $x_0^{n-1} \in \mathbb{F}_2^n$ is the corresponding binary vector. Equivalent to the C different subsampling operators, we can consider functions $h_i, i \in [C]$, where $h_i : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^b$ is given by

$$h_i(x_0^{n-1}) = (x_{ib}, x_{ib+1}, \dots, x_{ib+b-1}). \quad (3.9)$$

Now, let us consider two components of the signal x with binary indices x_0^{n-1} and y_0^{n-1} . Notice that if $x_t = y_t$ for $t = ib, ib+1, \dots, ib+b-1$, then these two variables are mapped to the same point in \mathbb{F}_2^b , i.e., $h_i(x_0^{n-1}) = h_i(y_0^{n-1})$. With this notation, it is easy to reinterpret the subsampling operator \mathcal{H}_i . This operator constructs a $B = 2^b$ dimensional signal from the $N = 2^n$ dimensional signal x as follows: simply consider a length b binary vector r and among all the elements in the initial signal x whose index-vector is mapped to r , select the one with maximum number of 0 in its index-vector.

It is also interesting to see what happens in the transform domain. Assume that X is the WHT of the signal x . Again with abuse of notation, we assume that the components of X are labelled by binary indices $X_0^{n-1} \in \mathbb{F}_2^n$. We know that under the i -th hash function h_i , we obtain a $B = 2^b$ dimensional signal subsampled from x . The WHT of this subsampled signal is a B -dimensional signal whose components can be labelled with \mathbb{F}_2^b . Let $r \in \mathbb{F}_2^b$ be an arbitrary index-vector in \mathbb{F}_2^b . Ignoring the multiplicative constants, it is seen from Equation (3.5) that to obtain the WHT

of the subsampled signal, one should sum up all the spectral component of X with $h_i(X_0^{n-1}) = r$. Intuitively, it is as if we have a bucket labelled with r and all the spectral components with X_0^{n-1} , $h_i(X_0^{n-1}) = r$, are mapped to this bucket. In particular, in order to obtain the value of the WHT at a specific label r , one sums up all the spectral components mapped to the bucket r .

In terms of computational complexity, to obtain the output of each hash bin, we only need to compute the WHT of a smaller subsampled signal of dimension B . Note that by hash construction, $K = B$ which implies that all the hash functions can be computed in $CK \log_2(K)$ operations. As we will explain later, we need at least $C = 3$ hashes for the peeling algorithm to work successfully and that is why this construction works for $\alpha \leq \frac{1}{3}$. For intermediate values of α , those not equal to $\frac{1}{C}$ for some integer C , one can construct $\lceil \frac{1}{\alpha} \rceil$ hashes with $B = 2^{\lceil n\alpha \rceil}$ output bins and one hash with smaller number of output bins, thus obtaining a computational complexity of order $(1 + \lceil \frac{1}{\alpha} \rceil)K \log_2(K)$.

3.6.2 Random Bipartite Graph Construction

Random Support Model

An N -dimensional signal $x \in \mathbb{R}^N$ is called K -sparse if $|\text{supp}(x)| = K$, where for $A \subset [N]$, $|A|$ denotes the cardinality of A . For a given (K, N) , RS1(K, N) is the class of all stochastic signals whose support is selected uniformly at random from the set of all $\binom{N}{K}$ possible supports of size K . We do not put any constraint on the non-zero components; they can be deterministic or random. Model RS1 is equivalent to selecting K out of N objects at random without replacement. If we assume that the selection of the indices for the support is done independently but with replacement, we obtain another model that we call RS2(K, N). In particular, if $V_i, i \in [K]$ are i.i.d. random variables uniformly distributed over $[N]$, a random support in RS2(K, N) is given by the random set $\{V_i : i \in [K]\}$. Obviously, the size of a random support in RS2(K, N) is not necessarily fixed but it is at most K . The following proposition, proved in Section 3.10.3, shows that in the sub-linear sparsity regime, RS1 and RS2 are essentially equivalent.

Proposition 3.3. *Consider the random support model RS2(K, N), where $K = N^\alpha$ for some fixed $0 < \alpha < 1$ and let H be the random size of the support set. Asymptotically as N tends to infinity $\frac{H}{K}$ converges to 1 in probability.*

‘Balls and Bins’ Model $\mathcal{G}(K, B, C)$

Consider C disjoint sets of check nodes S_1, S_2, \dots, S_C of the same size $|S_i| = B$. A graph in the ensemble of random bipartite graphs \mathcal{G} with K variable nodes at the left and $C \times B$ check nodes $\cup_{i=1}^C S_i$ at the right is generated as follows. Each variable node v in \mathcal{G} , independently from other variable nodes, is connected to check nodes $\{s_1, s_2, \dots, s_C\}$ where $s_i \in S_i$ is selected uniformly at random from S_i and the selection of s_i 's are independent of one another. Every edge e in \mathcal{G} can be labelled as (v, c) , where $v \in [K]$ is a variable node and c is a check node in one of S_1, S_2, \dots, S_C . For a variable node v , the neighbors of v , denoted by $\mathcal{N}(v)$, consist of C different check nodes connected to v , each of them from a different S_i . Similarly, for a check node $c \in \cup_{i=1}^C S_i$, $\mathcal{N}(c)$ is the set of all variable nodes connected to c .

By construction, all the resulting bipartite graphs in the ensemble are left regular with variable degree C but the check degree is not fixed. During the construction, it might happen that two variable nodes have exactly the same neighborhood. In that case, we consider them as equivalent variables and keep only one of them and remove the other, thus the number of variable nodes in a graph from the ensemble $\mathcal{G}(K, B, C)$ might be less than K .

This model is a variation of the Balls and Bins model, where we have K balls, C buckets of different color each containing B bins and every ball selects one bin from each bucket at random independent of the other balls.

Here, we also recall some terminology from graph theory that we will use later. A walk of size ℓ in graph \mathcal{G} starting from a node $v \in [K]$ is a set of ℓ edges e_1, e_2, \dots, e_ℓ , where v is one of the vertices of the edge e_1 and where consecutive edges are different, $e_i \neq e_{i+1}$, but incident with each other. A directed neighborhood of an edge $e = (v, c)$ of depth ℓ is the induced subgraph in \mathcal{G} consisting of all edges and associated check and variable nodes in all walks of size $\ell + 1$ starting from v with the first edge being $e_1 = (v, c)$. An edge e is said to have a tree neighborhood of depth ℓ if the directed neighborhood of e of depth ℓ is a tree.

Ensemble of Graphs Generated by Hashing

In the very sparse regime ($0 < \alpha < \frac{1}{3}$), in order to keep the computational complexity of the hashing algorithm around $O(K \log_2(K))$, we constructed $C = \frac{1}{\alpha}$ different surjective hash functions $h_i : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^b$, $i \in [C]$, where $b \approx n\alpha$ and where for an $x \in \mathbb{F}_2^n$ with binary representation x_0^{n-1} , $h_i(x_0^{n-1}) = (x_{ib}, x_{ib+1}, \dots, x_{ib+b-1})$. We also explained that in the spectral domain, this operation is equivalent to hashing each spectral component labeled with $X_0^{n-1} \in \mathbb{F}_2^n$ into the bin labelled with $h_i(X_0^{n-1})$. Notice that by this hashing scheme there is a one-to-one relation between a spectral element X and its bin number in different hashes $(h_0(X), h_1(X), \dots, h_{C-1}(X))$.

Let V be a uniformly distributed random variable over \mathbb{F}_2^n . It is easy to check that in the binary representation of V , V_0^{n-1} are like i.i.d. unbiased bits. This implies that $h_0(V), h_1(V), \dots, h_{C-1}(V)$ will be independent from one another because they depend on disjoint subsets of V_0^{n-1} . Moreover, $h_i(V)$ is also uniformly distributed over \mathbb{F}_2^b .

Assume that X_1, X_2, \dots, X_K are K different variables in \mathbb{F}_2^n denoting the position of non-zero spectral components. For these K variables and hash functions h_i , we can associate a bipartite graph as follows. We consider K variable nodes corresponding to X_1^K and C different set of check nodes S_0, S_1, \dots, S_{C-1} each of size $B = 2^b$. The check nodes in each S_i are labelled by elements of \mathbb{F}_2^b . For each variable X_i we consider C different edges connecting X_i to check nodes labelled with $h_j(X_i) \in S_j$, $j \in [C]$.

Proposition 3.4. *Let $h_i : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^b$, $i \in [C]$ be as in Equation (3.9). Let V_1, V_2, \dots, V_K be a set of variables generated from the ensemble $RS2(K, N)$, $N = 2^n$ denoting the position of non-zero components. The bipartite graph associated with variables V_1^K and hash functions h_i is a graph from the ensemble $\mathcal{G}(K, B, C)$, where $B = 2^b$.*

Proof. As V_1^K belong to the ensemble $RS2(N, K)$, they are i.i.d. variables uniformly distributed in $[N]$. This implies that for a specific V_i , $h_j(V_i)$, $j \in [C]$ are indepen-

dent from one another. Thus, every variable node selects its neighbor checks in S_0, S_1, \dots, S_{C-1} completely at random. Moreover, for any $j \in [C]$, the variables $h_j(V_1), \dots, h_j(V_K)$ are also independent, thus each variable selects its neighbor checks in S_j independent of all other variables. This implies that in the corresponding bipartite graph, every variable node selects its C check neighbors completely at random independent of all the other variable nodes, thus it belongs to $\mathcal{G}(K, B, C)$. \square

In Section 3.5, we explained the peeling decoder over the bipartite graph induced by the non-zero spectral components. It is easy to see that the performance of the algorithm always improves if we remove some of the variable nodes from the graph because it potentially reduces the number of colliding variables in the graph and there is more chance for the peeling decoder to succeed decoding.

Proposition 3.5. *Let $\alpha, C, K, h_i, i \in [C]$ be as in Proposition 3.4. Let \mathcal{G} be the bipartite graph induced by the random support set V_1^K generated from RS1 and hash functions h_i . For any $\epsilon > 0$, asymptotically as N tends to infinity, the average failure probability of the peeling decoder over \mathcal{G} is upper bounded by its average failure probability over the ensemble $\mathcal{G}(K(1 + \epsilon), B, C)$.*

Proof. Let \mathcal{G}_ϵ be a graph from ensemble $\mathcal{G}(K(1 + \epsilon), B, C)$. From Proposition 3.3, asymptotically the number of variable nodes in \mathcal{G}_ϵ is greater than K . If we drop some of the variable nodes at random from \mathcal{G}_ϵ to keep only K of them we obtain a graph from ensemble \mathcal{G} . From the explanation of the peeling decoder, it is easy to see that the performance of the decoder improves by removing some of the variable nodes because in that case less variables are collided together in different bins and there is more chance to peel them off. This implies that peeling decoder performs strictly better over \mathcal{G} rather than \mathcal{G}_ϵ . \square

Remark 3.4. *If we consider the graph induced by V_1^K from RS1 and hash functions h_i , the edge connection between variable nodes and check nodes is not completely random thus it is not compatible with Balls-and-Bins model explained before. Proposition 3.5 implies that asymptotically the failure probability for this model can be upper bounded by the failure probability of the peeling decoder for Balls-and-Bins model of slightly higher number of edges $K(1 + \epsilon)$.*

Edge Degree Distribution Polynomial

As we explained in the previous section, assuming a random support model for the non-zero spectral components in the very sparse regime $0 < \alpha < \frac{1}{3}$, we obtained a random graph from ensemble $\mathcal{G}(K, B, C)$. We also assumed that $n\alpha \in \mathbb{N}$ and we selected $b = n\alpha$, thus $K = B$. Let us call $\beta = \frac{K}{B}$ the average number of non-zero components per a hash bin. In our case, we designed hashes so that $\beta = 1$. As the resulting bipartite graph is left regular, all the variable nodes have degree C whereas for a specific check node the degree is random and depends on the graph realization.

Proposition 3.6. *Let $\mathcal{G}(K, B, C)$ be the random graph ensemble as before with $\beta = \frac{K}{B}$ fixed. Then asymptotically as N tends to infinity the check degree converges to a Poisson random variable with parameter β .*

Proof. Construction of the ensemble \mathcal{G} shows that any variable node has a probability of $\frac{1}{B}$ to be connected to a specific check node c , independent of all the other variable nodes. Let $Z_i \in \{0, 1\}$ be a Bernoulli random variable where $Z_i = 1$ if and only if variable i is connected to check node c . It is easy to check that the degree of c will be $Z = \sum_{i=1}^K Z_i$. The Characteristic function of Z can be easily obtained:

$$\begin{aligned}\Phi_Z(\omega) &= \mathbb{E}e^{j\omega Z} = \prod_{i=1}^K \mathbb{E}e^{j\omega Z_i} \\ &= \left(1 + \frac{1}{B}(e^{j\omega} - 1)\right)^{\beta B} \rightarrow e^{\beta(e^{j\omega} - 1)},\end{aligned}$$

showing the convergence of Z to a Poisson distribution with parameter β . \square

For a bipartite graph, the edge degree distribution polynomial is defined by $\rho(\alpha) = \sum_{i=1}^{\infty} \rho_i \alpha^{i-1}$ and $\lambda(\alpha) = \sum_{i=1}^{\infty} \lambda_i \alpha^{i-1}$, where ρ_i (λ_i) is the ratio of all edges that are connected to a check node (variable node) of degree i . Notice that we have $i - 1$ instead of i in the formula. This choice makes the analysis to be written in a more compact form as we will see.

Proposition 3.7. *Let \mathcal{G} be a random bipartite graph from the ensemble $\mathcal{G}(K, B, C)$ with $\beta = \frac{K}{B}$. Then $\lambda(\alpha) = \alpha^{C-1}$ and $\rho(\alpha)$ converges to $e^{-\beta(1-\alpha)}$ as N tends to infinity.*

Proof. From left regularity of a graph from ensemble \mathcal{G} , it results that all the edges are connected to variable nodes of degree C , thus $\lambda(\alpha) = \alpha^{C-1}$ and the number of edges is equal to CK . By symmetry of hash construction, it is sufficient to obtain the edge degree distribution polynomial for check nodes of the first hash. The total number of edges that are connected to the check nodes of the first hash is equal to K . Let N_i be the number of check nodes in this hash with degree i . By definition of ρ_i , it results that

$$\rho_i = \frac{iN_i}{K} = \frac{iN_i/B}{K/B}.$$

Let Z be the random variable as in the proof of Proposition 3.6 denoting the degree of a specific check node. Then, as N tends to infinity one can show that

$$\lim_{N \rightarrow \infty} \frac{N_i}{B} = \lim_{N \rightarrow \infty} \mathbb{P}(Z = i) = \frac{e^{-\beta} \beta^i}{i!} \text{ a.s.}$$

Thus ρ_i converges almost surely to $\frac{e^{-\beta} \beta^{i-1}}{(i-1)!}$. As $\rho_i \leq 1$, for any $\alpha : |\alpha| < 1 - \epsilon$, $|\rho_i \alpha^{i-1}| \leq (1 - \epsilon)^{i-1}$ and applying the *Dominated Convergence Theorem*, $\rho(\alpha)$ converges to $\sum_{i=1}^{\infty} \frac{e^{-\beta} \beta^{i-1}}{(i-1)!} \alpha^{i-1} = e^{-\beta(1-\alpha)}$. \square

Average Check Degree Parameter β

In the very sparse regime, as we explained that assuming that $b = n\alpha$ is an integer, we designed independent hashes with $B = 2^b$ output bins so that $\beta = \frac{K}{B} = 1$. As we will see the performance of the peeling decoder (described later by the DE equation in (3.10)) depends on the parameter β . The smaller β the better the performance of

the peeling decoder. Also notice that decreasing β via increasing B increases the time complexity $O(B \log_2(B))$ of computing the hash functions. For the general case, one can select B such that $\beta \in [1, 2)$ or at the cost of increasing the computational complexity, one can make β smaller, e.g., $\beta \in [\frac{1}{2}, 1)$, to obtain a better performance.

3.6.3 Performance Analysis of the Peeling Decoder

Assume that \mathcal{G} is the random bipartite graph resulting from applying C hashes to signal spectrum. As explained in Section 3.5, the iterative peeling algorithm starts by finding a singleton (check node of degree 1 which contains only one variable node or non-zero spectral component). The decoder peels off this variable node and removes all the edges connected to it from the graph. The algorithm continues by peeling off a singleton at each step until all the check nodes are zero-ton, i.e., all the non-zero variable nodes are decoded, or all the remaining unpeeled check nodes are multiton in which case the algorithm fails to completely decode all the spectral variables.

Wormald's Method

In order to analyze the behavior of the resulting random graphs under the peeling decoding, the authors in [50] applied Wormald's method to track the ratio of edges in the graph connected to check nodes of degree 1 (singleton). The essence of Wormald's method is to approximate the behavior of a stochastic system (here the random bipartite graph), after applying suitable time normalization, by a deterministic differential equation. The idea is that asymptotically as the size of the system becomes large (thermodynamic limit), the random state of the system is, uniformly for all times during the run of the algorithm, well concentrated around the solution of the differential equation. In [50], this method was applied to analyze the performance of the peeling decoder for bipartite graph codes over the BEC. We briefly explain the problem setting in [50] and how the results proved there can be extended to our case.

Assume that we have a bipartite graph \mathcal{G} with k variable nodes at the left, ck check nodes at the right and with edge degree polynomials $\lambda(x)$ and $\rho(x)$. We can define a channel code $\mathcal{C}(\mathcal{G})$ over this graph as follows. We assign k independent message bits to k input variable nodes. The output of each check node is the modulo 2 summation (XOR or summation over \mathbb{F}_2) of the message bits that are connected to it. Thus, the resulting code will be a systematic code with k message bits along with ck parity check bits. To communicate a k bit message over the channel, we send k message bits and all the check bits associated with them. While passing through the BEC, some of the message bits or check bits are erased independently. Assume a specific case in which the message bits and check bits are erased independently with probability δ and δ' respectively. Those message bits that pass perfectly through the channel are successfully transmitted, thus, the decoder tries to recover the erased message bits from the redundant information received via check bits. If we consider the induced graph after removing all variable nodes and check nodes corresponding to the erased ones from \mathcal{G} , we end up with another bipartite graph \mathcal{G}' . It is easy to see that over the new graph \mathcal{G}' , one can apply the peeling decoder to recover the erased bits.

In [50], this problem was fully analyzed for the case of $\delta' = 0$, where all the check bits are received perfectly but δ ratio of the message bits are erased independently from one another. In other words, the final graph \mathcal{G}' has on average $k\delta$ variable nodes to be decoded. Therefore, the analysis can be simply applied to our case, by assuming that $\delta \rightarrow 1$, where all the variable nodes are erased (they are all unknown and need to be identified). Notice that from the assumption $\delta' = 0$ no check bit is erased as is the case in our problem, i.e., we have access to all the hash outputs. In particular, Proposition 2 in [50] states that

Proposition 2 in [50]: Let \mathcal{G} be a bipartite graph with edge degrees specified by $\lambda(x)$ and $\rho(x)$ and with k message bits chosen at random. Let δ be fixed so that

$$\rho(1 - \delta\lambda(x)) > 1 - x, \quad \text{for } x \in (0, 1].$$

For any $\eta > 0$, there is some k_0 such that for all $k > k_0$, if the message bits of $\mathcal{C}(\mathcal{G})$ are erased independently with probability δ , then with probability at least $1 - k^{\frac{2}{3}} \exp(-\sqrt[3]{k}/2)$ the recovery algorithm terminates with at most ηk message bits erased.

Replacing $\delta = 1$ in the proposition above, we obtain the following performance guarantee for the peeling decoder.

Proposition 3.8. *Let \mathcal{G} be a bipartite graph from the ensemble $\mathcal{G}(K, B, C)$ induced by hashing functions $h_i, i \in [C]$ with $\beta = \frac{K}{B}$ and edge degree polynomials $\lambda(x) = x^{C-1}$ and $\rho(x) = e^{-\beta(1-x)}$ such that*

$$\rho(1 - \lambda(x)) > 1 - x, \quad \text{for } x \in (0, 1].$$

Given any $\epsilon \in (0, 1)$, there is a K_0 such that for any $K > K_0$ with probability at least $1 - K^{\frac{2}{3}} \exp(-\sqrt[3]{K}/2)$ the peeling decoder terminates with at most ϵK unrecovered non-zero spectral components.

Proposition 3.8 does not guarantee the success of the peeling decoder. It only implies that with very high probability, it can peel off any ratio $\eta \in (0, 1)$ of non-zero components but not necessarily all of them. However, using a combinatorial argument, it is possible to prove that with very high probability any graph in the ensemble \mathcal{G} is an expander graph, i.e., every small enough subset of left nodes has many check neighbors. This implies that if the peeling decoder can decode a specific ratio of variable nodes, it can proceed to decode all of them. A slight modification of Lemma 1 in [50] gives the following result proved in Section 3.10.4.

Proposition 3.9. *Let \mathcal{G} be a graph from the ensemble $\mathcal{G}(K, B, C)$ with $C \geq 3$. There is some $\eta > 0$ such that with probability at least $1 - O(\frac{1}{K^{\frac{1}{3}(C/2-1)}})$, the recovery process restricted to the subgraph induced by any η -fraction of the left nodes terminates successfully.*

Proof of Part 3 of Theorem 3.1 for $\alpha \in (0, \frac{1}{3}]$: In the very sparse regime $\alpha \in (0, \frac{1}{3}]$, we construct $C = \lceil \frac{1}{\alpha} \rceil \geq 3$ hashes each containing $2^{n\alpha}$ output bins. Combining Proposition 3.8 and 3.9, we obtain that the success probability of the peeling decoder is at least $1 - O(\frac{1}{K^{\frac{1}{3}(C/2-1)}})$ as mentioned in Remark 3.1.




Figure 3.6 – Tree-like neighborhood of an edge $e = (v, c)$. Dashed lines show the edges that have been removed before iteration t . The edge e is peeled off at iteration t because all the variable nodes v' connected to c are already decoded, thus, c is a singleton check.

Analysis based on Belief Propagation over Sparse Graphs

In this section, we give another method of analysis and further intuition about the performance of the peeling decoder and why it works very well in the very sparse regime. This method is based on the analysis of BP decoder over sparse locally tree-like graphs. The analysis is very similar to the analysis of the peeling decoder to recover non-zero frequency components in [47]. Consider a specific edge $e = (v, c)$ in a graph from ensemble $\mathcal{G}(K, B, C)$ and a directed neighborhood of this edge of depth ℓ as explained in Section 3.6.2. At the first stage, it is easy to see that this edge is peeled off from the graph assuming that all the edges (c, v') connected to the check node c are peeled off because in that case check c will be a singleton allowing to decode the variable v . This pictorially shown in Figure 3.6.

One can proceed in this way in the directed neighborhood to find the condition under which the variable v' connected to c can be peeled off and so on. Assuming that the directed neighborhood is a tree, all the messages that are passed from the leaves up to the head edge e are independent from one another. Let p_ℓ be the probability that edge e is peeled off depending on the information received from the directed neighborhood of depth ℓ assuming a tree up to depth ℓ . A simple analysis similar to [47], gives the following recursion

$$p_{j+1} = \lambda(1 - \rho(1 - p_j)), \quad j \in [\ell], \quad (3.10)$$

where λ and ρ are the edge degree polynomials of the ensemble \mathcal{G} . This iteration shows the progress of the peeling decoder in recovering unknown variable nodes. In [47], it was proved that for any specific edge e , asymptotically with very high probability the directed neighborhood of e up to any fixed depth ℓ is a tree. Specifically, if we start from a left regular graph \mathcal{G} from $\mathcal{G}(K, B, C)$ with KC edges, after ℓ steps of decoding, the average number of unpeeled edges is concentrated around KCp_ℓ . Moreover, a martingale argument was applied in [47] to show that not only the average of unpeeled edges is approximately KCp_ℓ but also with very high probability the number of those edges is well concentrated around KCp_ℓ .




Figure 3.7 – Density Evolution equation for $C = 3$ and different values of $\beta = \frac{K}{B}$.

Equation (3.10) is in general known as density evolution equation. Starting from $p_0 = 1$, this equation fully predicts the behavior of the peeling decoding over the ensemble \mathcal{G} . Figure 3.7 shows a typical behavior of this iterative equation for different values of the parameter $\beta = \frac{K}{B}$.

For very small values of β , this equation has only one fixed point at 0 which implies that asymptotically the peeling decoder can recover a fraction of variables very close to 1. However, for large values of β , e.g., $\beta \gtrsim 2.44$ for $C = 3$, this equation has a fixed point greater than 0. The largest fixed point is the place where the peeling decoder stops and can not proceed to decode the remaining variables. It is straightforward to check that the only fixed point is 0 provided that for any $p \in (0, 1]$, $p > \lambda(1 - \rho(1 - p))$. As $\lambda : [0, 1] \rightarrow [0, 1]$, $\lambda(x) = x^{C-1}$ is an increasing function of x , by change of variable $x = \lambda^{-1}(p)$, one obtains that $x > 1 - \rho(1 - \lambda(x))$ or equivalently

$$\rho(1 - \lambda(x)) > 1 - x, \quad \text{for } x \in (0, 1].$$

This is exactly the same result that we obtained by applying Wormald's method as in [50]. In particular, this analysis clarifies the role of x in Wormald's method.

Similar to Wormald's method, this analysis only guaranties that for any $\epsilon \in (0, 1)$, asymptotically as N tends to infinity, $1 - \epsilon$ fraction of the variable nodes can be recovered. An expander argument is still necessary to guarantee the full recovery of all the remaining variables.

3.7 Performance Analysis of the Less Sparse Regime

For the less sparse regime ($\frac{1}{3} < \alpha < 1$), similar to the very sparse case, we will first construct suitable hash functions that guarantee a low computational complexity of

order $O(K \log_2(K) \log_2(\frac{N}{K}))$ for the recovery of non-zero spectral values. Assuming a uniformly random support model in the spectral domain, similar to the very sparse case, one can represent the hashes by a regular bipartite graph. Over this graph, the peeling algorithm proceeds to find singleton checks and peel the associated variables from the graph until no singleton remains. The recovery is successful if all the variables are peeled off, thus, all the remaining checks are zero-ton. Otherwise some of the non-zero spectral values are not recovered and the recovery fails.

As we will explain, the structure of the induced bipartite graph in this regime is a bit different from the very sparse one. The following steps are used to analyze the performance of the peeling decoder:

1. Constructing suitable hash functions
2. Representing hash functions by their equivalent bipartite graph
3. Analyzing the performance of the peeling decoder over the resulting graph

For simplicity, we consider the case where $\alpha = 1 - \frac{1}{C}$ for some integer $C \geq 3$. In Section 3.7.4, we will explain how to deal with arbitrary values of C and α , specially those values of α in the range $(\frac{1}{3}, \frac{2}{3})$.

3.7.1 Hash Construction

Assume that $\alpha = 1 - \frac{1}{C}$ for some integer $C \geq 3$. Let x be an N -dimensional signal with $N = 2^n$ and let X denote its WHT. For simplicity, we label the components of X by a binary vector $X_0^{n-1} \in \mathbb{F}_2^n$. Let $t = \frac{n}{C}$ and let us divide the set of n binary indices X_0^{n-1} into C non-intersecting subsets r_0, r_1, \dots, r_{C-1} , where $r_i = X_{it}^{(i+1)t-1}$. Hence, there is a one-to-one relation between each binary vector $X_0^{n-1} \in \mathbb{F}_2^n$ and its representation $(r_0, r_1, \dots, r_{C-1})$. We construct C different hash functions $h_i, i \in [C]$ by selecting different subsets of $(r_0, r_1, \dots, r_{C-1})$ of size $C - 1$ and appending them together. For example

$$h_1(X_0^{n-1}) = (r_0, r_1, \dots, r_{C-2}) = X_0^{(C-1)t-1},$$

and the hash output is obtained by appending $C - 1$ first $r_i, i \in [C]$. One can simply check that $h_i, i \in [C]$ are linear surjective functions from \mathbb{F}_2^n to \mathbb{F}_2^b , where $b = (C - 1)t$. In particular, the range of each hash consists of $B = 2^b$ different elements of \mathbb{F}_2^b . Moreover, if we denote the null space of h_i by $\mathcal{N}(h_i)$, it is easy to show that for any $i, j \in [C], i \neq j, \mathcal{N}(h_i) \cap \mathcal{N}(h_j) = \mathbf{0} \in \mathbb{F}_2^n$.

Using the subsampling property of the WHT and similar to the hash construction that we had in Subsection 3.6.1, it is seen that subsampling the time domain signal and taking WHT of the subsampled signal is equivalent to hashing the spectral components of the signal. In particular, all the spectral components X_0^{n-1} with the same $h_i(X_0^{n-1})$ are mapped into the same bin in the hash i , thus, different bins of the hash can be labelled with B different elements of \mathbb{F}_2^b .

It is easy to see that, with this construction, the average number of non-zero elements per bin in every hash is kept at $\beta = \frac{K}{B} = 1$ and the complexity of computing all the hashes along with their $n - b$ shifts, which are necessary for collision detection/support estimation, is $CK \log_2(K) \log_2(\frac{N}{K})$. The sample complexity can also be easily checked to be $CK \log_2(\frac{N}{K})$.




Figure 3.8 – Bipartite graph representation for the less sparse case $\alpha = \frac{2}{3}$, $C = 3$

3.7.2 Bipartite Graph Representation

Similar to the very sparse regime, we can assign a bipartite graph with K left nodes (variable nodes) corresponding to non-zero spectral components and with CK right nodes corresponding to different bins of all the hashes. In particular, we consider C different set of check nodes S_1, S_2, \dots, S_C each containing B nodes labelled with the elements of \mathbb{F}_2^b and a specific non-zero spectral component labelled with X_0^{n-1} is connected to nodes $s_i \in S_i$ if and only if the binary label assigned to s_i is $h_i(X_0^{n-1})$. In the very sparse regime, we showed that if the support of the signal is generated according to RS2(K, N) model, where K random positions are selected uniformly at random independent from one another from $[N]$, then the resulting graph is a random left regular bipartite graph, where each variable nodes select its C neighbors in S_1, S_2, \dots, S_C completely independently. However, in the less sparse regime, the selection of the neighbor checks in different hashes is not completely random. To explain more, let us assume that $\alpha = \frac{2}{3}$, thus $C = 3$. Also assume that for a non-zero spectral variable labelled with X_0^{n-1} , r_i denotes $X_{it}^{(i+1)t-1}$, where $t = \frac{n}{C}$. In this case, this variable is connected to bins labelled with (r_0, r_1) , (r_1, r_2) and (r_0, r_2) in 3 different hashes. This has been pictorially shown in Figure 3.8.

If we assume that X_0^{n-1} is selected uniformly at random from \mathbb{F}_2^n then the bin numbers in each hash, e.g., (r_0, r_1) in the first hash, are individually selected uniformly at random among all possible bins. However, it is easily seen that the joint selection of bins is not completely random among different hashes. In other words, the associated bins in different hashes are not independent from one another. However, assuming the random support model, where K variable V_1^K are selected independently as the position of non-zero spectral variables, the bin association for different variables V_i is still done independently.

3.7.3 Performance Analysis of the Peeling Decoder

As the resulting bipartite graph is not a completely random graph, it is not possible to directly apply Wormald's method as we did for the very sparse case as in [50]. However, an analysis based on the DE for the BP algorithm can still be applied. In

other words, setting $p_0 = 1$ and

$$p_{j+1} = \lambda(1 - \rho(1 - p_j)), \quad j \in [\ell],$$

as in (3.10) with λ and ρ being the edge degree polynomials of the underlying bipartite graph, it is still possible to show that after ℓ steps of decoding the average number of unpeeled edges is approximately KCp_ℓ . A martingale argument similar to [47] can be applied to show that the number of remaining edges is also well concentrated around its average. Similar to the very sparse case, this argument asymptotically guarantees the recovery of any fraction of the variables between 0 and 1. Another argument is necessary to show that if the peeling decoder decodes a majority of the variables, it can proceed to decode all of them with very high probability. To formulate this, we use the concept of trapping sets for the peeling decoder.

Definition 3.1. *Let $\alpha = 1 - \frac{1}{C}$ for some integer $C \geq 3$ and let $h_i, i \in [C]$ be a set of hash functions as explained before. A subset of variables $T \subset \mathbb{F}_2^n$ is called a trapping set for the peeling decoder if for any $v \in T$ and for any $i \in [C]$, there is another $w \in T, v \neq w$ such that $h_i(v) = h_i(w)$, thus colliding with v in the i -th hash.*

Notice that a trapping set can not be decoded because all of its neighbor check nodes are multiton. We first analyze the structure of the trapping set and find the probability that a specific set of variables build a trapping set. Let X be a spectral variable in the trapping set with the corresponding binary representation X_0^{n-1} and assume that $C = 3$. As we explained, we can equivalently represent this variable with (r_0, r_1, r_2) , where $r_i = X_{it}^{(i+1)t-1}$ with $t = \frac{n}{C}$. We can consider a three dimensional lattice whose i -th axis is labelled by all possible values of r_i . In this space, there is a simple interpretation for a set T to be a trapping set, namely, for any $(r_0, r_1, r_2) \in T$ there are three other elements $(r'_0, r_1, r_2), (r_0, r'_1, r_2)$ and (r_0, r_1, r'_2) in T that can be reached from (r_0, r_1, r_2) by moving along exactly one axis. Notice that in this case each hash is equivalent to projecting (r_0, r_1, r_2) onto two dimensional planes spanned by different coordinates, for example, $h_1(r_0, r_1, r_2) = (r_0, r_1)$ is a projection on the plane spanned by the first and second coordinate axes of the lattice. A similar argument holds for other values of $C > 3$, thus, larger values of α .

For $C \geq 3$, the set of all C -tuples $(r_0, r_1, \dots, r_{C-1})$ is a C -dimensional lattice. We denote this lattice by L . The intersection of this lattice by the hyperplane $R_i = r_i$ is a $(C - 1)$ -dimensional lattice defined by

$$L(R_i = r_i) = \{(r_0, \dots, r_{i-1}, r_{i+1}, \dots, r_{C-1}) : (r_0, r_1, \dots, r_{i-1}, r_i, r_{i+1}, \dots, r_{C-1}) \in L\}.$$

Similarly, for $S \subset L$, we have the following definition

$$S(R_i = r_i) = \{(r_0, \dots, r_{i-1}, r_{i+1}, \dots, r_{C-1}) : (r_0, r_1, \dots, r_{i-1}, r_i, r_{i+1}, \dots, r_{C-1}) \in S\}.$$

Obviously, $S(R_i = r_i) \subset L(R_i = r_i)$. We have the following proposition whose proof simply follows from the definition of the trapping set.

Proposition 3.10. *Assume that T is a trapping set for the C -dimensional lattice representation L of the non-zero spectral domain variables as explained before. Then for any r_i on the i -th axis, $T(R_i = r_i)$ is either empty or a trapping set for the $(C - 1)$ -dimensional lattice $L(R_i = r_i)$.*

Proposition 3.11. *The size of the trapping set for a C -dimensional lattice is at least 2^C .*

Proof. We use a simple proof using the induction on C . For $C = 1$, we have a one dimensional lattice along a line labelled with r_0 . In this case, there must be at least two variables on the line to build a trapping set. Consider a trapping set T of dimension C . There are at least two points $(r_0, r_1, \dots, r_{C-1})$ and $(r'_0, r_1, \dots, r_{C-1})$ in T . By Proposition 3.10, $T(R_0 = r_0)$ and $T(R_0 = r'_0)$ are two $(C - 1)$ -dimensional trapping sets each consisting of at least 2^{C-1} elements by induction hypothesis. Thus, T has at least 2^C elements. \square

Remark 3.5. *The bound $|T| \geq 2^C$ on the size of the trapping set is actually tight. For example, for $i \in [C]$ consider r_i, r'_i where $r_i \neq r'_i$ and let*

$$T = \{(a_0, a_1, \dots, a_{C-1}) : a_i \in \{r_i, r'_i\}, i \in [C]\}.$$

It is easy to see that T is a trapping set with 2^C elements corresponding to the vertices of a C -dimensional cube.

We now prove the following proposition showing that if the peeling decoder can decode all the variable nodes except a fixed number of them, with high probability, it can continue to decode all of them.

Proposition 3.12. *Let s be a fixed positive integer. Assume that $\alpha = 1 - \frac{1}{C}$ for some integer $C \geq 3$ and consider a hash structure with C different hashes as explained before. If the peeling decoder decodes all except a set of variables of size s , it can decode all the variables with very high probability.*

Proof. The proof follows from a similar proof in [47]. Let T be a trapping set of size s . By Proposition 3.11, we have $s \geq 2^C$. Let p_i be the number of distinct values taken by elements of T along the R_i axis and let $p_{\max} = \max_{i \in [C]} p_i$. Without loss of generality, let us assume that the R_0 axis is the one having the maximum p_i . Consider $T(R_0 = r_0)$ for those p_{\max} values of r_0 along the R_0 axis. Proposition 3.10 implies that each $T(R_0 = r_0)$ is a trapping set which has at least 2^{C-1} elements according to Proposition 3.11. This implies that $s \geq 2^{C-1} p_{\max}$ or $p_{\max} \leq \frac{s}{2^{C-1}}$. Moreover, T being the trapping set implies that there are subsets T_i consisting of elements from axes R_i and all the elements of T are restricted to take their i -th coordinate values along R_i from the set T_i . Considering the way we generate the position of the non-zero variables X_0^{n-1} with the equivalent representation $(r_0, r_1, \dots, r_{C-1})$, the coordinate of any variable is selected uniformly and completely independently from one another and from the coordinates of the other variables. This implies that

$$\begin{aligned} \mathbb{P}(F_s) &\leq \mathbb{P}\{\text{For any variables in } T, r_i \in T_i, i \in [C]\} \\ &\leq \prod_{i=0}^{C-1} \binom{\mathcal{P}_i}{p_i} \left(\frac{p_i}{\mathcal{P}_i}\right)^s \leq \prod_{i=0}^{C-1} \binom{\mathcal{P}_i}{s/2^{C-1}} \left(\frac{s}{2^{C-1}\mathcal{P}_i}\right)^s, \end{aligned}$$

where F_s is the event that the peeling decoder fails to decode a specific subset of variables of size s and where \mathcal{P}_i denotes the number of all possible values for the

54 Fast Hadamard Transform: Construction for Signals with Sub-linear Sparsity

i -th coordinate of a variable. By our construction, all \mathcal{P}_i are equal to $P = 2^{n/C} = 2^{n(1-\alpha)} = N^{(1-\alpha)}$, thus we obtain that

$$\begin{aligned} \mathbb{P}(F_s) &\leq \binom{P}{s/2^{C-1}}^C \left(\frac{s}{2^{C-1}P}\right)^{sC} \leq \left(\frac{2^{C-1}Pe}{s}\right)^{sC/2^{C-1}} \left(\frac{s}{2^{C-1}P}\right)^{sC} \\ &\leq \left(\frac{se^{1/(2^{C-1}-1)}}{2^{C-1}P}\right)^{sC(1-1/2^{C-1})}. \end{aligned}$$

Taking the union bound over all $\binom{K}{s}$ possible ways of selecting s variables out of K variables, we obtain that

$$\begin{aligned} \mathbb{P}(F) &\leq \binom{K}{s} \mathbb{P}(F_s) \leq \left(\frac{eP^{C-1}}{s}\right)^s \left(\frac{se^{1/(2^{C-1}-1)}}{2^{C-1}P}\right)^{sC(1-1/2^{C-1})} = O(1/P^{s(1-\frac{C}{2^{C-1}})}) \\ &\leq O(1/P^{(2^C-2^C)}) = O(1/N^{\frac{2^C}{C}-2}). \end{aligned}$$

For $C \geq 3$, this gives an upper bound of order $O(N^{-\frac{2}{3}})$ vanishing asymptotically. \square

3.7.4 Generalized Hash Construction

The hash construction that we explained only covers values of $\alpha = 1 - \frac{1}{C}$ for $C \geq 3$ which belongs to the region $\alpha \in [\frac{2}{3}, 1)$. We will explain a hash construction that extends to any value of C and $\alpha \in (0, 1)$, which is not necessarily of the form $1 - \frac{1}{C}$. This construction reduces to the very sparse and less sparse regimes when $\alpha = \frac{1}{C}$, $\alpha \in (0, 1/3]$, and $\alpha = 1 - \frac{1}{C}$, $\alpha \in [2/3, 1)$, respectively.

In the very sparse regime $\alpha = \frac{1}{3}$, we have $C = 3$ different hashes and for a non-zero spectral variable X with index $X_0^{n-1} = (r_0, r_1, r_2)$, $h_i(X_0^{n-1}) = r_i$, thus the output of different hashes depend on non-overlapping parts of the binary index of X whereas for $\alpha = \frac{2}{3}$ the hash outputs are (r_0, r_1) , (r_1, r_2) and (r_0, r_2) which overlap on a portion of binary indices of length $\frac{n}{3}$. Intuitively, it is clear that in order to construct different hashes for $\alpha \in (\frac{1}{3}, \frac{2}{3})$, we should start increasing the overlapping size of different hashes from 0 for $\alpha = \frac{1}{3}$ to $\frac{n}{3}$ for $\alpha = \frac{2}{3}$. We give the following construction for the hash functions

$$h_i(X_0^{n-1}) = X_{it}^{i+b}, i \in [C],$$

where $t = \frac{n}{C}$ and the values of the indices are computed modulo n , for example $X_n = X_0$. In the terminology of Section 3.4, we pick $\mathcal{H}_i = \Psi_b^T \Sigma_i^T \in \mathbb{F}_2^{k \times n}$, where $\Sigma_i \in \mathbb{F}_2^{n \times n}$ is the identity matrix with columns circularly shifted by $(i+1)b$ to the left. It is clear that each hash is a surjective map from \mathbb{F}_2^n into $\mathbb{F}_2^{n\alpha}$. Therefore, if we pick $b = n\alpha$, the number of output bins in each hash is $B = 2^{n\alpha} = N^\alpha = K$, thus the average number of non-zero variables per bin in every hash is equal to $\beta = \frac{K}{B} = 1$. In terms of decoding performance for the intermediate values of $\alpha \in (\frac{1}{3}, \frac{2}{3})$, one expects that the performance of the peeling decoder for this regime is between the very sparse regime $\alpha = \frac{1}{3}$ and the less sparse one $\alpha = \frac{2}{3}$.




Figure 3.9 – Probability of success of the algorithm as a function of α and C for deterministic hash construction. The dimension of the signal is $N = 2^{22}$. The black line corresponds to $\alpha = \frac{1}{C}$ and $\alpha = 1 - \frac{1}{C}$ in the very sparse and less sparse regimes respectively. We fix $\beta = 1$. The hashing matrices are deterministically picked as described in Section 3.7.4.

3.8 Simulation Results

In this section, we empirically evaluate the performance of the SFHT algorithm for a variety of design parameters. The simulations are implemented in C programming language and the success probability of the algorithm has been estimated via sufficient number of trials. We also provide a comparison of the run time of our algorithm and the standard Hadamard transform.

- *Experiment 1:* We fix the signal size to $N = 2^{22}$ and run the algorithm 1000 times to estimate the success probability for all range of $\alpha \in (0, 1)$. For every value of C , we use the deterministic hashing scheme as described in Section 3.7.4. Figure 3.9 shows the simulation result. The solid line shows our proposed design with optimal computational complexity, where in the very sparse regime $\alpha = \frac{1}{C}$ and in the less sparse regime $\alpha = 1 - \frac{1}{C}$ for some $C \geq 3$. It is seen that over the solid line, the success probability is very close to 1. Moreover, if the computational complexity is not important, the generalized hash design (as explained in Section 3.7.4) for $C = 6$ guarantees a good performance over all values of $\alpha \in (0, 1)$.
- *Experiment 2:* We repeat experiment 1, but instead of deterministic hashing matrices designed according to Section 3.7.4, we now pick $\Sigma_i, i \in [C]$, for hash construction uniformly at random from $\text{GL}(n, \mathbb{F}_2)$. The result is shown in Figure 3.10. We observe that the performance of this scheme is comparable to the deterministic one.




Figure 3.10 – Probability of success of the algorithm as a function of α and C for random hash construction. The dimension of the signal is $N = 2^{22}$. The black line corresponds to $\alpha = \frac{1}{C}$ and $\alpha = 1 - \frac{1}{C}$ in the very sparse and less sparse regimes respectively. We fix $\beta = 1$. The hashing matrices are picked uniformly at random for every trial.

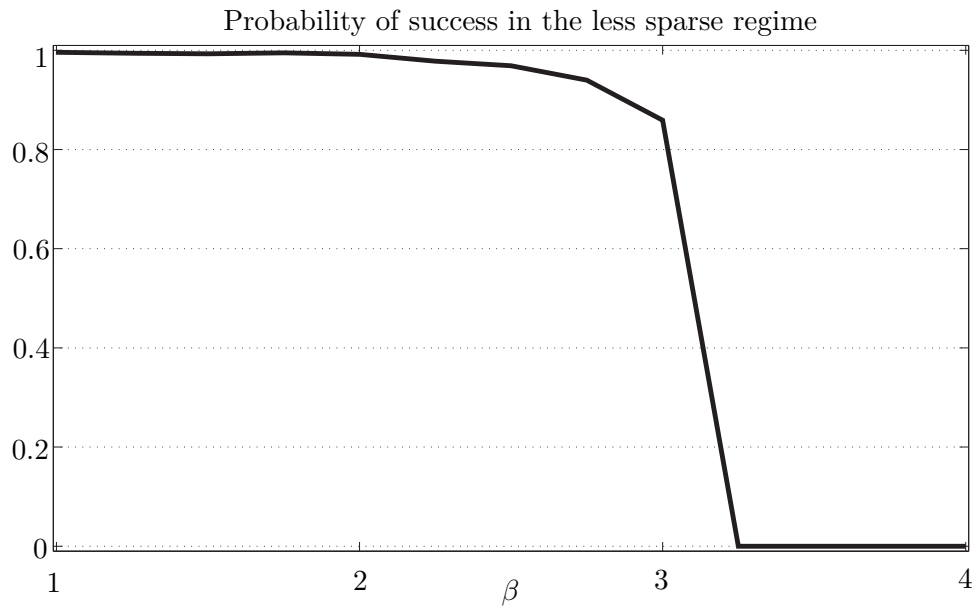


Figure 3.11 – Probability of success of the algorithm in the less sparse regime as a function of $\beta = K/B$. We fix $N = 2^{22}$, $B = 2^{17}$, $C = 4$, and vary α in the range 0.7 to 0.9.




Figure 3.12 – Comparison of the Median runtime in ms of the SFHT and conventional WHT for $N = 2^{15}$ and for different values of α .

- *Experiment 3:* In this experiment, we investigate the sensitivity of the algorithm to the value of the parameter $\beta = K/B$; the average number of non-zero coefficients per bin. As we explained, in our hash design we use $\beta \approx 1$. However, using larger values of β is appealing from a computational complexity point of view. For the simulation, we fix $N = 2^{22}$, $B = 2^{17}$, $C = 4$, and vary α between 0.7 and 0.9, thus changing K and as a result β . Figure 3.11 shows the simulation results. It is seen that the success probability has a sharp transition for $\beta \approx 3$.
- *Runtime measurement:* We compare the runtime of the SFHT algorithm with the traditional Walsh-Hadamard transform. The result is shown in Figure 3.12 for $N = 2^{15}$. SFHT performs much faster for $0 < \alpha < 2/3$.

It is also interesting to identify the range of α for which SFHT has a better runtime than the traditional WHT. We define α^* as the largest value of α such that SFHT is faster than WHT, i.e.,

$$\alpha^* = \sup_{\alpha \in (0,1)} \{\alpha : T_{\text{WHT}}(n) > T_{\text{SFHT}}(\alpha', n)\},$$

where T_{WHT} and T_{SFHT} are the runtimes of the conventional WHT and SFHT, respectively. We plot α^* as a function of $n = \log_2 N$ in Figure 3.13. It is seen that as n increases, SFHT performs better than the WHT over a larger range of α .

For computing the the run-time complexity of the SFHT in Section 3.5.2, we have assumed that matrix-vector multiplications in \mathbb{F}_2^n can be done in $O(1)$. The reason is that the deterministic hashing scheme of the algorithm is nothing but a circular

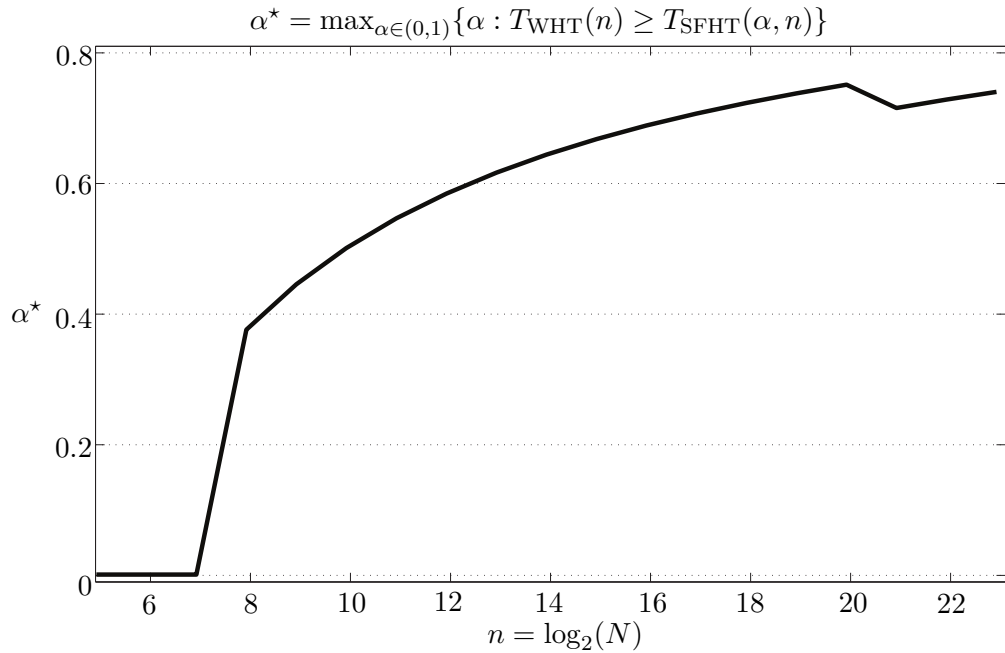


Figure 3.13 – In this figure, we plot $n = \log_2 N$ against α^* , the largest value of α such that SFHT runs faster than the conventional WHT. When WHT is always faster, we simply set $\alpha^* = 0$.

bit shift that can be implemented in a constant number of operations, independently of the vector size n .

If one is given Σ , some matrix from $\mathbb{F}_2^{n \times n}$, and its inverse transpose Σ^{-T} , the overall complexity of the algorithm would nonetheless be unchanged. First, we observe that it is possible to compute the inner product of two vectors in constant time using bitwise operations and a small look-up table². Now, given the structure of Ψ_b , computing $\Sigma \Psi_b m$ in Algorithm 1 only requires $\log_2 K$ inner products. Thus the complexity of Algorithm 1 is unchanged. Finally, (3.8) can be split into pre-computing $\Sigma^{-T} \Psi_b k$ at the same time as we subsample the signal (in $O(\log_2 K)$), and computing the inner product between \hat{v} and the $n - b$ first columns of Σ when doing the decoding ($O(\log_2 \frac{N}{K})$).

3.9 Conclusion

In this chapter, we presented a low-complexity iterative algorithm to compute the Walsh-Hadamard transform of a signal of length N . In particular, we assumed that the signal is K -sparse in the Hadamard domain with $K = O(N^\alpha)$ scaling sub-linearly with N for some $\alpha \in (0, 1)$. This equivalently provides a deterministic matrix construction along with a low-complexity recovery algorithm for compressed sensing of the signals with sub-linear sparsity. The algorithm has complexity $O(K \log_2 K \log_2 \frac{N}{K})$ and only requires $O(K \log_2 \frac{N}{K})$ time-domain samples (measurements in the compressed sensing setting). We showed that the algorithm reconstructs the Hadamard transform of the signal with a very high probability asymptotically approaching one.

²<http://graphics.stanford.edu/~seander/bithacks.html#ParityLookupTable>

The performance of the algorithm is also evaluated empirically through simulation, and its speed is compared to that of the conventional fast Hadamard transform.

Our algorithm works for the noiseless case where there is no measurement noise as can be seen from the the statement of Proposition 3.2. To make the algorithm fully practical, a robust estimator is needed to replace Proposition 3.2. This can be a direction for future work.

3.10 Proof of the Auxiliary Results

3.10.1 Proof of the Properties of the WHT

Proof of Property 1

$$\sum_{m \in \mathbb{F}_2^n} (-1)^{\langle k, m \rangle} x_{m+p} = \sum_{m \in \mathbb{F}_2^n} (-1)^{\langle k, m+p \rangle} x_m.$$

The proof follows by taking $(-1)^{\langle k, p \rangle}$ out of the sum and recognizing the Hadamard transform of x_m . ■

Proof of Property 2

As we explained, it is possible to assign an $N \times N$ matrix Π to the permutation π as follows

$$(\Pi)_{i,j} = \begin{cases} 1 & \text{if } j = \pi(i) \Leftrightarrow i = \pi^{-1}(j) \\ 0 & \text{otherwise.} \end{cases}$$

Let π_1 and π_2 be the permutations associated with Π_1 and Π_2 . Since $(H_N)_{i,j} = (-1)^{\langle i, j \rangle}$, the identity (3.2) implies that

$$(-1)^{\langle \pi_2(i), j \rangle} = (-1)^{\langle i, \pi_1^{-1}(j) \rangle}.$$

Therefore, for any $i, j \in \mathbb{F}_2^n$, π_1, π_2 must satisfy $\langle \pi_2(i), j \rangle = \langle i, \pi_1^{-1}(j) \rangle$. By linearity of the inner product, one obtains that

$$\begin{aligned} \langle \pi_2(i+k), j \rangle &= \langle i+k, \pi_1^{-1}(j) \rangle = \langle i, \pi_1^{-1}(j) \rangle + \langle k, \pi_1^{-1}(j) \rangle \\ &= \langle \pi_2(i), j \rangle + \langle \pi_2(k), j \rangle. \end{aligned}$$

As $i, j \in \mathbb{F}_2^n$ are arbitrary, this implies that π_2 , and by symmetry π_1 , are both linear operators. Hence, all the permutations satisfying (3.2) are in one-to-one correspondence with the elements of $\text{GL}(n, \mathbb{F}_2)$. ■

Proof of Property 3

Since Σ is non-singular, then Σ^{-1} exists. It follows from the definition of the WHT that

$$\sum_{m \in \mathbb{F}_2^n} (-1)^{\langle k, m \rangle} x_{\Sigma m} = \sum_{m \in \mathbb{F}_2^n} (-1)^{\langle k, \Sigma^{-1} m \rangle} x_m = \sum_{m \in \mathbb{F}_2^n} (-1)^{\langle \Sigma^{-T} k, m \rangle} x_m.$$

This completes the proof. ■

Proof of Property 4

$$\begin{aligned} \sum_{m \in \mathbb{F}_2^b} (-1)^{\langle k, m \rangle} x_{\Psi_b m} &= \frac{1}{\sqrt{N}} \sum_{m \in \mathbb{F}_2^b} (-1)^{\langle k, m \rangle} \sum_{p \in \mathbb{F}_2^n} (-1)^{\langle \Psi_b m, p \rangle} X_p \\ &= \frac{1}{\sqrt{N}} \sum_{p \in \mathbb{F}_2^n} X_p \sum_{m \in \mathbb{F}_2^b} (-1)^{\langle m, k + \Psi_b^T p \rangle}. \end{aligned}$$

In the last expression, if $p = \Psi_b k + i$ with $i \in \mathcal{N}(\Psi_b^T)$ then it is easy to check that the inner sum is equal to B , otherwise it is equal to zero. Thus, by proper renormalization of the sums one obtains the proof. \blacksquare

3.10.2 Proof of Proposition 3.2

We first show that if multiple coefficients fall in the same bin, it is very unlikely that part (1) is fulfilled. Let $\mathcal{I}_k = \{j \mid \mathcal{H}j = k\}$ be the set of variable indices hashed to bin k . This set is finite and its element can be enumerated as $\mathcal{I}_k = \{j_1, \dots, j_{\frac{N}{B}}\}$. We show that a set $\{X_j\}_{j \in \mathcal{I}_k}$ does not pass the ratio test unless it contains only one non-zero element. Without loss of generality, we consider $\sum_{j \in \mathcal{I}_k} X_j = 1$. Such $\{X_j\}_{j \in \mathcal{I}_k}$ is a solution of

$$\begin{bmatrix} 1 & \cdots & 1 \\ (-1)^{\langle \sigma_1, j_1 \rangle} & \cdots & (-1)^{\langle \sigma_1, j_{\frac{N}{B}} \rangle} \\ \vdots & \ddots & \vdots \\ (-1)^{\langle \sigma_{n-b}, j_1 \rangle} & \cdots & (-1)^{\langle \sigma_{n-b}, j_{\frac{N}{B}} \rangle} \end{bmatrix} \begin{bmatrix} X_{j_1} \\ \vdots \\ X_{j_{\frac{N}{B}}} \end{bmatrix} = \begin{bmatrix} 1 \\ \pm 1 \\ \vdots \\ \pm 1 \end{bmatrix},$$

where $\sigma_i, i \in \{1, \dots, n\}$ denotes the i -th column of the matrix Σ . The left hand side matrix in the expression above, is $(n-b+1) \times 2^{n-b}$. As $\sigma_1, \dots, \sigma_{n-b}$ form a basis for \mathcal{I}_k , all the columns are different and are (omitting the top row) the exhaustive list of all 2^{n-b} possible ± 1 vectors. Thus the right vector is always one of the columns of the matrix and there is a solution with only one non-zero component (1-sparse solution) to this system whose support can be uniquely identified. Adding any vector from the null space of the matrix to this initial solution yields another solution. However, as we will show, due to its structure this matrix is full rank and thus its null space has dimension $2^{n-b} - n + b - 1$. Assuming a continuous distribution on the non-zero components X_i , the probability that $\{X_i\}_{i \in \mathcal{I}_k}$ falls in this null space is zero.

To prove that the matrix is indeed full rank, let us first focus on the rank of the sub-matrix obtained by removing the first row. This submatrix itself always contains $M = -2I + \mathbf{1}\mathbf{1}^T$, where I is the identity matrix of order $n-b$ and $\mathbf{1}$ is the all-one vector of dimension $(n-b)$. One can simply check that M is a symmetric matrix, thus by spectral decomposition, it has $n-b$ orthogonal eigen-vectors $v_i, i \in [n-b]$. It is also easy to see that the normalized all-one vector $v_0 = \frac{\mathbf{1}}{\sqrt{n-b}}$ of dimension $n-b$ is an eigen-vector of M with eigen-value $\lambda_0 = n-b-2$. Moreover, assuming the orthonormality of the eigen-vectors, it results that $v_i^T M v_i = \lambda_i = -2$, where we used $v_i^T \mathbf{1} = v_i^T v_0 = 0$ for $i \neq 0$. Thus, for $n-b \neq 2$ all the eigen-values are non-zero and M is invertible, which implies that the sub-matrix resulted after removing the first row is full rank. In the case where $n-b = 2$, one can notice that the Hadamard

matrix of size 2 will be contained as a submatrix, and thus the matrix will be full rank.

Now it remains to prove that initial matrix is also full rank with a rank of $n - b + 1$. Assume that the columns of the matrix are arranged in the lexicographical order such that neglecting the first row, the first and the last column are all 1 and all -1 . If we consider any linear combination of the rows except the first one, it is easy to see that the first and the last element in the resulting row vector have identical magnitudes but opposite signs. This implies that the all-one row cannot be written as a linear combination of the other rows of the matrix. Therefore, the rank of the matrix must be $n - b + 1$.

To prove (3.8), let Σ_L and Σ_R be the matrices containing respectively the first $n - b$ and the last b columns of Σ , such that $\Sigma = [\Sigma_L \Sigma_R]$. If there is only one coefficient in the bin, then (3.7) implies that $\hat{v} = [(j^T \Sigma_L) \ 0]^T$. Using definitions (3.3) and (3.6), we obtain that $\Psi_b \mathcal{H}j = [0 \ (j^T \Sigma_R)]^T$. We observe that they sum to $\Sigma^T j$ and the proof follows. ■

3.10.3 Proof of Proposition 3.3

For $t \in [K]$, let H_t denote the size of the random set obtained by picking t objects from $[N]$ independently and uniformly at random with replacement. Let a_t and v_t denote the average and the variance of H_t for $t \in [K]$. It is easy to see that $\{H_t\}_{t \in [K]}$ is a Markov process. Thus, we have

$$\mathbb{E}\{H_{t+1} - H_t | H_t\} = (1 - H_t/N),$$

because the size of the random set increases if and only if we choose an element from $[N] \setminus H_t$. This implies that $a_{t+1} = 1 + \gamma a_t$, where $\gamma = 1 - \frac{1}{N}$. Solving this equation we obtain that

$$a_t = \sum_{r=0}^t \gamma^r = \frac{1 - \gamma^{t+1}}{1 - \gamma} = N(1 - \gamma^{t+1}). \quad (3.11)$$

In particular, $a_K = N(1 - (1 - \frac{1}{N})^K)$, which implies that $\mathbb{E}\{\frac{H_K}{K}\} = \frac{N}{K}(1 - (1 - \frac{1}{N})^K)$. One can check that for $K = N^\alpha$, $0 < \alpha < 1$, as N tends to infinity $\mathbb{E}\frac{H_K}{K}$ converges to 1. To find the variance of H_t , we use the formula

$$\text{Var}(H_{t+1}) = \text{Var}(H_{t+1} | H_t) + \text{Var}(\mathbb{E}\{H_{t+1} | H_t\}). \quad (3.12)$$

Therefore, we obtain that

$$\text{Var}(\mathbb{E}H_{t+1} | H_t) = \text{Var}(1 + \gamma H_t) = \gamma^2 v_t. \quad (3.13)$$

Moreover, for the first part in (3.12), we have

$$\begin{aligned} \text{Var}(H_{t+1} | H_t) &= \mathbb{E}_{H_t} \{ \text{Var}(H_{t+1} | H_t = h_t) \} = \mathbb{E}_{H_t} \{ \text{Var}(H_{t+1} - H_t | H_t = h_t) \} \\ &\stackrel{(I)}{=} \mathbb{E} \left\{ \frac{H_t}{N} \left(1 - \frac{H_t}{N} \right) \right\} = \frac{a_t}{N} + \frac{a_t^2 + v_t}{N^2}, \end{aligned} \quad (3.14)$$

where in (I) we used the fact that given H_t , $H_{t+1} - H_t$ is a Bernoulli random variable with probability $\frac{H_t}{N}$, thus its variance is equal to $\frac{H_t}{N}(1 - \frac{H_t}{N})$. Combining (3.13) and

(3.14), we obtain that

$$v_{t+1} = \left(\gamma^2 + \frac{1}{N^2} \right) v_t + \frac{a_t}{N} \left(1 + \frac{a_t}{N} \right). \quad (3.15)$$

From (3.11), it is easy to see that a_t is increasing in t . Moreover, from (3.15), it is seen that v_{t+1} is increasing function of a_t , thus, if we consider the following recursion

$$w_{t+1} = \left(\gamma^2 + \frac{1}{N^2} \right) w_t + \frac{a_K}{N} \left(1 + \frac{a_K}{N} \right),$$

then for any $t \in [K]$, $v_t \leq w_t$. As w_t is also an increasing sequence of t , we have

$$\begin{aligned} v_K &\leq w_K \leq w_\infty = \frac{a_K}{N} \left(1 + \frac{a_K}{N} \right) / \left(1 - \gamma^2 - \frac{1}{N^2} \right) \\ &= \frac{a_K}{2} \left(1 + \frac{a_K}{N} \right) / \left(1 - \frac{1}{N} \right). \end{aligned}$$

Using Chebyshev's inequality, we obtain that for any $\epsilon > 0$

$$\mathbb{P}\left\{ \frac{H_K}{K} \geq (1 + \epsilon) \right\} \leq \frac{v_K}{K^2(\epsilon + 1 - \frac{a_K}{K})^2} = \Theta\left(\frac{1}{\epsilon^2 K} \right).$$

Obviously, $\frac{H_K}{K} \leq 1$, thus $\frac{H_K}{K}$ converges to 1 in probability as N and as a result K tend to infinity. \blacksquare

3.10.4 Proof of Proposition 3.9

Let S be any set of variable nodes of size at most ηK , where we will choose η later. the average degree of variable nodes in S is C . Let $\mathcal{N}_i(S), i \in [C]$ be the check neighbors of \mathcal{G} in hash i . If for at least one of the hashes $i \in [C]$, $|\mathcal{N}_i(S)| > \frac{|S|}{2}$, it results that there is at least one check node of degree 1 (a singleton) among the neighbors, which implies that the peeling decoder can still proceed to decode further variable nodes.

Let \mathcal{E}_s^i denote the event that a specific subset A of size s of variable nodes has at most $\frac{s}{2}$ check neighbors in hash i . Also let $\mathcal{E}_s = \cap_{i=1}^C \mathcal{E}_s^i$. By construction of \mathcal{G} , it is easy to see that $\mathbb{P}\{\mathcal{E}_s\} = \prod_{i=1}^C \mathbb{P}\{\mathcal{E}_s^i\}$. Let T be any subset of check nodes in hash i of size $\frac{s}{2}$. The probability that all the neighbors of A in hash i belong to a specific set T of size $\frac{s}{2}$ is equal to $(\frac{s}{2B})^s$. Taking a union bound over $\binom{B}{s/2}$ of all such sets, it is seen that $\mathbb{P}\{\mathcal{E}_s\} \leq \binom{B}{s/2} (\frac{s}{2B})^s$, which implies that $\mathbb{P}\{\mathcal{E}_s^i\} \leq \left(\binom{B}{s/2} (\frac{s}{2B})^s \right)^C$. Taking a union bound over all possible subsets of size s of variables, we obtain that

$$\begin{aligned} \mathbb{P}\{F_s\} &\leq \binom{K}{s} \mathbb{P}\{\mathcal{E}_s\} \leq \binom{K}{s} \left(\binom{B}{s/2} \left(\frac{s}{2B} \right)^s \right)^C \\ &\leq \left(\frac{eK}{s} \right)^s \left(\frac{2eB}{s} \right)^{sC/2} \left(\frac{s}{2B} \right)^{sC} \leq \frac{u^s s^{s(C/2-1)}}{K^{s(C/2-1)}}, \end{aligned}$$

where $u = e^{C/2+1} (\frac{\beta}{2})^{C/2}$ and where F_s denotes the event that the peeling decoder fail to decode a set of variables of size s . We also used the fact that for $n \geq m$,

$\binom{n}{m} \leq \left(\frac{ne}{m}\right)^m$ and $\mathbb{P}\{F_1\} = \mathbb{P}\{F_2\} = 0$. Selecting $\eta = \frac{1}{2u^{2/(C-2)}}$ and applying the union bound, we obtain that

$$\begin{aligned} \mathbb{P}\{F\} &\leq \sum_{s=1}^{\eta K} \mathbb{P}\{F_s\} = \sum_{s=3}^{\eta K} \mathbb{P}\{F_s\} = \sum_{s=3}^{\eta K} \frac{u^s s^{s(C/2-1)}}{K^{s(C/2-1)}} \\ &= O\left(\frac{1}{K^{3(C/2-1)}}\right) + \sum_{s=4}^{\eta K} \left(\frac{1}{2}\right)^s = O\left(\frac{1}{K^{3(C/2-1)}}\right), \end{aligned}$$

where F is the event that the peeling decoder fails to decode all the variables. This completes the proof. \blacksquare

Rényi Polarization: Hadamard Construction for Signals with Linear Sparsity

4

In this chapter¹, we give a new Hadamard construction for capturing the information of a memoryless sources X with a given probability distribution p_X , recently known as *Analog to Analog* compression. The probability distribution of the source p_X can be continuous, mixture or discrete. For a source with a discrete distribution, this construction immediately gives the absorption phenomenon that we studied in Chapter 2. Although the construction works for general memoryless cases, we are mostly interested to the case where $p_X = (1 - \gamma) \delta_0 + \gamma p_c$ is a mixture distribution where p_c is a continuous distribution, e.g., Gaussian, Laplace, etc., and δ_0 denotes a delta-measure at point 0. In a typical realization of the source for a very large block-length N , approximately $N(1 - \gamma)$ of the components are 0 and the remaining $K = N\gamma$ of the components are generated according to the continuous distribution p_c . We call this the *linear sparsity* regime in contrast to the sublinear case studied in Chapter 3, where the number of nonzero components $K = O(N^\alpha)$ scales sub-linearly with N for some $\alpha \in (0, 1)$. In particular, from the memoryless assumption of the source, it is seen that the support of the resulting sparse signal is uniformly random.

The structure of this chapter is as follows. In Section 4.1, we briefly overview some of the related work, make a connection with our results and introduce the notation for the rest of the chapter. We introduce the *Rényi Information Dimension* as a suitable information measure for our problem in Section 4.2. The main results of the chapter are introduced in Section 4.3. Proof techniques and some further intuition about the problem are given in Section 4.4. In Section 4.5, we discuss two different but closely related aspects of our proposed construction for compressed sensing, i.e., the *informational* aspect versus the *operational* one. In particular, we emphasize the operational implications of the proposed construction, i.e., if we use the constructed matrices to take the measurements and run different recovery algorithms to recover the signal, how the the resulting performance will change in terms of the measurement rate. Simulation results are presented in Section 4.6. Section 4.7 concludes the chapter and discusses some of the benefits of the new Hadamard construction in terms of computational complexity and ease of implementation.

¹This chapter is the result of collaboration with Emmanuel Abbe.

4.1 Related Work

Analog to analog (A2A) compression of signals has recently gathered interest as an information theoretic counterpart of the compressed sensing problem [51–54]. In A2A compression, a high-dimensional analog signal $x^n \in \mathbb{R}^n$ is encoded into a lower-dimensional analog signal $y^m = f_n(x^n) \in \mathbb{R}^m$ where $m \leq n$. The goal is to design the encoder so as to preserve in y^m all the information about x^n , and to successfully reconstruct x^n from y^m for a given distortion measure such as mean square of the error (MSE) or error probability. In particular, the encoding may be corrupted by noise. It is worth mentioning that when the alphabet of x and y is finite, this framework falls into traditional topics of information theory such as lossless and lossy data compression, or joint source-channel coding. The novelty of A2A compression is to consider x and y to be real-valued and to impose regularity constraints on the encoder and the decoder, e.g., the linearity of the encoder as motivated by compressed sensing [1–4].

The challenge and practicality of A2A compression is to obtain dimensionality reduction, i.e., $\frac{m}{n} \ll 1$, by exploiting a prior knowledge on the signal, e.g., sparsity as usually used in signal processing. For k -sparse signals, and without any stability or complexity considerations, it is not hard to see that the dimensionality reduction can be of order $\frac{k}{n}$. A measurement rate of order $\frac{k}{n} \log(\frac{n}{k})$ has been shown to be sufficient to obtain stable recovery by solving tractable optimization algorithms such as convex programming (ℓ_1 -minimization). This remarkable achievement has gathered tremendous amount of attention with a large variety of algorithmic solutions deployed over the past years. The vast majority of the research has, however, capitalized on a common sparsity model.

Several works have explored connections between information theory and compressed sensing², in particular [58–63], however it is only recently in [51] that a foundation of A2A compression has been developed, shifting the attention to probabilistic signal models beyond the sparsity structure. It is shown in [51] that under linear encoding and Lipschitz-continuous decoding, the fundamental limit of A2A compression is the Rényi information dimension (RID), a measure whose operational meaning had remained marginal in information theory until [51]. In the case of a nonsingular mixture distribution, the RID is given by the mass of the continuous part, and for the specific case of sparse mixture distributions, this gives the dimensionality reduction $\frac{k}{n}$. It is natural to ask whether this improvement on compressed sensing is due to potentially complex or non-robust coding strategies. [52] shows that robustness to noise is not a limitation of the framework in [51]. Two other works [53, 54] have corroborated the fact that complexity may not be a limitation either. In [53] spatially-coupled matrices are used for the encoding of the signal, leveraging on the analytical ground of spatially-coupled codes and the predictions of [64]. In particular, [53] shows that the RID is achieved using an approximate message passing algorithm with block diagonal Gaussian measurement matrices. However, the size of the blocks is increasing as the measurement rate approaches the RID.

In this chapter, we give a new approach to A2A compression by means of a polarization theory over the reals. The use of polarization techniques for sparse recovery was proposed in [65] for discrete signals, relying on coding strategies over

²[55–57] investigate LDPC coding techniques for compressed sensing

finite fields. We show that using the RID, one obtains a natural counter-part over the reals of the entropy polarization phenomenon over the finite alphabets [27, 28]. Specifically, we show that the RID of an i.i.d. sequence of mixture random variables polarizes to the two extremal values 0 and 1 (discrete and continuous distributions). To get to this result, we develop properties of the RID in vector setting and other related information measures. We then show that the RID polarization is, as opposed to the entropy polarization, obtained with an analytical pattern. In other words, there is no need to rely on algorithms to compute the set of components that tend to 0 or 1, since this is given by a known pattern equivalent to the BEC channel polarization [27]. This is then used to construct partial Hadamard matrices for A2A compression. Numerical simulations provide evidence that efficient recovery algorithms such as ℓ_1 -minimization or approximate message passing (AMP) can be used in conjunction to the constructed matrices. In particular, using the partial Hadamard matrices constructed in this chapter allows to tremendously speed up these recovery algorithms.

Table 4.1 gives a summary of the notations in this chapter.

Table 4.1 – Summary of Notations

\mathbb{Z}	the set of integers	$[m]$	$\{1, 2, \dots, m\}$
\mathbb{Z}_+	the set of positive integers	X_i^j	$\{X_i, X_{i+1}, \dots, X_j\}$
\mathbb{N}	strictly positive integers	H	discrete entropy
\mathbb{R}	the set of reals	H_N	Hadamard matrix $N = 2^n$
$I(X; Y)$	mutual information of X and Y	$H(X)$	entropy of discrete R.V. X
$I(X; Y z)$	mutual information given $Z = z$	$H(p)$	entropy of p
$I(X; Y Z)$	$\mathbb{E}_Z\{I(X; Y z)\}$	$h(X)$	differential entropy of X
$b(q) \succeq a(q)$	$\lim_{q \rightarrow \infty} \frac{b(q) - a(q)}{\log_2(q)} \geq 0$	$h(p)$	differential entropy of p
$a(q) \preceq b(q)$	$b(q) \succeq a(q)$	$d(X)$	Rényi ID of X
$a(q) \doteq b(q)$	$b(q) \succeq a(q), b(q) \preceq a(q)$	$d(p)$	Rényi ID of p
$[x]_q$	q -ary quantization of $x : \frac{\lfloor qx \rfloor}{q}$		

All probability distributions are assumed to be nonsingular. Hence, in general, for a random variable X , the distribution of X can be decomposed as $p_X = \delta p_c + (1 - \delta)p_d$, where p_c and p_d are the continuous and the discrete part of the distribution and $0 \leq \delta \leq 1$ is the weight of the continuous part. Thus, $\delta = 0$ and $\delta = 1$ correspond to the fully discrete and the fully continuous case respectively. For such a probability distribution, the RID is interchangeably denoted by $d(p_X)$ or $d(X)$ and is equal to the weight of the continuous part δ . Assume that U is a continuous random variable with the probability distribution p_c and V is a discrete random variable with the probability distribution p_d . Let $\Theta \in \{0, 1\}$ be a binary valued random variable, independent of U and V with $\mathbb{P}(\Theta = 1) = \delta$. The random variable X can be written as $X = \Theta U + \bar{\Theta} V$, where $\bar{\Theta} = 1 - \Theta$ and the equality holds in distribution. In this case, the random variable X will have the distribution $p_X = \delta p_c + (1 - \delta)p_d$. Also, if X_1^n is a sequence of such random variables with the corresponding binary random variables Θ_1^n , $C_\Theta = \{i \in [n] : \Theta_i = 1\}$ is a random set consisting of the position of continuous components of the signal. Similarly, $\bar{C}_\Theta = [n] \setminus C_\Theta$ denotes the position of the discrete components.

For a matrix Φ of dimension $m \times n$ and a set $S \subset [n]$, Φ_S is a sub-matrix of dimension $m \times |S|$ consisting of those columns of Φ having index in S . Similarly, for a vector of random variables X_1^n , the vector $X_S = \{X_i : i \in S\}$ is a sub-vector of X_1^n consisting of those random variables having index in S . For two matrices A and B of dimensions $m_1 \times n$ and $m_2 \times n$, $[A; B]$ denotes the $(m_1 + m_2) \times n$ matrix obtained by stacking A on top of B (vertically concatenating A and B).

An ensemble of measurement matrices will be denoted by $\{\Phi_N\}$, where N belongs to a labeling set which is a subset of \mathbb{N} . The dimension of the family for a specific N will be denoted by $m_N \times N$, where m_N is the number of measurements taken by Φ_N . The asymptotic measurement rate of the ensemble is defined by $\rho = \limsup_{N \rightarrow \infty} \frac{m_N}{N}$.

4.2 Rényi information dimension

Let X be a random variable with a probability distribution p_X over \mathbb{R} . The upper and the lower RID of this random variable are defined as follows:

$$\bar{d}(X) = \limsup_{q \rightarrow \infty} \frac{H([X]_q)}{\log_2(q)},$$

$$\underline{d}(X) = \liminf_{q \rightarrow \infty} \frac{H([X]_q)}{\log_2(q)}.$$

By Lebesgue decomposition or Jordan decomposition theorem, any probability distribution over \mathbb{R} such as p_X can be written as a convex combination of a discrete part, a continuous part and a singular part, namely,

$$p_X = \alpha_d p_d + \alpha_c p_c + \alpha_s p_s,$$

where p_d , p_c and p_s denote the discrete, the continuous and the singular part of the distribution and where $\alpha_d, \alpha_c, \alpha_s \geq 0$ and $\alpha_d + \alpha_c + \alpha_s = 1$.

The continuous and the discrete distributions are usually well-known and they are frequently used for modeling the random phenomena in most of the engineering and signal processing applications. However, the singular distributions are rarely encountered in applications. Very briefly, a singular distribution is a probability measure concentrated on a set of Lebesgue measure zero. This is trivially true for a discrete distribution because its support is countable thus has zero Lebesgue measure. However, the difference is that in contrast to a discrete distribution a singular measure assigns zero probability to each individual point. These distributions are sometimes called singular continuous distributions. Such distributions are not absolutely continuous with respect to Lebesgue measure. As a result, they do not have a probability density function with respect to Lebesgue measure since the Lebesgue integral of any such function would be zero. An example is the Cantor distribution which can be described in the following way. Let W_1, W_2, \dots be a sequence of i.i.d. Binomial random variables taking values $\{0, 1\}$ with equal probability. Then the random variable $W = \sum_{i=1}^{\infty} \frac{W_i}{3^i}$ is a well-defined bounded random variable in $[0, \frac{1}{2}]$ with a singular probability distribution over $[0, \frac{1}{2}]$.

In [66], Rényi showed that if $\alpha_s = 0$, i.e., if there is no singular part in the distribution, thus $p_X = (1 - \delta) p_d + \delta p_c$ for some $\delta \in [0, 1]$, then the RID is well-defined and $d(X) = \bar{d}(X) = \underline{d}(X) = \delta$. Moreover, he proved that if X_1^n is a continuous

random vector then $\lim_{q \rightarrow \infty} \frac{H([X_1^n]_q)}{\log_2(q)} = n$, i.e., the RID of any an n -dimensional continuous random vector is n .

Our objective is to extend the definition of RID to arbitrary vector random variables, which are not necessarily continuous. To do so, we first restrict ourselves to a rich space of random variables with a well-defined RID. Over this space, it will be possible to give a full characterization of the RID as we will see in a moment.

Definition 4.1. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a standard probability space and let \mathcal{L}_1 be the set of all nonsingular random variables measurable with respect to \mathcal{F} . The space $\mathcal{L}(\Omega, \mathcal{F}, \mathbb{P})$ is defined as $\mathcal{L} = \cup_{n=1}^{\infty} \mathcal{L}_n$, where for $n \in \mathbb{N} \setminus \{1\}$, \mathcal{L}_n is the space of n -dimensional random vectors defined as

$$\mathcal{L}_n = \{X_1^n : \text{there exist } k \in \mathbb{N}, A \in \mathbb{R}^{n \times k} \text{ and } Z_1^k \text{ independent and nonsingular random variables form } \mathcal{L}_1 \text{ such that } X_1^n = AZ_1^k\}.$$

We call \mathcal{L} the space of all *linearly-correlated* nonsingular random variables. Although it might seem that the linear structure of \mathcal{L} is restrictive for modeling purposes, it is not difficult to see that all n -dimensional vector random variables, singular or nonsingular, can be well approximated in the space \mathcal{L} , e.g., in ℓ_1 or in ℓ_2 sense. However, this is not sufficient to fully characterize the RID. Specially, the RID is discontinuous in ℓ_p topology, $p \geq 1$. For example, we can construct a sequence of discrete random variables in \mathcal{L} converging to a continuous random variable in ℓ_p , whereas the RID of the sequence is 0 and does not converge to 1. Although we have such a mathematical difficulty in characterizing the RID, the space \mathcal{L} is rich enough to model the cases that we encounter in applications.

Over \mathcal{L} , we will generalize the definition of the RID to include joint RID, conditional RID and Rényi information defined as follows.

Definition 4.2. Let X_1^n be a random vector in \mathcal{L} . The joint RID of X_1^n , provided that it exists, is defined as

$$d(X_1^n) = \lim_{q \rightarrow \infty} \frac{H([X_1^n]_q)}{\log_2(q)}.$$

Definition 4.3. Let (X_1^n, Y_1^m) be a random vector in \mathcal{L} . The conditional RID of X_1^n given Y_1^m and the Rényi information of Y_1^m about X_1^n , provided that they exist, are defined as follows:

$$d(X_1^n | Y_1^m) = \lim_{q \rightarrow \infty} \frac{H([X_1^n]_q | Y_1^m)}{\log_2(q)}$$

$$I_R(X_1^n; Y_1^m) = d(X_1^n) - d(X_1^n | Y_1^m).$$

Generally, it is difficult to give a characterization of the RID for a general multi-dimensional distribution because it can contain probability mass over complicated subsets or sub-manifolds of lower-dimension. However, we will show that the vector Rényi information dimension is well-defined for the space \mathcal{L} . In order to give a closed-form formula for the RID over \mathcal{L} , we also need to define some concepts from linear algebra of matrices, namely, for two matrices of appropriate dimensions, we give the following definition of “influence” of one matrix on another matrix and “residual” of one matrix given another matrix.

Definition 4.4. Let A and B be two arbitrary matrices of dimension $m_1 \times n$ and $m_2 \times n$. Also, let $K \subset [n]$. The influence of the matrix B on the matrix A and the residual of the matrix A given B over the column set K are defined as follows:

$$\begin{aligned} I(A; B)[K] &= \text{rank}([A; B]_K) - \text{rank}(A_K), \\ R(A; B)[K] &= \text{rank}([A; B]_K) - \text{rank}(B_K). \end{aligned}$$

Recall that the $[A; B]$ is the matrix obtained by stacking the matrix A on top of the matrix B , thus the number of rows of $[A; B]$ is equal to the sum of the number of rows of A and B . Moreover, for a matrix Φ , we denote by Φ_K a submatrix of Φ consisting of those columns of Φ in the set K . It is easy to check that $I(A; B)[K]$ is the amount of increase of the rank of the matrix A_K by adding the rows of the matrix B_K and $R(A; B)[K]$ is the residual rank of the matrix A_K knowing the rows of the matrix B_K . Moreover, one can easily check that $I(A; B)[K] = R(B; A)[K]$. The next theorem provides closed-form expressions for the joint and conditional RID's. The proof is given in Section 4.4.1.

Theorem 4.1. Let (X_1^n, Y_1^m) be a random vector in the space \mathcal{L} , namely, there are i.i.d. nonsingular random variables Z_1^k and two matrices A and B of dimension $n \times k$ and $m \times k$ such that $X_1^n = A Z_1^k$ and $Y_1^m = B Z_1^k$. Let $Z_i = \Theta_i U_i + \bar{\Theta}_i V_i$ be the representation for Z_i , $i \in [k]$. Then, we have

1. $d(X_1^n) = \mathbb{E}\{\text{rank}(A_{C_\Theta})\}$,
2. $d(X_1^n | Y_1^m) = \mathbb{E}\{R(A; B)[C_\Theta]\}$,

where $C_\Theta = \{i \in [k] : \Theta_i = 1\}$ is the random set consisting of the position of continuous components.

Remark 4.1. Notice that the results intuitively make sense, namely, for a specific realization θ_1^k , if $\theta_i = 0$ we can neglect Z_i because it is fully discrete and does not affect the RID. Moreover, over the continuous components the resulting contribution to the RID is equal to the rank of the matrix A_{C_θ} , which is the effective dimension of the space over which the continuous random variable $A_{C_\theta} U_{C_\theta}$ is distributed. Finally, all these contributions are averaged over all possible realizations of Θ_1^k .

Using Theorem 4.1, we prove a list of properties of the RID in the next theorem. The proof of the theorem is given in Section 4.4.1.

Theorem 4.2. Let (X_1^n, Y_1^m) be a random vector in \mathcal{L} as in Theorem 4.1. Then, we have the following properties:

1. $d(X_1^n) = d(M X_1^n)$ for any arbitrary invertible matrix M of dimension $n \times n$.
2. $d(X_1^n, Y_1^m) = d(X_1^n) + d(Y_1^m | X_1^n)$.
3. $I_R(X_1^n; Y_1^m) = I_R(Y_1^m; X_1^n)$.
4. $I_R(X_1^n; Y_1^m) \geq 0$ and $I_R(X_1^n; Y_1^m) = 0$ if and only if X_1^n and Y_1^m are independent after removing discrete common parts, namely, those $Z_i, i \in [k]$ that are fully discrete.

Further investigation shows that there is a nice duality between the discrete entropy and the RID as depicted in Table 4.2.

Discrete random variables Discrete entropy H Conditional entropy Mutual information Deterministic Chain rule	Random variables in \mathcal{L} RID d Conditional RID Rényi mutual information Discrete Chain rule
Single-terminal source coding Multi-terminal source coding	Single-terminal A2A compression Multi-terminal A2A compression

Table 4.2 – Duality between H and d

4.3 Main results

In this section, we will give a brief overview of the results proved in this chapter. Subsection 4.3.1 is devoted to the results obtained for the polarization of the RID. These results are used in Subsections 4.3.2 to study *A2A compression* problem from an information theoretic point of view.

4.3.1 Polarization of the Rényi information dimension

Before stating the polarization result for the RID, we define the *erasure process*.

Definition 4.5. Let $\alpha \in [0, 1]$. An “erasure process” with initial value α is defined as follows:

1. Let $e_0 = \alpha$ and let $e^+ = e_0^+ = 2\alpha - \alpha^2$ and $e^- = e_0^- = \alpha^2$.
2. Let $e_n = e^{b_1 b_2 \dots b_n}$, for some arbitrary $\{+, -\}$ -valued sequence b_1^n . Define

$$\begin{aligned} e_n^+ &= e^{b_1 b_2 \dots b_n^+} = 2e_n - e_n^2, \\ e_n^- &= e^{b_1 b_2 \dots b_n^-} = e_n^2. \end{aligned}$$

Remark 4.2. Notice that using the $\{+, -\}$ labeling, we can construct a binary tree where each leaf of the tree is labelled by a specific $\{+, -\}$ -valued sequence and is assigned the corresponding erasure value.

Let $\{B_n\}_{n=1}^\infty$ be a sequence of i.i.d. uniform $\{+, -\}$ -valued random variables. Assume that $e_n = e^{B_1 B_2 \dots B_n}$ is the stochastic process induced by the random sequence $\{B_n\}_{n=1}^\infty$ and let \mathcal{F}_n be the σ -field generated by B_1^n . Using the BEC polarization [27, 67], we have the following results:

1. (e_n, \mathcal{F}_n) is a positive bounded martingale.
2. e_n converges to $e_\infty \in \{0, 1\}$ with $\mathbb{P}(e_\infty = 1) = \alpha$.
3. For any $0 < \beta < \frac{1}{2}$, $\liminf_{n \rightarrow \infty} \mathbb{P}(e_n \leq 2^{-N^\beta}) = 1 - \alpha$, where $N = 2^n$ is the number of all possible cases that e_n can take.

Let $n \in \mathbb{N}$ and $N = 2^n$. Assume that X_1^N is a sequence of i.i.d. nonsingular random variables with a RID equal to $d(X)$ and let $Z_1^N = H_N X_1^N$, where H_N is the Hadamard matrix of order N . For $i \in [N]$, let us define $I_n(i) = d(Z_i | Z_1^{i-1})$. Assume that b_1^n is the binary expansion of $i - 1$. By replacing 0 by + and 1 by -, we can equivalently represent $I_n(i)$ be a sequence of $\{+, -\}$ values, namely, $I_n(i) = I^{b_1 b_2 \dots b_n}$. Similar to the erasure process, we can convert I_n to a stochastic process $I_n = I^{B_1 B_2 \dots B_n}$ by using i.i.d. uniform $\{+, -\}$ -valued random variables B_1^n . We have the following theorem whose proof is given in Section 4.4.2.

Theorem 4.3 (Polarization of the RID). *$(I_n, \mathcal{F}_n, \mathbb{P})$ is an erasure stochastic process with initial value $d(X)$ polarizing to $\{0, 1\}$.*

Figure 4.1 shows the polarization of the RID for a random variable with RID 0.5 and for a block-length $N = 512$.

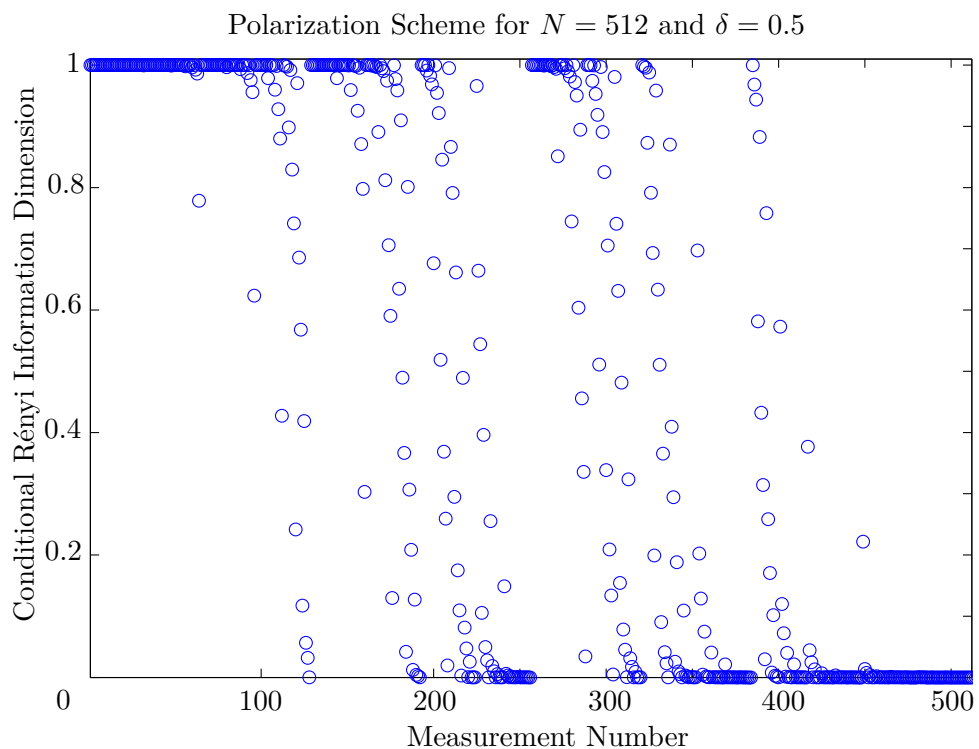


Figure 4.1 – Polarization of the RID for $N = 512$ and $d(X) = 0.5$

4.3.2 A2A compression

In this subsection, we will use the properties of the RID developed in Section 4.2 to study the A2A compression of memoryless sources. We assume that we have a memoryless source with a given nonsingular probability distribution. The idea is to capture the information of the source, to be made clearer in a moment, by taking some linear measurements. As is usual in information theory, we are mostly interested in the asymptotic regime for large block lengths. To do so, we will use an ensemble of measurement matrices to analyze the asymptotic behavior. We will also define the notion of REP (Restricted iso-Entropy Property) for an ensemble of measurement matrices.

Definition 4.6. Let X_1^N be a sequence of i.i.d. random variables with a probability distribution p_X (discrete, mixture or continuous) over \mathbb{R} , and let $D_1^N = [X_1^N]_q$ for $q \in \mathbb{N}$. The family of measurement matrices $\{\Phi_N\}$, indexed with a subsequence of \mathbb{N} and with dimension $m_N \times N$, is ϵ -REP(p_X) with measurement rate ρ if

$$\limsup_{q \rightarrow \infty} \frac{H(D_1^N | \Phi_N X_1^N)}{H(D_1^N)} \leq \epsilon, \quad \limsup_{N \rightarrow \infty} \frac{m_N}{N} = \rho. \quad (4.1)$$

To give some intuitive justification for the REP definition, let us assume that all the measurements are captured with a device with a finite precision $\frac{1}{q_0}$ for some $q_0 \in \mathbb{N}$. In that case, although the potential information content of the signal can be very high, what we effectively observe through the finite precision device is only $H([X_1^N]_{q_0})$. In such a setting, the fraction of the information we lose after taking the measurements is exactly what we have in the definition of REP, namely,

$$\frac{H(D_1^N | \Phi_N X_1^N)}{H(D_1^N)}, \quad (4.2)$$

where we assume that $D_1^N = [X_1^N]_{q_0}$. This might be a reasonable model for application because pretty much this is what happens in reality. The problem with this model is that it is not invariant under some obvious transformations like scaling. For example, assume that we are scaling the signal by some real number. In that case, through some simple examples, it is possible to show that the ratio in Equation (4.2) can change considerably. There are two approaches to cope with this problem. One is to scale the signal with a desired factor to match it to the finite precision quantizer, which in its own can be very interesting to analyze but probably will be too complicated. The other way is to take our approach and develop a theory for the case in which the resolution is high enough so that the quality measure proposed in (4.2) is not affected by the shape of the distribution of the signal.

Remark 4.3. Notice that in the fully discrete case, the REP definition is simplified to the equivalent form that we studied in Chapter 2.

$$\frac{H(X_1^N | \Phi_N X_1^N)}{H(X_1^N)} \leq \epsilon, \quad \limsup_{N \rightarrow \infty} \frac{m_N}{N} \leq \rho.$$

For a non discrete source with strictly positive RID, $d(X) > 0$, if we divide the numerator and the denominator in the expression (4.1) by $\log_2(q)$, take the limit as q tends to infinity and use the definition of the RID, we get the equivalent form

$$\frac{d(X_1^N | \Phi_N X_1^N)}{d(X_1^N)} \leq \epsilon.$$

Interestingly, this implies that in the high resolution regime that we are considering for analysis, as q tends to infinity, the information isometry (keeping more than $1 - \epsilon$ fraction of the information of the signal) is equivalent to the Rényi isometry. Moreover, from the properties of the RID, it is easy to see that this REP measure meets some of the invariance requirements that we expect. For example, it is scale invariant and any invertible linear transformation of the input signal X_1^N keeps the ϵ -REP measure unchanged.

We can also extend the definition when the probability distribution of the source is not known exactly but it is known to belong to a given collection of distributions Π .

Definition 4.7. Assume Π is a class of nonsingular probability distributions over \mathbb{R} . The family of measurement matrices $\{\Phi_N\}$, indexed with a subsequence of \mathbb{N} and with dimension $m_N \times N$, is ϵ -REP(Π) with measurement rate ρ if it is ϵ -REP(π) for every $\pi \in \Pi$.

Now that we have the required tools and definitions, we give a characterization of the required measurement rate in order to preserve the information isometry. As customary in information theory, we do this using the “converse” and the “achievability” parts.

Theorem 4.4 (Converse result). Let X_1^N be a sequence of i.i.d. random variables in \mathcal{L} . Suppose $\{\Phi_N\}$ is a family of ϵ -REP(p_X) measurement matrices of dimension $m_N \times N$ and with measurement rate ρ . Then, $\rho \geq d(X_1)(1 - \epsilon)$.

Proof. The proof is simple considering the fact that information isometry in the definition of ϵ -REP is equivalent to the RID isometry, i.e., for an ϵ -REP ensemble $\{\Phi_N\}$, we have $\frac{d(X_1^N | \Phi_N X_1^N)}{d(X_1^N)} \leq \epsilon$, which implies that

$$I_R(X_1^N; \Phi_N X_1^N) \geq d(X_1^N)(1 - \epsilon).$$

Therefore, we obtain that

$$m_N \geq d(\Phi_N X_1^N) \geq I_R(X_1^N; \Phi_N X_1^N) \geq d(X_1^N)(1 - \epsilon) = Nd(X_1)(1 - \epsilon),$$

which implies that $\frac{m_N}{N} \geq d(X_1)(1 - \epsilon)$. Taking the limit as N tends to infinity, we get the desired result $\rho = \limsup_{N \rightarrow \infty} \frac{m_N}{N} \geq d(X_1)(1 - \epsilon)$. \square

This result implies that to capture the information of the signal, the asymptotic measurement rate must be approximately greater than the RID of the source. This, in some sense, is similar to the single-terminal source coding problem in which the encoding rate must be greater than the entropy of the source, which again emphasizes the analogy between H and d . Moreover, in the discrete case, where $d(X) = 0$, Theorem 4.4 gives the trivial result $\rho \geq 0$.

Remark 4.4. It was proved in [51] that under linear encoding and block error probability distortion, the measurement rate must be higher than the RID of the source, $\rho \geq d(X)$. Theorem 4.4 strengthens this result by stating that $\rho \geq d(X)$ must hold even under the milder ϵ -REP restriction on the measurement ensemble.

Theorem 4.4 puts a lower bound on the measurement rate in order to keep the ϵ -REP property. However, it might happen that there is no measurement family achieving this bound. Fortunately, as we will see, it is possible to deterministically truncate the family of Hadamard matrices to obtain a measurement family with ϵ -REP property and measurement rate $d(X)$. This is summarized in the following two theorems. Notice that in the fully continuous case as Theorem 4.4 implies, the feasible measurement rate, for small ϵ , is approximately 1 which, for example, can be achieved with any complete orthonormal family, thus no explicit construction is necessary. For the non-continuous case, we will only deal with the mixture (non discrete) case.

Theorem 4.5 (Achievability result). *Let X_1^N be a sequence of i.i.d. random variables in \mathcal{L} with $d(X_1) > 0$ and with a common probability distribution p_X over \mathbb{R} . Then, for any $\epsilon > 0$, there is a family of ϵ -REP(p_X) partial Hadamard matrices of dimension $m_N \times N$, for $N = 2^n$ with $\rho = d(X_1)$.*

We also have the general result in Theorem 4.6, which implies that one can construct a family of truncated Hadamard matrices which is ϵ -REP for a class of distributions.

Theorem 4.6 (Achievability result). *Let Π be a family of probability distributions with strictly positive RID. Then, for any $\epsilon > 0$, there is a family of ϵ -REP(Π) partial Hadamard matrices of dimension $m_N \times N$, for $N = 2^n$, with $\rho = \sup_{\pi \in \Pi} d(\pi)$.*

Theorem 4.6 implies that there is a fixed ensemble of measurement matrices capable of capturing the information of all the distributions in the family Π . This is very useful in applications because usually taking the measurements is costly and most of the time, we do not have the exact distribution of the signal. If each distribution needs its own specific measurement matrix, we might need to do several rounds of measurements each time taking the measurements compatible with one specific distribution and do the recovery process for that specific distribution. The benefit of Theorem 4.6 is that one measurement ensemble works for all the distributions. It is also good to notice that although the measurement ensemble is fixed, the recovery (decoding) process might need to know the exact distribution of the signal in order to have successful recovery.

4.4 Proof Techniques

In this section, we will give a brief overview of the techniques used to prove the results. We will divide this section into two subsections. In Subsection 4.4.1, we will overview the proof techniques for the RID. Subsection 4.4.3 will be devoted to proof ideas and intuitions about the A2A compression problem.

4.4.1 Rényi Information Dimension

In this part, we will prove Theorem 4.1 and 4.2 which together give a full characterization of the RID over the space \mathcal{L} .

Proof of Theorem 4.1. To prove the first part of the theorem, notice that

$$H([X_1^n]_q) \doteq H([X_1^n]_q, \Theta_1^k) \doteq H([X_1^n]_q | \Theta_1^k),$$

because $H(\Theta_1^k) \leq k \doteq 0$. As $\Theta_1^k \in \{0, 1\}^k$ and takes finitely many values, it is sufficient to show that for any realization θ_1^k ,

$$\lim_{q \rightarrow \infty} \frac{H([X_1^n]_q | \theta_1^k)}{\log_2(q)} = \text{rank}(A_{C_\theta}). \quad (4.3)$$

Taking the expectation over Θ_1^k , we will get the result. To prove (4.3), notice that

$$H([X_1^n]_q | \theta_1^k) = H([A_{C_\theta} U_{C_\theta} + A_{\bar{C}_\theta} V_{\bar{C}_\theta}]_q) \doteq H([A_{C_\theta} U_{C_\theta} + A_{\bar{C}_\theta} V_{\bar{C}_\theta}]_q | V_{\bar{C}_\theta}) \quad (4.4)$$

$$\doteq H([A_{C_\theta} U_{C_\theta}]_q), \quad (4.5)$$

where we used $H(V_{\bar{C}_\theta}) \leq NH(V_1) \doteq 0$. We also used the fact that knowing $V_{\bar{C}_\theta}$, $[A_{C_\theta}U_{C_\theta}]_q$ and $[A_{C_\theta}U_{C_\theta} + A_{\bar{C}_\theta}V_{\bar{C}_\theta}]_q$ are equal up to a finite uncertainty. Specifically, suppose L is the minimum number of lattices of size $\frac{1}{q}$ required to cover $A_{\bar{C}_\theta} \times [0, \frac{2}{q}]^{|\bar{C}_\theta|}$, which is a finite number. Then

$$H([A_{C_\theta}U_{C_\theta}]_q|V_{\bar{C}_\theta}, [A_{C_\theta}U_{C_\theta} + A_{\bar{C}_\theta}V_{\bar{C}_\theta}]_q) \leq \log_2(L),$$

which implies (4.4) and (4.5).

Generally A_{C_θ} is not full rank. Assume that the rank of A_{C_θ} is equal to m and let A_m be a subset of linearly independent rows. It is not difficult to see that knowing $[A_mU_{C_\theta}]_q$ there is only finite uncertainty in the remaining components of $[A_{C_\theta}U_{C_\theta}]_q$, which is negligible compared with $\log_2(q)$ as q tends to infinity. Therefore, we obtain

$$H([X_1^n]_q|\theta_1^k) \doteq H([A_{C_\theta}U_{C_\theta}]_q) \doteq H([A_mU_{C_\theta}]_q) \doteq m \log_2(q).$$

Thus, taking the limit as q tends to infinity, we obtain

$$\lim_{q \rightarrow \infty} \frac{H([X_1^n]_q|\theta_1^k)}{\log_2(q)} = \text{rank}(A_{C_\theta}).$$

Also, taking the expectation with respect to Θ_1^k , we obtain $d(X_1^n) = \mathbb{E}\{\text{rank}(A_{C_\Theta})\}$, which is the desired result.

To prove the second part of the theorem, notice that $H([X_1^n]_q|Y_1^m) \doteq H([X_1^n]_q|Y_1^m, \Theta_1^k)$. For a specific realization θ_1^k , we have

$$\begin{aligned} H([X_1^n]_q|Y_1^m, \theta_1^k) &= H([A_{C_\theta}U_{C_\theta} + A_{\bar{C}_\theta}V_{\bar{C}_\theta}]_q|B_{C_\theta}U_{C_\theta} + B_{\bar{C}_\theta}V_{\bar{C}_\theta}) \\ &\doteq H([A_{C_\theta}U_{C_\theta} + A_{\bar{C}_\theta}V_{\bar{C}_\theta}]_q|B_{C_\theta}U_{C_\theta} + B_{\bar{C}_\theta}V_{\bar{C}_\theta}, V_{\bar{C}_\theta}) \\ &\doteq H([A_{C_\theta}U_{C_\theta}]_q|B_{C_\theta}U_{C_\theta}). \end{aligned}$$

Generally, A_{C_θ} is not full-rank. Let A_m be the set of all linearly independent rows of A_{C_θ} of size m . Then $H([A_{C_\theta}U_{C_\theta}]_q|B_{C_\theta}U_{C_\theta}) \doteq H([A_mU_{C_\theta}]_q|B_{C_\theta}U_{C_\theta})$.

It may happen that some of the rows of A_m can be written as a linear combination of rows of B_{C_θ} . Let A_r be the remaining matrix after dropping $m - r$ predictable rows of A_m . Given, $B_{C_\theta}U_{C_\theta}$, $A_rU_{C_\theta}$ has a continuous distribution thus

$$H([A_rU_{C_\theta}]_q|B_{C_\theta}U_{C_\theta}) \doteq r \log_2(q).$$

It is easy to check that r is exactly $R(A; B)[C_\theta]$. Therefore, taking the expectation with respect to Θ_1^k , we get $d(X_1^n|Y_1^m) = \mathbb{E}\{R(A; B)[C_\Theta]\}$. \square

Using the results of Theorem 4.1, we can prove Theorem 4.2.

Proof of Theorem 4.2. For part 1, the proof is simple by considering the rank characterization. We know that $X_1^n = AZ_1^k$ and $d(X_1^n) = \mathbb{E}\{\text{rank}(A_{C_\Theta})\}$. Moreover, $MX_1^n = MAZ_1^k$, thus $d(X_1^n) = \mathbb{E}\{\text{rank}(MA_{C_\Theta})\}$. As M is invertible $\text{rank}(A_{C_\Theta}) = \text{rank}(MA_{C_\Theta})$, and we get the result.

For part 2, notice that for any realization θ_1^k and the corresponding set C_θ ,

$$\text{rank}([A; B]_{C_\theta}) = \text{rank}(A_{C_\theta}) + R(B; A)[C_\theta] = \text{rank}(B_{C_\theta}) + R(A; B)[C_\theta].$$

Taking the expectation over Θ_1^k , we get the desired result

$$d(X_1^n, Y_1^m) = d(X_1^n) + d(Y_1^m | X_1^n) = d(Y_1^m) + d(X_1^n | Y_1^m).$$

For part 3, using the chain rule result from part 2 and applying the definition of $I_R(X_1^n; Y_1^m)$, we get

$$I_R(X_1^n; Y_1^m) = d(X_1^n) + d(Y_1^m) - d(X_1^n, Y_1^m),$$

which shows the symmetry of I_R with respect to X_1^n and Y_1^m .

For part 4, notice that for a specific realization θ_1^k , a simple rank check shows that $R(A; B)[C_\theta] \leq \text{rank}(A_{C_\theta})$. Taking the expectation over Θ_1^k , we get $d(X_1^n | Y_1^m) \leq d(X_1^n)$.

If X_1^n and Y_1^m are independent, the equality follows from the definition. For the converse part, notice that if X_1^n is fully discrete then $d(X_1^n | Y_1^m) \leq d(X_1^n) = 0$. Similarly, if Y_1^m is fully discrete then $d(Y_1^m | X_1^n) \leq d(Y_1^m) = 0$ and using the identity $d(X_1^n) - d(X_1^n | Y_1^m) = d(Y_1^m) - d(Y_1^m | X_1^n)$, we get the equality. This case is fine because after removing the discrete $Z_i, i \in [k]$, either X_1^n or Y_1^m is equal to 0, namely, a deterministic value, and the independence holds.

Assume that none of X_1^n or Y_1^m is fully discrete. Without loss of generality, let $Z_1^r, r \leq k$, be the non-discrete random variables among Z_1^k and let \tilde{X}_1^n and \tilde{Y}_1^m be the resulting random vectors after dropping the discrete constituents, namely, we have $\tilde{X}_1^n = A_r Z_1^r$ and $\tilde{Y}_1^m = B_r Z_1^r$, where A_r and B_r are the matrices consisting of the first r columns of A and B respectively. It is easy to check that $d(X_1^n) = d(\tilde{X}_1^n)$ and $d(\tilde{X}_1^n | \tilde{Y}_1^m) = d(X_1^n | Y_1^m)$. Thus, it remains to show that \tilde{X}_1^n and \tilde{Y}_1^m are independent. As we have dropped all the discrete components, the resulting $\Theta_i, i \in [r]$, are 1 with strictly positive probability. This implies that for any realization of θ_1^n and the corresponding C_θ , $R(A_r; B_r)[C_\theta] = \text{rank}(A_{r, C_\theta})$. In particular, this holds for any C_θ of size 1, namely, for any column of A_r and B_r , which implies that if A_r has a non-zero column the corresponding column in B_r must be zero and if B_r has a non-zero column then the corresponding column in A_r must be zero. This implies that \tilde{X}_1^n and \tilde{Y}_1^m depend on disjoint subsets of the random variables Z_1^r . Therefore, they must be independent. \square

4.4.2 Polarization of the RID

In this part, we will prove the polarization of the RID in the single and multi-terminal case as stated in Theorem 4.3. The main idea is to use the recursive structure of the Hadamard matrices and the rank characterization of the RID in the space \mathcal{L} .

Proof of Theorem 4.3. For the initial value, we have $I_0(1) = d(X_1)$. Let $n \in \mathbb{N}$ and $N = 2^n$. To simplify the proof, instead of the Hadamard matrices, H , we will use shuffled Hadamard matrices, \tilde{H} , constructed as follows: $\tilde{H}_1 = H_1$ and \tilde{H}_{2N} is constructed from \tilde{H}_N as follows

$$\begin{pmatrix} \tilde{h}_1 \\ \vdots \\ \tilde{h}_N \end{pmatrix} \rightarrow \begin{pmatrix} \tilde{h}_1 & , & \tilde{h}_1 \\ \tilde{h}_1 & , & -\tilde{h}_1 \\ \vdots & , & \vdots \\ \tilde{h}_i & , & \tilde{h}_i \\ \tilde{h}_i & , & -\tilde{h}_i \\ \vdots & , & \vdots \end{pmatrix},$$

where \tilde{h}_i , $i \in [N]$ denotes the i -th row of the \tilde{H}_N . Let X_1^N be as in Theorem 4.3 and let $\tilde{Z}_1^N = \tilde{H}_N X_1^N$, where H_N has been replaced by \tilde{H}_N . Also, let $\tilde{I}_n(i) = d(\tilde{Z}_i | \tilde{Z}_1^{i-1})$, $i \in [N]$. We first prove that \tilde{I} is also an erasure process with initial value $d(X_1)$ and evolves as follows

$$\begin{aligned}\tilde{I}_n(i)^+ &= \tilde{I}_{n+1}(2i-1) = 2\tilde{I}_n(i) - \tilde{I}_n(i)^2 \\ \tilde{I}_n(i)^- &= \tilde{I}_{n+1}(2i) = \tilde{I}_n(i)^2,\end{aligned}$$

where $i \in [N]$ with the corresponding $\{+, -\}$ -labeling b_1^n . Let \tilde{H}^{i-1} and \tilde{H}^i denote the first $i-1$ and the first i rows of \tilde{H}_N . Also, let \tilde{h}_i denote the i -th row of \tilde{H}_N . We have $\tilde{Z}_1^i = \tilde{H}^i X_1^N$ and $\tilde{Z}_1^{i-1} = \tilde{H}^{i-1} X_1^N$. As X_1^N are i.i.d. nonsingular random variables, it results that \tilde{Z}_1^i belong to the space \mathcal{L} generated by X_1^N . Notice that using the rank characterization for the RID over \mathcal{L} , we have

$$d(\tilde{Z}_i | \tilde{Z}_1^{i-1}) = \mathbb{E}\{I(\tilde{H}^{i-1}; \tilde{h}_i)[C_\Theta]\},$$

where $I(\tilde{H}^{i-1}; \tilde{h}_i)[C_\Theta] \in \{0, 1\}$ is the amount of increase of rank of $\tilde{H}_{C_\Theta}^{i-1}$ by adding \tilde{h}_i . Consider the stage $n+1$, where we have the shuffled Hadamard matrix \tilde{H}_{2N} . Consider the row i^+ , which corresponds to the row $2i-1$ of \tilde{H}_{2N} . If we look at the first block of the new matrix, we simply notice that adding \tilde{h}_i has the same effect in increasing the rank of this block as it had in \tilde{H}_N . A similar argument holds for the second block. Moreover, adding \tilde{h}_i increases the rank of the matrix if it increases the rank of either the first or the second block or both. Let $\mathbf{1}_i(\Theta_1^N) \in \{0, 1\}$ denote the random rank increase in \tilde{H}^{i-1} by adding \tilde{h}_i , then we have

$$\mathbf{1}_{2i-1}(\Theta_1^{2N}) = \mathbf{1}_i(\Theta_1^N) + \mathbf{1}_i(\Theta_{N+1}^{2N}) - \mathbf{1}_i(\Theta_1^N)\mathbf{1}_i(\Theta_{N+1}^{2N}).$$

Θ_1^N and Θ_{N+1}^{2N} are i.i.d. random variables and a simple check shows that $\mathbf{1}_i(\Theta_1^N)$ and $\mathbf{1}_i(\Theta_{N+1}^{2N})$ are also i.i.d. Taking the expectation value, we obtain

$$\tilde{I}_n(i)^+ = \tilde{I}_{n+1}(2i-1) = 2\tilde{I}_n(i) - \tilde{I}_n(i)^2. \quad (4.6)$$

Moreover, if we denote $\tilde{W}_1^N = \tilde{H}_N X_{N+1}^{2N}$, then by the structure of \tilde{H}_N , it is easy to see that $\tilde{I}_n(i)^+$ and $\tilde{I}_n(i)^-$ can be written as follows:

$$\begin{aligned}\tilde{I}_n(i)^+ &= \tilde{I}_{n+1}(2i-1) = d(\tilde{Z}_i + \tilde{W}_i | \tilde{Z}_1^{i-1}, \tilde{W}_1^{i-1}), \\ \tilde{I}_n(i)^- &= \tilde{I}_{n+1}(2i) = d(\tilde{Z}_i - \tilde{W}_i | \tilde{Z}_i + \tilde{W}_i, \tilde{Z}_1^{i-1}, \tilde{W}_1^{i-1}).\end{aligned}$$

Using the chain rule for the RID, we have

$$\begin{aligned}\frac{\tilde{I}_n(i)^+ + \tilde{I}_n(i)^-}{2} &= \frac{1}{2}d(\tilde{Z}_i - \tilde{W}_i, \tilde{Z}_i + \tilde{W}_i | \tilde{Z}_1^{i-1}, \tilde{W}_1^{i-1}) \\ &= \frac{1}{2}d(\tilde{Z}_i, \tilde{W}_i | \tilde{Z}_1^{i-1}, \tilde{W}_1^{i-1}) = d(\tilde{Z}_i, | \tilde{Z}_1^{i-1}) = \tilde{I}_n(i),\end{aligned}$$

which along with (4.6), implies that $\tilde{I}_n(i)^- = \tilde{I}_n(i)^2$. Therefore, \tilde{I} evolves like an erasure process with the initial value $d(X)$.

Notice that the only difference between H_N and \tilde{H}_N is the permutation of the rows, namely, there is a row shuffling matrix B_N such that $\tilde{H}_N = B_N H_N$. It was proved in [28] that B_N and H_N commute, which implies that $\tilde{H}_N X_1^N = H_N B_N X_1^N$.

However, notice that X_1^N is an i.i.d. sequence and $B_N X_1^N$ is again an i.i.d. sequence with the same distribution as X_1^N . In particular, adding or removing B_N does not change the RID values, which implies that for $Z_1^N = H_N X_1^N$ and $I_n(i) = d(Z_i | Z_1^{i-1})$, $I_n(i) = \tilde{I}_n(i)$. Therefore, I is also an erasure process with initial value $d(X)$, which polarizes to $\{0, 1\}$. \square

4.4.3 A2A Compression

In this part, we prove the achievability part for the single-terminal case.

Proof of Theorem 4.5. We will give an explicit construction of the measurement ensemble. Let $n \in \mathbb{N}$ and $N = 2^n$. Assume that X_1^N is a sequence of i.i.d. nonsingular random variables with RID equal to $d(X)$. Let $Z_1^N = H_N X_1^N$, where H_N is the Hadamard matrix of order N . Also, assume that $I_n(i) = d(Z_i | Z_1^{i-1})$, $i \in [N]$. As we proved in Theorem 4.3, I is an erasure process with initial value $d(X)$. We will construct the measurement matrix Φ_N by selecting all the rows of H_N with the corresponding I_n value greater than or equal to $\epsilon d(X)$. Therefore, we can construct the measurement ensemble $\{\Phi_N\}$ labelled with all N that are a power of 2. Assume that the dimension of Φ_N is $m_N \times N$. It remains to prove that the ensemble $\{\Phi_N\}$ is ϵ -REP with measurement rate $d(X)$. This will complete the proof of Theorem 4.5.

We first show that the family $\{\Phi_N\}$ has measurement rate $d(X)$. Notice that the process I_n converges almost surely. Thus, it also converges in probability. Specifically, considering the uniform probability assumption, we have

$$\begin{aligned} \limsup_{N \rightarrow \infty} \frac{m_N}{N} &= \limsup_{N \rightarrow \infty} \frac{\#\{i \in [N] : I_n(i) \geq \epsilon d(X)\}}{N} \\ &= \limsup_{n \rightarrow \infty} \mathbb{P}(I_n \geq \epsilon d(X)) = \mathbb{P}(I_\infty \geq \epsilon d(X)) = d(X). \end{aligned}$$

It remains to prove that $\{\Phi_N\}$ is ϵ -REP. Let $S = \{i \in [N] : I_n(i) \geq \epsilon d(X)\}$ denote the selected rows to construct Φ_N and let $Z_1^N = H_N X_1^N$ be the full measurements. It is easy to check that $\Phi_N X_1^N = Z_S$. Also, let $B_i = S^c \cap [i-1]$ denote all the indices in S^c before i . We have

$$\begin{aligned} d(X_1^N | Z_S) &= d(Z_1^N | Z_S) = d(Z_{S^c} | Z_S) = \sum_{i \in S^c} d(Z_i | Z_{B_i}, Z_S) \\ &\leq \sum_{i \in S^c} d(Z_i | Z_1^{i-1}) = \sum_{i \in S^c} I_n(i) \leq N \epsilon d(X) = \epsilon d(X_1^N), \end{aligned}$$

which shows the ϵ -REP property for $\{\Phi_N\}$. \square

The last step is to prove Theorem 4.6, i.e., to show that for a family of mixture distributions Π with strictly positive RID, there is a fixed measurement family $\{\Phi_N\}$, which is ϵ -REP for all the distributions in Π with a measurement rate vector lying in the Rényi information region of the family.

Proof of Theorem 4.6. The proof is simple considering the fact that the construction of the family $\{\Phi_N\}$ in the proof of Theorem 4.5 depends only on the erasure pattern. Also, the erasure pattern is independent of the shape of the distribution and only depends on its RID. Moreover, it can be shown that the erasure patterns for

different values of δ are ordered, namely, for $\delta > \delta'$, $I_n^\delta(i) \geq I_n^{\delta'}(i)$, $i \in [N]$. Considering the method we use to construct the family $\{\Phi_N\}$, this implies that an ϵ -REP measurement family designed for a specific RID δ is ϵ -REP for any distribution with RID less than δ . Thus, if we design $\{\Phi_N\}$ for $\sup_{\pi \in \Pi} d(\pi)$, it will be ϵ -REP for any distribution in the family. \square

4.5 Operational vs. Informational Characterization

Up to now, we defined the notion of ϵ -REP for an ensemble of measurement matrices. This definition is what we call an “informational” characterization, in the sense that taking measurements by the ensemble potentially keeps more than $1 - \epsilon$ fraction of the information of the source. One can ask the natural question whether this has some “operational” implication, in the sense that after having the linear measurements, is it possible to recover the source up to an acceptable distortion (for example mean square error distortion)? Notice that, at the end of the day, it is the operational implication that matters because it practically deals with taking the measurements and reconstructing the signal via a recovery algorithm. More importantly, it takes into account the computational complexity of the recovery algorithm which is completely missing in the informational point of view. However, the informational point of view has its own advantages, namely, it allows to find the underlying fundamental limits of the problem without dealing with algorithmic issues. In this section, our goal is to briefly clarify these two different aspects for our proposed Hadamard construction.

Let us start from an example from polar codes for binary source compression which has lots of similarities with what we have studied in this section. As shown in [28], for a binary memoryless source with $\mathbb{P}(0) = p$, for a large block length n , there is a matrix G_n of an approximate dimension $nh_2(p) \times n$ such that the linear measurement of the source by this matrix over \mathbb{F}_2 faithfully captures the randomness of the source. This in its own only solves the encoding part of problem without directly addressing the decoding part, namely, it does not imply the existence of a decoder to recover the source from the measurements up to a negligible distortion (in this case error probability). Therefore, the operational picture is not complete yet. Fortunately, in the case of polar codes the successive cancellation decoder (or other decoders proposed in the literature) fills up the gap and shows that the *informational* characterization implies the *operational* one.

In the compressed sensing setting, the operational fundamental limits have been studied extensively for different measurement matrices and different recovery algorithms. In particular, for random Gaussian measurement matrices, the asymptotic sparsity-measurement rate behavior of different recovery algorithms, such as AMP and ℓ_1 -minimization, has been vastly studied. We will mainly focus on the results for the Gaussian matrices because, generally speaking, they give better measurement rates than other families of matrices. For the probabilistic model that we studied in this chapter, it has been observed that the required measurement rate of the Gaussian matrices is far from optimal. More precisely, for the sparsity δ very close to zero, the required measurement rate for successful recovery of the source under low-complexity algorithms such as ℓ_1 -minimization and AMP scales like $\delta \log_2(\frac{1}{\delta})$, which suffers from an oversampling of order $\log_2(\frac{1}{\delta})$ compared with the optimal measurement rate δ predicted information theoretically. In [53], it was shown that one can compensate

this extra factor by using spatially coupled Gaussian matrices and running AMP. This specifically shows that the operational limits meet the informational predictions.

Recently, using extensive numerical simulations, it was shown that the optimal measurement rate δ still seems to be achievable by spatially coupling of partial Hadamard matrices, where the rows of the sub-matrices embedded in each block are selected completely at random from the rows of a Hadamard matrix [68]. Interestingly, it was observed that for this random construction, the resulting measurement rate of the Lasso and AMP recovery algorithms are comparable with (even slightly better than) that of random Gaussian matrices. It will be very interesting to find out if it is possible to derandomize this construction to obtain a deterministic partial Hadamard family of matrices with a close to optimal performance.

In the rest of this chapter, our goal is to operationally compare the performance of our proposed partial Hadamard matrices with that of the random Gaussian matrices. We will restrict ourselves to the dense Gaussian matrices without using the spatial coupling. To build our proposed measurement matrices, we use the REP criterion that we developed in this chapter. More precisely, we select those rows of the corresponding Hadamard matrix with a significant RID as we will explain further in Section 4.6. Since it is difficult to theoretically analyze the performance of the constructed matrices, we use numerical simulations to assess the performance for different signal models and different off-the-shelf algorithms from compressed sensing. Very briefly, the simulation results show that for our construction, the resulting measurement rate is comparable with (and even slightly better than) the random Gaussian matrices but it still suffers from the oversampling factor $\log_2(\frac{1}{\delta})$ for small sparsity values δ . This shows that even in our construction there is a gap between the informational and operational characterization and it seems that an extra spatial coupling as in [68] is still necessary to meet the optimal informational predictions.

4.6 Simulation Results

In this section, we assess the operational performance of the partial Hadamard matrices proposed in this chapter via numerical simulations. As explained in Section 4.5, this allows to numerically compare the gap that exists between the informational and operational characterizations.

4.6.1 Signal Model and the Recovery Algorithm

For simulations, we use a zero mean and unit variance sparse distribution

$$p_X(x) = (1 - \delta)\delta_0(x) + \delta p_c(x), \quad (4.7)$$

where $\delta_0(x)$ is the unit delta measure at point zero, p_c is the distribution of the continuous part and $\delta \in \{0.0, 0.1, \dots, 0.9, 1.0\}$ is the RID of the signal. We use the mean square error (MSE) as distortion measure. The simulations are done with the Hadamard matrix of order $N = 512$. To build the measurement matrix A , we select those rows of H_N with highest conditional RID, as stated in Section 4.4.3, until we get an acceptable recovery distortion. One of the algorithms that we use to recover the signal is the ℓ_1 -minimization algorithm:

$$\hat{x}(y) = \arg \min_{w \in \mathbb{R}^N} \|w\|_1 \quad \text{subject to } y = Aw, \quad (4.8)$$

where $y = Ax$ is the vector of measurements taken from the signal x . Another algorithm that we use is the AMP algorithm given by the following iteration:

$$z_t = y - A\hat{x}_t + \frac{1}{\gamma}z_{t-1}\langle\eta'_{t-1}(A^*z_{t-1} + \hat{x}_{t-1})\rangle, \quad (4.9)$$

$$\hat{x}_{t+1} = \eta_t(A^*z_t + \hat{x}_t), \quad (4.10)$$

where $y = Ax$ denotes the vector of linear measurements taken by A , γ is the measurement rate, $\langle a_1^n \rangle = \sum_{i=1}^n a_i/n$, $\eta_t(u) = (\eta_{t,1}(u_1), \dots, \eta_{t,N}(u_N))$, where $\eta_{t,i}(u_i) = \mathbb{E}\{X|u_i = X + \tau_t W\}$, with $W \sim \mathcal{N}(0, 1)$ independent of the signal X and τ_t given by the state evolution equation for AMP, is the soft-thresholding function designed for the known distribution of X . For initialization, we use $\hat{x}_0 = 0$ and $z_0 = 0$. The behavior of the AMP algorithm was rigorously analyzed for random Gaussian matrices in [21]. Specifically, it was shown that the behavior of the AMP is fully characterized by a closed-form *State Evolution* equation.

4.6.2 Sensitivity to Signal Distribution

In this chapter, we proposed the Hadamard construction for a memoryless source with a given probability distribution. However, we showed that the polarization of the RID and as a result the matrix construction only depends on the RID of the source and not the detail of the distribution of the source. To assess how sensitive the construction is to the distribution of the signal, we do the simulations using three different distribution for the continuous part of signal distribution p_c in Equation (4.7). To recover the signal, we use ℓ_1 -minimization algorithm as in Equation (4.8). Figure 4.2 shows the boundary of the low-distortion region for the ℓ_1 -minimization algorithm where for the boundary we use 0.01 of the signal power as the threshold. The results show that the required measurement rate is not sensitive to the distribution of the signal.

4.6.3 Comparison of the Performance of ℓ_1 -minimization and AMP

In this part, we compare the performance of the two algorithms for a Bernoulli-Gaussian distribution in which p_c is the normal distribution. Knowing the exact distribution of the signal, we use MMSE soft-thresholding function for AMP as in Equation 4.10. Figure 4.3 shows the simulation results. Although AMP, with the thresholding function η_t designed for the known distribution of the signal, performs slightly better than ℓ_1 -minimization, there is still a gap compared with the optimal line.

4.6.4 Comparison with Random Gaussian Matrices

In this section, we compare the performance of the Hadamard construction with the traditional random Gaussian matrices extensively used in compressed sensing.

The simulation results for ℓ_1 -minimization are depicted in Figures 4.4 and 4.5. A visual comparison shows that the Hadamard construction works slightly better than the Gaussian matrices, i.e., for a given measurement rate has less recovery distortion.

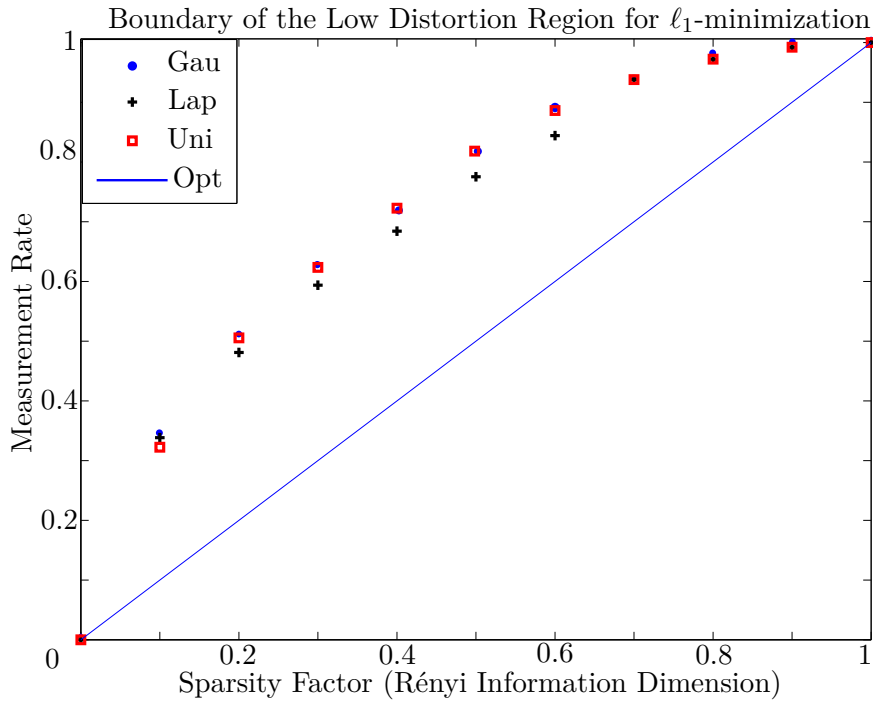


Figure 4.2 – Boundary of the Low-Distortion Region for ℓ_1 -minimization for Different Signal Distributions. It is seen that the performance of the ℓ_1 -minimization is not sensitive to the distribution of the signal.

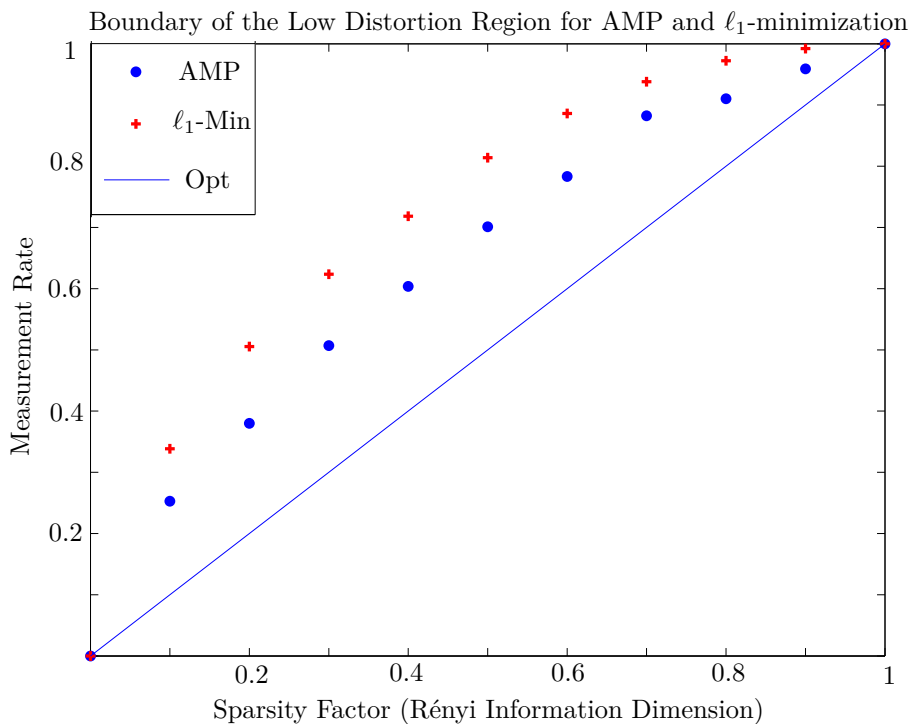


Figure 4.3 – Boundary of the Low-Distortion Region for AMP and ℓ_1 -minimization. The solid line shows the optimal boundary. Below this line no algorithm can work with a low distortion. As seen from the figure, AMP performs better than ℓ_1 -minimization, requiring lower measurement rate for low-distortion recovery.

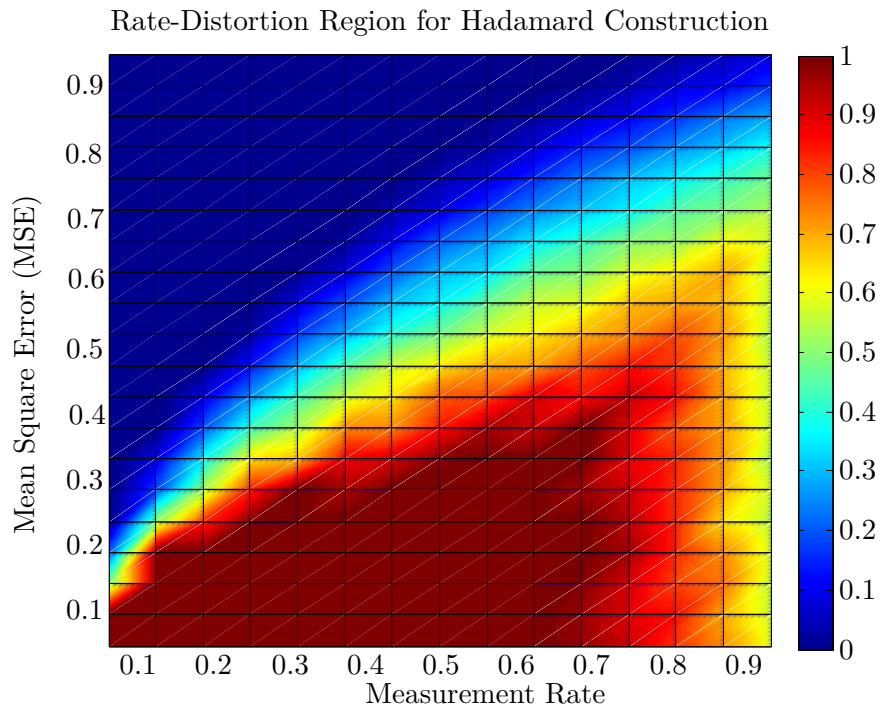


Figure 4.4 – Rate-Distortion Region for Hadamard Construction and ℓ_1 -minimization

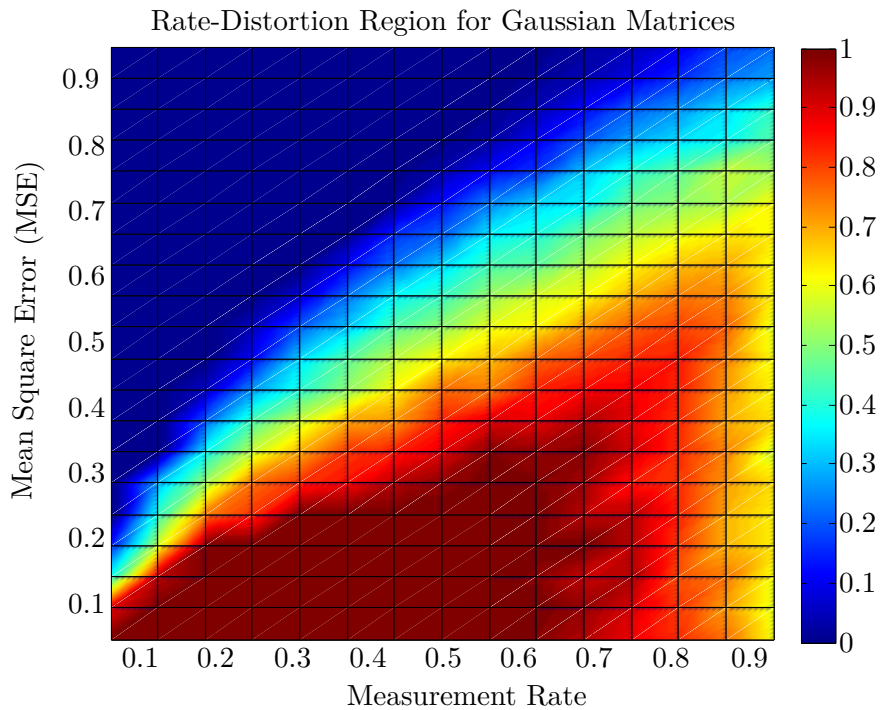


Figure 4.5 – Rate-Distortion Region for Random Gaussian Matrices and ℓ_1 -minimization

4.7 Conclusion and Further Discussion

In this chapter, using ideas from information theory and polarization theory, we proposed a new Hadamard construction which was able to capture the information of a memoryless signal with information theoretically optimal measurement rate. We also simulated the resulting measurement matrices for different signal models and recovery algorithm. In particular, a comparison with the performance of traditional Gaussian matrices revealed that the constructed Hadamard matrices perform equally or sometimes slightly better. As a conclusion, it is worth to mention some of the benefits of the new construction over the usual Gaussian matrices:

- **Implementation Gain:** as the components of the constructed matrices are $\{+1, -1\}$, generally speaking, it is very easy to implement them on sensors or measurement devices. In a practical scenario, even after careful adjustment and calibration of the measurement devices, there is still a mismatch with the intended measurement. Using the two valued matrices as in our proposed construction, one can hope to reduce the mismatch significantly.
- **Storage Gain:** there is no need to store the measurement matrices in the software. One only needs to store the indices of the rows selected from the Hadamard matrix and the whole matrix can be simply generated by a closed-form formula. This is in particular important for the recovery algorithm because one does not need to save the matrix in the computer, thus the algorithm can run for very high-dimensional signals.
- **Computational Gain:** after taking the measurements $y = Ax$ from the signal x via the measurement matrix A , it is necessary to run the recovery algorithm to recover the resulting signal. In more or less all the recovery algorithms, one has the *matched-filter* phase of computation where it is necessary to compute A^*y , i.e., the correlation of the columns of the matrix with the measurement y . Applying the new construction and using the recursive structure of the Hadamard matrices, it is possible to reduce the complexity of this phase of computation by approximately $O(\frac{N}{\log_2(N)})$ compared with the traditional matrix multiplication, where N denotes the dimension of the signal. Even for a typical value $N = 10^3$, this is around 100 times faster. Nowadays, one of the main difficulties in using compressed sensing in real-world applications is the dimensionality scaling problem in the sense that for high-dimensional signals, the computational complexity of most of the algorithms prohibits their application for recovery. By replacing the unstructured measurement matrices by structured ones, e.g., partial Hadamard matrices constructed in this chapter, one can get a huge gain in computational complexity and much better dimensionality scaling without any loss in performance, e.g., measurement rate or recovery distortion.

Multi-Terminal Compressed Sensing

5

In this chapter, we follow two purposes. First, we show that the information theoretic setting and the Hadamard construction that we developed in Chapter 4, which we call single-terminal construction, can be extended in a natural way to a distributed or multi-terminal scenario in which one observes a distributed signal via multiple observation points (terminals). This is a typical scenario in a sensor network, where different sensors observe, for example, the temperature or humidity in multiple points in the environment and the goal is to recover the distributed signal by taking sufficiently many measurements from different terminals. We show that in the information theoretic limit, in order to be able to recover the signal, the required measurement rate from each terminal can be fully characterized in terms of joint and conditional Rényi information dimension of the multi-terminal signal.

Second, as it is difficult to analyze and fully characterize the required measurement region for the Hadamard construction, we study the Gaussian case where the measurement matrices are random Gaussian independent across different terminals. In order to analytically study the problem, we extend the ‘*Approximate Message Passing*’ (AMP) algorithm, initially developed for compressed sensing of signals under i.i.d. Gaussian matrices, to a multi-terminal setting (MAMP algorithm). In particular, we show that similar to its single-terminal counterpart, the behavior of MAMP algorithm is fully characterized by a ‘*State Evolution*’ (SE) equation for large block-lengths. We use this equation to obtain the rate-distortion curve of a multi-terminal memoryless source. We also extend the analysis to the spatially coupled measurement matrices and show that the measurement rate region corresponding to a low distortion (approximately zero distortion) regime is fully characterized by the joint and conditional Rényi information dimension (RID) of the multi-terminal source essentially proving that the information theoretic characterization and the operational characterization meet each other. Simulations have been done to investigate the empirical behavior of MAMP algorithm. It is observed that simulation results match very well with predictions of SE equation for reasonably large block-lengths.

The structure of this chapter is as follows. In Section 5.1, we introduce the problem and review some of the related works. We extend the Hadamard construction for multi-terminal settings in Section 5.2. In Section 5.3, we develop a message passing

algorithm for Gaussian measurement matrices and approximate it to get the MAMP algorithm and study the asymptotic behavior of this algorithm as predicted by a state evolution equation. Section 5.4 extends the results to spatially coupled Gaussian matrices. Finally, in Section 5.5, we simulate the MAMP and compare the simulation results with the theoretical results developed in this chapter.

5.1 Introduction and Related Work

Let (x^n, y^n) be a realization of a two-terminal memoryless sources (X, Y) with a probability distribution $p_{X,Y}$ over \mathbb{R}^2 and assume that one is interested in recovering the signal in both terminals (x^n, y^n) by taking sufficiently many linear measurements $\mathbf{u} = Ax^n$ and $\mathbf{v} = By^n$, where A and B denote the measurement matrices in T_X (terminal X) and T_Y (terminal Y) respectively. In particular, it is implicitly assumed that the measurements are taken separately from each terminal whereas for the recovery, one has access to the measurements (\mathbf{u}, \mathbf{v}) from both terminals. One can compare this setting with a more general case where (x^n, y^n) is considered as a $2n$ dimensional signal and each measurement is a linear mixture of components of both x^n and y^n . The difference is that in the former each terminal takes measurements from its own signal (*separate measurements*) but in the latter some coordination between the terminals is necessary for taking the measurements (*joint measurements*). Because of this separate measurement scheme, one can reasonably define the measurement rate in a specific terminal as the number of measurements taken solely from that terminal normalized by the signal dimension. Moreover, in a general multi-terminal setup, we define the measurement rate from a subset of terminals to be the sum of individual measurement rates in those terminals.

This problem in its general multi-terminal form arises in many distributed processing systems. For example, in an ad hoc sensor network, a collection of sensors measure a distributed environmental signal such as temperature, humidity, etc. One can imagine a particular sensor as a terminal that takes a collection of linear measurements and transmits the gathered data to a data fusion center by routing via the other sensors. Because of limited communication and low processing power of sensors, it is difficult to take joint measurements from two or several different terminals even if they are very close to one another. Therefore, one can reasonably assume that the measurements are taken separately from each terminal and processed jointly in a data fusion center to recover the distributed signal. Usually there is a high correlation among the terminals and one can exploit it to reduce the required number of measurements. In particular, in a very low energy scenario such as a sensor network this results in a saving in the energy consumption of devices, which in turn, increases the life time of the network.

There are two different kinds of correlations that should be considered: temporal and spatial. In a sensor network scenario, temporal correlations result because of the slow changes of the natural phenomenon such as temperature, humidity, etc. Temporal correlations usually can be moderated by suitable sampling time and preprocessing of the signal before transmission. Spatial correlations are more important and much more difficult to deal with. If the sensors are densely distributed in the environment for precise data acquisition, the resulting measurements from different terminals will be highly redundant thus the network energy resources are

wasted without any significant gain. Therefore, it is always desirable to reduce the number of sensors to a minimum possible and still be able to recover the environmental distributed signal. Compared with a densely distributed sensor network, this is as if no sensor is assigned to some of the terminals and as a result the measurement rate from those terminals is 0. In both cases, one needs to characterize the required measurement rate region of the terminals for recovery with an acceptable distortion.

This problem has been vastly studied under different signal structures and recovery algorithms (in particular [22, 23]) as an extension of the traditional single-terminal compressed sensing introduced in [1, 2]. Specially, it has been attempted to make a connection between the multi-terminal compressed sensing and the distributed source coding (Slepian-Wolf) counterpart in information theory; refer to [24] for extra references. The main difficulty is that, unlike the single-terminal case where the sparsity model works very well and provides fruitful results and intuitions from a signal processing point of view, in the multi-terminal case it is difficult to find a comprehensive model that on one hand models the sparsity of the signal in each terminal and on the other hand takes into account the spatial correlation existing among the different terminals.

In this chapter, for simplicity, we address a two-terminal scenario for a memoryless distributed source (X^n, Y^n) . The memoryless property of the source implies that there is no temporal correlation between samples of the signals along the time. The spatial correlation is modeled by assuming that the samples of the signals (X_i, Y_i) are generated by a probability distribution $p_{X,Y}$. The extension to more than two terminals is also straightforward.

We essentially follow the same notation as Chapter 4. In particular, we will use $d(X) = d(p_X)$ to denote the RID of a random variable with probability distribution p_X . A closely related parameter to the RID is the MMSE dimension of X defined in [69]. Let

$$\text{mmse}(s) = \mathbb{E}(X - \mathbb{E}(X|Y))^2, \quad Y = \sqrt{s}X + Z,$$

where $Z \sim \mathcal{N}(0, 1)$ is a Gaussian random variable independent of X . The upper and lower MMSE dimension of X are defined by

$$\begin{aligned} \overline{D}(p_X) &= \overline{D}(X) = \limsup_{s \rightarrow \infty} s \text{mmse}(s) \\ \underline{D}(p_X) &= \underline{D}(X) = \liminf_{s \rightarrow \infty} s \text{mmse}(s), \end{aligned}$$

and if both limits coincide then we define $D(X) = \overline{D}(X) = \underline{D}(X)$. In [69], it was proved that if $H(\lfloor X \rfloor) < \infty$ then

$$\underline{D}(X) \leq d(X) \leq \overline{d}(X) \leq \overline{D}(X).$$

Hence, if $D(X)$ exists so does $d(X)$ and they are equal. For simplicity, we will restrict ourselves to the space of linearly correlated random variables \mathcal{L} introduced in Chapter 4, where a k dimensional random vector S is linearly correlated if there is a sequence of independent non-singular variables Z^n and a $k \times n$ matrix A such that $S = AZ^n$. This space is rich enough for most of the applications. Furthermore, over this space it is possible to give a full characterization of the joint and the conditional RID as explained in Chapter 4.

5.2 Hadamard Construction for Multi-terminal A2A Compression

In this section, our goal is to extend the A2A compression theory from the single-terminal case to the multi-terminal one. In the multi-terminal setting, we have a memoryless source distributed in more than one terminal and we are going to take linear measurements from different terminals in order to capture the information of the source. We are again interested in an asymptotic regime for large block lengths. We will use an ensemble of distributed measurement matrices that we will introduce in a moment. Similar to the single-terminal case, we are interested in the measurement rate region of the problem, namely, the number of measurements that we need from different terminals in order to capture the signal faithfully. For simplicity, we will analyze the problem for the two-terminal case.

Definition 5.1. Let $\{(X_i, Y_i)\}_{i=1}^N$ be a two-terminal memoryless source with (X_1, Y_1) being in \mathcal{L} and having a distribution $p_{X,Y}$ over \mathbb{R}^2 . The family of distributed measurement matrices $\{\Phi_N^x, \Phi_N^y\}$, indexed with a subsequence of \mathbb{N} , is ϵ -REP($p_{X,Y}$) with measurement rate (ρ_x, ρ_y) if

$$\begin{aligned} \limsup_{q \rightarrow \infty} \frac{H([X_1^N]_q, [Y_1^N]_q | \Phi_N^x X_1^N, \Phi_N^y Y_1^N)}{H([X_1^N]_q, [Y_1^N]_q)} &\leq \epsilon, \\ \limsup_{N \rightarrow \infty} \frac{m_N^x}{N} = \rho_x, \quad \limsup_{N \rightarrow \infty} \frac{m_N^y}{N} = \rho_y. \end{aligned} \quad (5.1)$$

Remark 5.1. If (X, Y) is a random vector in \mathcal{L} with $d(X, Y) > 0$, similar to what did in the single-terminal case, dividing the numerator and the denominator in the expression (5.1) by $\log_2(q)$ and taking the limit as q tends to infinity, we get the equivalent definition

$$\frac{d(X_1^N, Y_1^N | \Phi_N^x X_1^N, \Phi_N^y Y_1^N)}{d(X_1^N, Y_1^N)} \leq \epsilon,$$

which implies the equivalence of the information isometry and the Rényi isometry.

Remark 5.2. Notice that in the fully discrete case, the definition above is simplified to the equivalent form

$$\begin{aligned} \frac{H(X_1^N, Y_1^N | \Phi_N^x X_1^N, \Phi_N^y Y_1^N)}{H(X_1^N, Y_1^N)} &\leq \epsilon, \\ \limsup_{N \rightarrow \infty} \frac{m_N^x}{N} = \rho_x, \quad \limsup_{N \rightarrow \infty} \frac{m_N^y}{N} = \rho_y. \end{aligned}$$

Similar to the single-terminal case, we are interested in the rate region of the problem. We have the following converse and achievability results.

Theorem 5.1 (Converse result). Let $\{(X_i, Y_i)\}_{i=1}^N$ be a two-terminal memoryless source with (X_1, Y_1) being in \mathcal{L} . Assume that the distributed family of measurement matrices $\{\Phi_N^x, \Phi_N^y\}$ is ϵ -REP with a measurement rate (ρ_x, ρ_y) . Then,

$$\begin{aligned} \rho_x + \rho_y &\geq d(X_1, Y_1)(1 - \epsilon), \\ \rho_x &\geq d(X_1 | Y_1) - \epsilon d(X_1, Y_1), \quad \rho_y \geq d(Y_1 | X_1) - \epsilon d(X_1, Y_1). \end{aligned}$$

Proof. In the two-terminal case, ϵ -REP implies that $\frac{d(X_1^N, Y_1^N | \Phi_N^x X_1^N, \Phi_N^y Y_1^N)}{d(X_1^N, Y_1^N)} \leq \epsilon$, which is equivalent to $I_R(X_1^N, Y_1^N; \Phi_N^x X_1^N, \Phi_N^y Y_1^N) \geq d(X_1^N, Y_1^N)(1 - \epsilon)$. Therefore, we get

$$m_N^x + m_N^y \geq d(\Phi_N^x X_1^N, \Phi_N^y Y_1^N) \geq I_R(X_1^N, Y_1^N; \Phi_N^x X_1^N, \Phi_N^y Y_1^N) \geq d(X_1^N, Y_1^N)(1 - \epsilon),$$

which implies that $\frac{m_N^x}{N} + \frac{m_N^y}{N} \geq d(X_1, Y_1)(1 - \epsilon)$. Taking the limit as N tends to infinity, we get $\rho_x + \rho_y \geq d(X_1, Y_1)(1 - \epsilon)$. Similarly, we have

$$\begin{aligned} m_N^x &\geq d(\Phi_N^x X_1^N) \geq d(\Phi_N^x X_1^N | \Phi_N^y Y_1^N) = d(\Phi_N^x X_1^N, \Phi_N^y Y_1^N) - d(\Phi_N^y Y_1^N) \\ &\geq I_R(X_1^N, Y_1^N; \Phi_N^x X_1^N, \Phi_N^y Y_1^N) - d(Y_1^N) \geq d(X_1^N, Y_1^N)(1 - \epsilon) - d(Y_1^N), \end{aligned}$$

where in the last inequality we used the ϵ -REP property. This implies that $\frac{m_N^x}{N} \geq d(X_1|Y_1) - \epsilon d(X_1, Y_1)$. Taking the limit as N tends to infinity, we get the result $\rho_x \geq d(X_1|Y_1) - \epsilon d(X_1, Y_1)$. The other inequality follows by symmetry. \square

Remark 5.3. *This rate region is very similar to the rate region of the distributed source coding (Slepian & Wolf) problem with the only difference that the discrete entropy has been replaced by the RID, which again emphasizes the analogy between the discrete entropy and the RID. Similar to the Slepian & Wolf problem, we call $\rho_x + \rho_y = d(X_1, Y_1)$ the dominant face of the measurement rate region.*

Theorem 5.2 (Achievability result). *Let $\{(X_i, Y_i)\}_{i=1}^N$ be a two-terminal memoryless source with (X_1, Y_1) belonging to \mathcal{L} with $d(X_1, Y_1) > 0$. Given any (ρ_x, ρ_y) satisfying*

$$\rho_x + \rho_y \geq d(X_1, Y_1), \rho_x \geq d(X_1|Y_1), \rho_y \geq d(Y_1|X_1),$$

there is a family of ϵ -REP partial Hadamard matrices with measurement rate (ρ_x, ρ_y) .

We have also the achievability result for the fully discrete distributions obtained by extending the absorption phenomenon introduced in Chapter 2.

Theorem 5.3 (Achievability result). *Let $\{(X_i, Y_i)\}_{i=1}^N$ be a discrete two-terminal memoryless source with finite entropy. Then, there is a family of ϵ -REP partial Hadamard matrices $\{\Phi_N^x, \Phi_N^y\}$ with $(\rho_x, \rho_y) = (0, 0)$.*

The proofs of Theorem 5.2 and Theorem 5.3 are given in Section 5.6.1.

5.3 Gaussian Measurement Matrices

Let $(x^n, y^n) = \{(x_i, y_i)\}_{i=1}^n$ be a realization of a two-terminal memoryless source (X, Y) with a probability distribution $p_{X,Y}$. Let $\mathbf{u} = Ax^n$ and $\mathbf{v} = By^n$ be the measurement vectors, where A is an $m_x \times n$ and B is an $m_y \times n$ matrix whose components are i.i.d. zero mean Gaussian random variables with variance $\frac{1}{m_x}$ and $\frac{1}{m_y}$ respectively. We define $\rho_x = \frac{m_x}{n}$ and $\rho_y = \frac{m_y}{n}$ as the measurement rates of the two terminals.

In order to recover the initial signal (x^n, y^n) , we propose the following joint message passing algorithm which is an extension of the single-terminal message passing proposed in [20]. We assign a variable node to each component of x^n and y^n and a check node to every measurement. Figure 5.1 shows the resulting graphical




Figure 5.1 – Graphical Model Representation for Two-Terminal Compressed Sensing. The external check nodes correspond to measurements whereas the internal check nodes between x^n and y^n represent the joint distribution $p_{X,Y}$ between (X_i, Y_i) .

model, where the internal check node between variable nodes (x_i, y_i) denotes the joint probability distribution p_{XY} . This is used to model the correlation between x_i and y_i .

Let $a, b \in [m_x]$ and $i, j \in [n]$ be the indices for check and variable nodes in T_X (X terminal) and let $c, d \in [m_y]$ and $k, l \in [n]$ be the corresponding indices for T_Y (Y terminal). The multi-terminal message passing algorithm is given by

$$r_{a \rightarrow i}^t = u_a - \sum_{j \in [n] \setminus i} A_{aj} x_{j \rightarrow a}^t, \quad (5.2)$$

$$s_{c \rightarrow k}^t = v_c - \sum_{l \in [n] \setminus k} B_{cl} y_{l \rightarrow c}^t, \quad (5.3)$$

$$x_{i \rightarrow a}^{t+1} = \eta_t^x \left(\sum_{b \in [m_x] \setminus a} A_{bi} r_{b \rightarrow i}^t, \sum_{d \in [m_y]} B_{di} s_{d \rightarrow i}^t \right), \quad (5.4)$$

$$y_{k \rightarrow c}^{t+1} = \eta_t^y \left(\sum_{b \in [m_x]} A_{bk} r_{b \rightarrow k}^t, \sum_{d \in [m_y] \setminus c} B_{kd} s_{d \rightarrow k}^t \right). \quad (5.5)$$

Notice that the only interaction between the messages in T_X and T_Y is via the threshold functions η_t^x and η_t^y . In particular, if η_t^x only depends on the first argument and if η_t^y only depends on the second argument, this message passing algorithm is transformed into two independent message passing algorithms one running on T_X and the other one on T_Y . As the measurement matrices A and B are dense matrices with columns that have ℓ_2 norms close to 1, it is possible to approximate the above message passing algorithm. This has been done heuristically in Appendix 5.6.2. The resulting MAMP (multi-terminal approximate message passing) algorithm is as follows, initialized with $r^{-1} = 0$, $s^{-1} = 0$ and $x^0 = y^0 = 0$:

$$r^t = \mathbf{u} - Ax^t - \frac{\langle \partial_1 \eta_t^x(A^* r^{t-1} + x^{t-1}, B^* s^{t-1} + y^{t-1}) \rangle}{\rho_x} r^{t-1}, \quad (5.6)$$

$$s^t = \mathbf{v} - By^t - \frac{\langle \partial_2 \eta_t^y(A^* r^{t-1} + x^{t-1}, B^* s^{t-1} + y^{t-1}) \rangle}{\rho_y} s^{t-1}, \quad (5.7)$$

$$x^{t+1} = \eta_t^x(A^* r^t + x^t, B^* s^t + y^t), \quad (5.8)$$

$$y^{t+1} = \eta_t^y(A^* r^t + x^t, B^* s^t + y^t), \quad (5.9)$$

where $r^t \in \mathbb{R}^{m_x}$ and $s^t \in \mathbb{R}^{m_y}$ are the residual terms and $x^t, y^t \in \mathbb{R}^n$ are estimates of the signals at time t and where for a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, $\partial_1 f$ and $\partial_2 f$ denote the partial derivative of f with respect to the first and the second argument respectively. Moreover, with some abuse of notation, we assume that $\eta_t(g^l, h^l) = (\eta_t(g_1, h_1), \dots, \eta_t(g_l, h_l))$ applies component-wise. Also for an n -dimensional vector u^n , $\langle u^n \rangle = \frac{1}{n} \sum_{i=1}^n u_i$ denotes the average of the elements of u^n .

It is also important to mention the appearance of Onsager terms in the Equations (5.6) and (5.7) as also mentioned in [20, 21]. This term can be considered as a second order correction for the mean field approximation of the message passing algorithm whose addition removes the correlation that exists between the fixed measurement matrices A and B and the estimated signal (x^t, y^t) in the thermodynamic limit as the system size n tends to infinity, which specially allows to completely describe the system state with a state evolution (SE) equation.

Theorem 5.4. *Let (x^n, y^n) be a realization of a memoryless source and assume that $(x^t, y^t)_{t \geq 0}$ is the output of the MAMP algorithm as in Equations (5.6)-(5.9) with Lipschitz continuous threshold functions η_t^x and η_t^y . Let $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a pseudo-Lipschitz function. Asymptotically as n tends to infinity*

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \psi(x_i, x_i^t) &\rightarrow \mathbb{E} \psi(X, \eta_t^x(X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y)), \\ \frac{1}{n} \sum_{i=1}^n \psi(y_i, y_i^t) &\rightarrow \mathbb{E} \psi(Y, \eta_t^y(X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y)) \end{aligned}$$

almost surely, where $(\tau_x^t, \tau_y^t)_{t \geq 0}$ satisfy the equation

$$\begin{aligned} \tau_x^{t+1} &= \sigma_x^2 + \frac{1}{\rho_x} \mathbb{E} (X - \eta_t^x(X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y))^2, \\ \tau_y^{t+1} &= \sigma_y^2 + \frac{1}{\rho_y} \mathbb{E} (Y - \eta_t^y(X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y))^2, \end{aligned}$$

with $\tau_x^{(-1)} = \tau_y^{(-1)} = \infty$, with Z_x, Z_y zero mean unit variance Gaussian variables independent of each other and X and Y and with σ_x^2 and σ_y^2 denoting the measurement noise variance in X and Y terminals.

Proof. The proof follows from the Bolthausen's conditioning technique used in [21] with the only difference that one should apply the conditioning to both terminals instead of a single-terminal. \square

Remark 5.4. *Theorem 5.4 provides a single letter characterization of the asymptotic behavior of the MAMP, in the sense that to estimate a specific variable (X_k, Y_k) , the effect of all the other variables is equivalent to adding a Gaussian noise with variance (τ_t^x, τ_t^y) . Moreover, replacing $\psi(a, b) = (a - b)^2$, one gets the mean square error (MSE) of the estimator*

$$\begin{aligned} \frac{\|x^{t+1} - x\|_2^2}{n} &\rightarrow \mathbb{E}(X - \eta_t^x(X + \sqrt{\tau_x^t}Z_x, Y + \sqrt{\tau_y^t}Z_y))^2, \\ \frac{\|y^{t+1} - y\|_2^2}{n} &\rightarrow \mathbb{E}(Y - \eta_t^y(X + \sqrt{\tau_x^t}Z_x, Y + \sqrt{\tau_y^t}Z_y))^2. \end{aligned}$$

We will also consider a noiseless case where $\sigma_x = \sigma_y = 0$, which using the SE equation implies that the empirical error after t iteration is given by $\rho_x \tau_x^t$ and $\rho_y \tau_y^t$. One can also simply check that choosing (η_t^x, η_t^y) to be the MMSE estimator minimizes the resulting error. We will always assume that the distribution of the signal is known and we will use the MMSE estimator for (η_t^x, η_t^y) , thus the resulting SE equation is

$$\tau_x^{t+1} = \frac{1}{\rho_x} \text{mmse}(X|X + \sqrt{\tau_x^t}Z_x, Y + \sqrt{\tau_y^t}Z_y) \quad (5.10)$$

$$\tau_y^{t+1} = \frac{1}{\rho_y} \text{mmse}(Y|X + \sqrt{\tau_x^t}Z_x, Y + \sqrt{\tau_y^t}Z_y). \quad (5.11)$$

The behavior of MAMP depends on the stable set of the SE equation. Proposition 5.1 states that for the special choice of MMSE estimators for η_t^x and η_t^y , this stable set is a fixed point.

Proposition 5.1. *For a given ρ_x, ρ_y and starting from $\tau_x^{(-1)} = \tau_y^{(-1)} = \infty$, the state vector (τ_x^t, τ_y^t) given by the SE equations in (5.10), (5.11) converges to a well-defined fixed point.*

Proof. It is sufficient to prove that the resulting sequence is non-increasing thus converging to a well-defined fixed point. We use induction on t . For $t = 0$, this obviously holds because $\tau_x^0 \leq \frac{\mathbb{E}(X^2)}{\rho_x} < \tau_x^{(-1)} = \infty$ and the same holds for τ_y^0 . From the Data Processing inequality, $(\tau_x^{t+1}, \tau_y^{t+1})$ are increasing function of (τ_x^t, τ_y^t) . Therefore, using the induction hypothesis $\tau_x^t \leq \tau_x^{t-1}$ and $\tau_y^t \leq \tau_y^{t-1}$, it immediately results that $\tau_x^{t+1} \leq \tau_x^t$ and $\tau_y^{t+1} \leq \tau_y^t$. \square

5.4 Spatially Coupled Gaussian Measurement Matrices

In the single-terminal case, it has been already observed that with traditional Gaussian matrices, the required measurement rate for complete recovery of the signal is far from the optimal rate given by the RID and spatial coupling is necessary to reduce the required measurement rate down to RID. The situation is very similar to coding theory where the BP threshold associated to a message passing algorithm is different from the optimal MAP threshold and extra spatial coupling is necessary to approach the optimal rate [70].

We briefly describe the structure of a spatially coupled measurement matrix as in [53]. We consider a band diagonal weighting matrix W of dimension $L_r \times L_c$

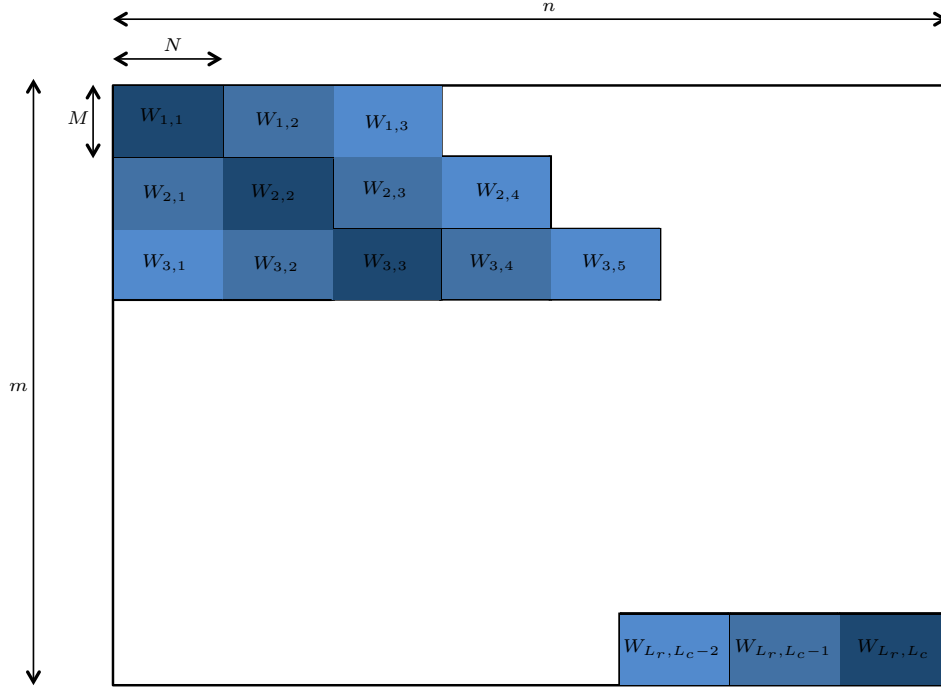


Figure 5.2 – The Structure of a band-diagonal Gaussian matrix with non-homogenous entry variances.

which is roughly row stochastic, i.e. $\frac{1}{2} \leq \sum_c W_{r,c} \leq 2$. In order to obtain the final measurement matrix, we replace every entry $W_{r,c}$ by a i.i.d. $M \times N$ Gaussian matrix with entries having variance $\frac{W_{r,c}}{M}$, thus the final matrix will be $m \times n$ where $m = ML_r$ and $n = NL_c$ and the resulting measurement rate is $\rho = \frac{m}{n} = \frac{ML_r}{NL_c}$. Figure 5.2, borrowed from [53], shows a typical structure of a band diagonal matrix.

Each component of $W_{r,c}$ corresponds to one block containing an $M \times N$ matrix. Following the notations of [53], let $\mathbf{C} = \{1, 2, \dots, L_c\}$ and $\mathbf{R} = \{1, 2, \dots, L_r\}$ denote the row and column indices of these blocks. Let us define the following operators

$$\begin{aligned} \text{mmse}_x(s_x, s_y) &= \text{mmse}(X | \sqrt{s_x}X + Z_x, \sqrt{s_y}Y + Z_y), \\ \text{mmse}_y(s_x, s_y) &= \text{mmse}(Y | \sqrt{s_x}X + Z_x, \sqrt{s_y}Y + Z_y). \end{aligned}$$

In the two-terminal case, for simplicity, we will use the same weight matrix in both terminals and the final measurement rate for each terminal can be controlled by the aspect ratios $\delta_x = \frac{M_x}{N_x}$ and $\delta_y = \frac{M_y}{N_y}$ of the corresponding sub matrices.

Definition 5.2. For a roughly stochastic matrix of dimension $L_r \times L_c$, the state evolution sequence $\{\phi^x(t), \psi^x(t)\}_{t \geq 0}$ and $\{\phi^y(t), \psi^y(t)\}_{t \geq 0}$, $\phi^o(t) = (\phi_a^o(t))_{a \in \mathbf{R}}$, $\psi^o(t) = (\psi_i^o(t))_{i \in \mathbf{C}}$ with $o \in \{x, y\}$ is defined as follows: $\psi_i^o(0) = \infty$, $i \in \mathbf{C}$, and for all $t \geq 0$,

$$\phi_a^o(t) = \sigma_o^2 + \frac{1}{\delta_o} \sum_{i \in \mathbf{C}} W_{a,i} \psi_i^o(t), \quad (5.12)$$

$$\psi_i^o(t+1) = \text{mmse}_o \left(\sum_{b \in \mathbf{R}} W_{b,i} \phi_b^x(t)^{-1}, \sum_{b \in \mathbf{R}} W_{b,i} \phi_b^y(t)^{-1} \right), \quad (5.13)$$

where σ_o^2 is the variance of the measurement noise and $\delta_o = \frac{M_o}{N_o}$ is the measurement rate of the sub-matrices for terminal $o \in \{x, y\}$.

Quantities $\psi_i(t)$ and $\phi_a(t)$ correspond to the asymptotic MSE of the MAMP. In particular, $\psi_i(t)$ is the asymptotic MSE of the variables located in block $i \in \mathbf{C}$ and $\phi_a(t)$ is the noise variance in the residual terms corresponding to row $a \in \mathbf{R}$ as we will explain later. Using a $\{\phi, \psi\}$ sequence for each terminal, it is possible to define the following MAMP algorithm. Let Q^t be an $m \times n$ matrix whose i, j component is given by

$$Q_{ij}^t = \frac{\phi_r(t)^{-1}}{\sum_{k=1}^{L_r} W_{kc} \phi_k(t)^{-1}} \quad (5.14)$$

where r is the row index of the measurement i and c is the column index of the variable j , thus it is a block constant matrix. We also define the MMSE threshold functions as

$$\eta_{t,i}^x(g_i, h_i) = \mathbb{E}(X | X + s_i^x(t)^{-1} Z_x = g_i, Y + s_i^y(t)^{-1} Z_y = h_i),$$

with $s_i^o(t) = \sum_{u \in \mathbf{R}} W_{u,c} \phi_u^o(t)^{-1}$, where c is the column index of variable i . The MMSE estimator for T_Y is defined similarly. We also assume that both estimators apply component-wise, i.e. $\eta_t(g^l, h^l) = (\eta_{t,1}(g_1, h_1), \dots, \eta_{t,l}(g_l, h_l))$. With these notations, the MAMP can be written as follows

$$x^{t+1} = \eta_t^x(x^t + (Q_x^t \odot A)^* r_x^t, y^t + (Q_y^t \odot B)^* r_y^t), \quad (5.15)$$

$$r_x^t = \mathbf{u} - Ax^t + b_x^t \odot r_x^{t-1}, \quad (5.16)$$

$$y^{t+1} = \eta_t^y(x^t + (Q_x^t \odot A)^* r_x^t, y^t + (Q_y^t \odot B)^* r_y^t), \quad (5.17)$$

$$r_y^t = \mathbf{v} - By^t + b_y^t \odot r_y^{t-1}, \quad (5.18)$$

where A and B denote the spatially coupled measurement matrices, $\mathbf{u} = Ax$, $\mathbf{v} = By$ are the measurements, r_x and r_y are residual terms, Q_x and Q_y are defined according to Equation (5.14) and b_x and b_y are defined as follows. Let $C(c)$ denote all the variables i with column index $c \in \mathbf{C}$ and let

$$\langle \partial_1 \eta_t^x \rangle_c = \left\langle \partial_1 \eta_t^x(x_i^t + ((Q_x^t \odot A)^* r_x^t)_i, y_i^t + ((Q_y^t \odot B)^* r_y^t)_i) \right\rangle$$

where the average is taken over all variables belonging to the the column block c . We define b_x^t as a column vector of length m which takes the same value for all components belonging to a row block $r \in \mathbf{R}$ and is defined as follows

$$b_{x,i}^t = \frac{1}{\delta_x} \sum_{c \in \mathbf{C}} W_{r_i,c} \tilde{Q}_{a_i,c}^{t-1} \langle \partial_1 \eta_{t-1}^x \rangle_c,$$

where r_i is the row block that i belongs to and \tilde{Q}^t is a $L_r \times L_c$ matrix defined by $\tilde{Q}_{r,c}^t = Q_{x,ij}^t$ for any i that belongs to the row block r and any column that belongs to the column block c . Notice that Q^t itself is also block-constant, therefore it is not important which i or j is taken from the block. A similar expression holds for the b_y^t by replacing $\partial_1 \eta_x^t$ with $\partial_2 \eta_y^t$, Q_x^t with Q_y^t and δ_x by δ_y .

Using similar steps as in [53], it is possible to show that the performance of the MAMP algorithm can be described by the state evolution formula given in Equation (5.12) and (5.13), where the number of states is equal to $2(L_r + L_c)$.

Theorem 5.5. *Let (x^n, y^n) be a two-terminal signal and let $\mathbf{u} = Ax^n$ and $\mathbf{v} = By^n$, where A and B are spatially coupled matrices with the same weight matrix W . Let (x^t, y^t) be the output of the MAMP algorithm described by Equations (5.15)-(5.18), where $\{\phi^o(t), \psi^o(t)\}_{t \geq 0, o \in \{x, y\}}$ is obtained from the SE equation (5.12)-(5.13). Asymptotically, as N_x, N_y go to infinity*

$$\frac{1}{N_x} \sum_{j \in C(i)} (x_j - x_j^t)^2 \rightarrow \psi_i^x(t), \quad \frac{1}{N_y} \sum_{j \in C(i)} (y_j - y_j^t)^2 \rightarrow \psi_i^y(t).$$

Based on the results proved in [53] for the single-terminal case and the lower bound proved in Theorem 5.1, it is possible to give the following characterization for the achievable measurement rate region in the multi-terminal case.

Theorem 5.6. *Let (X, Y) be a linearly correlated two-terminal source and let $\rho_x, \rho_y \in [0, 1]$ be such that*

$$\rho_x > d(X|Y), \quad \rho_y > d(Y|X), \quad \rho_x + \rho_y > d(X, Y). \quad (5.19)$$

There is an ensemble of spatially coupled measurement matrices that separately captures the signals in the two-terminals and an MAMP algorithm that jointly recovers the signals in each terminal with a negligible distortion.

Remark 5.5. *The optimal measurement rate region given by Equation (5.19) is very similar to the Slepian-Wolf rate region for distributed source coding where the RID in the compressed sensing setting plays a role similar to the discrete entropy in distributed source coding.*

Proof. We prove that the corner points $(d(X), d(Y|X))$ and $(d(Y), d(X|Y))$ are achievable under MAMP. In the single-terminal case, if $\rho_x > d(X)$ asymptotically the signal in T_X can be recovered with a negligible distortion. In the multi-terminal case, if we consider only the terms related to T_X , from Equation (5.12)-(5.13), we have

$$\begin{aligned} \phi_a^x(t) &= \frac{1}{\rho_x} \sum_{i \in C} W_{a,i} \psi_i^x(t), \\ \psi_i^x(t+1) &= \text{mmse}_x \left(\sum_{b \in R} W_{b,i} \phi_b^x(t)^{-1}, \sum_{b \in R} W_{b,i} \phi_b^y(t)^{-1} \right). \end{aligned}$$

As η_x, η_y are MMSE estimators, from the Data Processing inequality, one can check that $\text{mmse}_x(s_x, s_y)$ and $\text{mmse}_y(s_x, s_y)$ are decreasing functions of s_x and s_y . This implies that $\text{mmse}_x(\sum_{b \in R} W_{b,i} \phi_b^x(t)^{-1}, \sum_{b \in R} W_{b,i} \phi_b^y(t)^{-1})$ is less than or equal to $\text{mmse}_x(\sum_{b \in R} W_{b,i} \phi_b^x(t)^{-1}, 0)$, which is equal to the variance of the MMSE estimator for X which does not use the information of Y . One can also check that the SE equation is increasing with respect to $\psi_i^x(t)$, which implies that the ψ^x sequence for the MAMP is dominated by the ψ^x sequence of a single-terminal AMP, which converges to 0 for any $\rho_x > d(X)$ as proved in [53]. If $\psi_i^x(t)$ converges to zero so does the ϕ_a^x sequence, thus the SE equation for T_Y will be as follows

$$\begin{aligned} \phi_a^y(t) &= \frac{1}{\rho_y} \sum_{i \in C} W_{a,i} \psi_i^y(t), \\ \psi_i^y(t+1) &= \text{mmse}_y(\infty, \sum_{b \in R} W_{b,i} \phi_b^y(t)^{-1}), \end{aligned}$$

which using the same steps as in the single-terminal case, can be proved to converge to zero provided that

$$\rho_y > \limsup_{s \rightarrow \infty} s \text{mmse}(Y|X, \sqrt{s}Y + Z_y) = d(Y|X),$$

where Z_y is a zero mean unit variance Gaussian noise and where we used the fact that for the class of linearly correlated signals that we use, $d(Y|X)$ is well defined.

Similarly, it is possible to prove that $(\rho_x, \rho_y) = (d(X|Y), d(Y))$ is also achievable. Furthermore, any point on the dominant face is also achievable because if we consider two ensembles of measurement matrices (A_1, B_1) and (A_2, B_2) with rate vectors $\vec{R}_1 = (d(X), d(Y|X))$ and $\vec{R}_2 = (d(Y), d(X|Y))$ achieving the two corner points respectively, by diagonally concatenating r copies of the former with s copies of the latter, one can get an ensemble with measurement rate $\frac{r}{r+s}\vec{R}_1 + \frac{s}{r+s}\vec{R}_2$ and a negligible distortion.

The other points on the region are also achieved because their measurement rate is larger than or equal to the measurement rate of at least one point on the dominant face, thus their distortion will be asymptotically negligible as well. \square

5.5 Simulation Results

5.5.1 Signal Model

For simulation, we will use a linearly correlated random vector from \mathcal{L}_2 whose independent constituents are random variables with Bernoulli-Gaussian distribution. Let Z^k be a sequence of independent random variables with a probability distribution $p_i(z) = (1 - \alpha_i)\delta_0(z) + \alpha_i\mathbf{N}(0, \frac{1}{\alpha_i}, z)$ where δ_0 is a delta measure at point zero and $\mathbf{N}(0, \sigma^2, z)$ denotes the distribution of a zero mean Gaussian distribution with variance σ^2 . One can simply check that $\text{Var}(Z_i) = 1$ and $d(Z_i) = \alpha_i$. Let Φ be a $2 \times k$ real-valued matrix. The two-terminal linearly correlated source is given by ΦZ^k . For this class of signals, the joint and the conditional RID are well-defined. Notice that depending on the values of α_i and the structure of the matrix Φ , this model can cover a wide variety of correlations between the signals in two terminals. In Appendix 5.6.4, we have obtained a closed-form expression for the MMSE estimator (η^x, η^y) of this source in the presence of the Gaussian measurement noise which we will use as a denoising (threshold) function in MAMP algorithm.

5.5.2 Performance without Spatial Coupling

In this section, we use the message passing algorithm given by Equations (5.6)-(5.9) to recover a linearly correlated Bernoulli-Gaussian signal for the noiseless measurements taken from both terminals.

Comparison of the Empirical Results and SE predictions

We consider a very simple case where Z_1, Z_2, Z_3 are three Bernoulli-Gaussian random variables with $d(Z_1) = d(Z_3) = 0.2$ and $d(Z_2) = 0.3$. The signal for the two terminals is given by $X = Z_1 + Z_2$ and $Y = Z_2 + Z_3$, thus Z_1 and Z_3 are the private parts of the signals and Z_2 is the common part which creates correlation between X and Y . It is easy to check that $d(X) = d(Y) = 0.44$ and $d(X|Y) = d(Y|X) = 0.248$.

Figures 5.3, 5.4 show the simulation results for $\rho_x = 0.5, \rho_y = 0.6$. It is seen that there is a good match between the empirical variance of the estimator and the predictions of the SE. Moreover, the algorithm can not fully recover the signal which means that the SE equation has a fixed point other than $(\tau_x, \tau_y) = (0, 0)$. The simulations has been repeated by increasing the measurement rate of the T_Y from $\rho_y = 0.6$ to 0.7. The simulation results have been depicted in Figures 5.5 and 5.6. Plots show that this time the MAMP algorithm successfully recovers the signal of both terminals. It is also important to notice that because of the correlation between the terminals, increasing ρ_y is helpful for recovering the signal in T_X .




Figure 5.3 – Empirical and State Evolution result for T_X for $\rho_x = 0.5, \rho_y = 0.6$

Rate-Distortion Region

In this part, we run the MAMP algorithm for the same signal as in Section 5.5.2 for different measurement rates. As a distortion measure, we consider the average of the mean square error of the two terminals. Figures 5.7, 5.8, and 5.9 show a contour plot of the Rate-Distortion curve for three sources with the same individual but different conditional RID. The dashed lines show the boundary of the optimal pentagon. Low distortion recovery is not possible outside this region.

In the extreme case where the signals in two terminals are independent from each other, i.e., there is no common signal, the pentagon region reduces to a square region. However, if there is no private signal then the signals in both terminals are the same and the problem is reduced to a simple single-terminal problem. Obviously in this case, because of the independence of measurement matrices in the two terminals, individual measurement rates are not important as far as their sum is large enough.

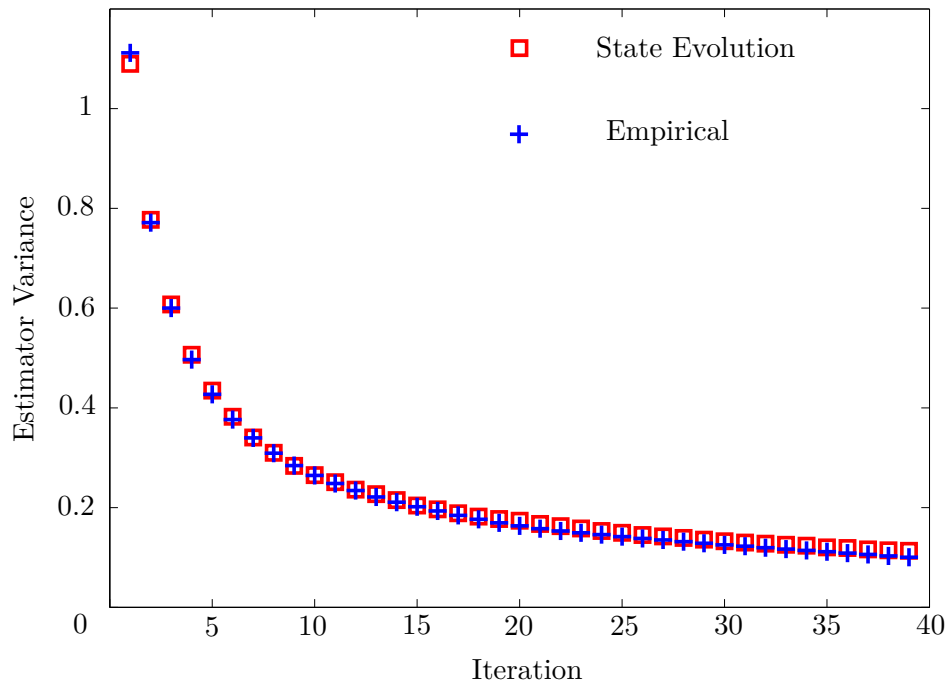


Figure 5.4 – Empirical and State Evolution result for T_Y for $\rho_x = 0.5, \rho_y = 0.6$




Figure 5.5 – Empirical and State Evolution result for T_X for $\rho_x = 0.5, \rho_y = 0.7$

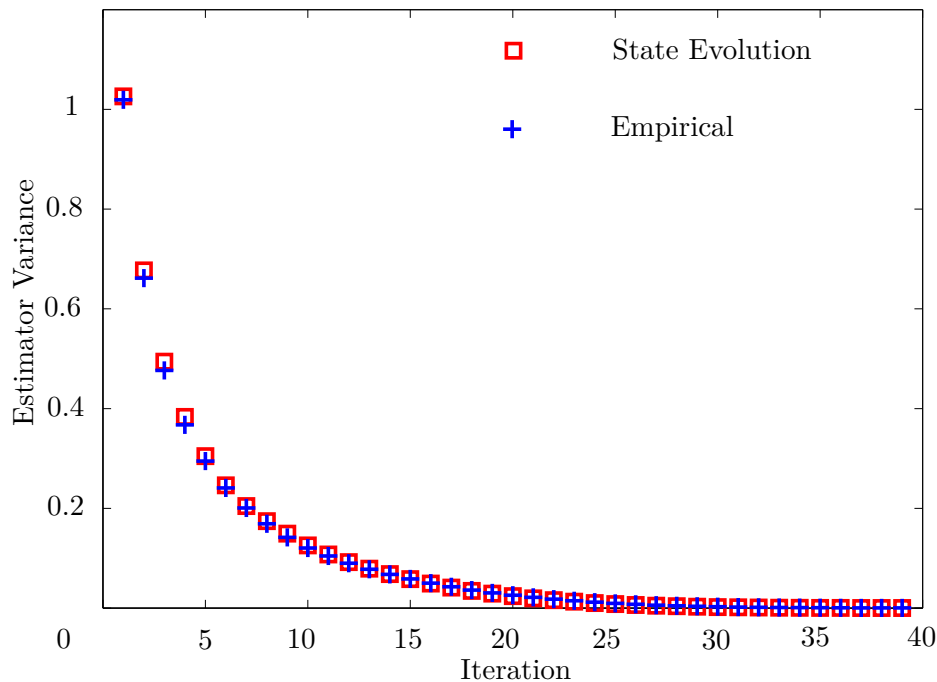


Figure 5.6 – Empirical and State Evolution result for T_Y for $\rho_x = 0.5, \rho_y = 0.7$

This can be seen from Figures 5.7, 5.8, and 5.9, where we keep $d(X) = d(Y) = 0.44$ but gradually increase the share of the common signal, thus, as a result $d(X|Y)$ and $d(Y|X)$ start to decrease. It is observed that the contour lines gradually become parallel with $\rho_x + \rho_y = \text{constant}$.

Notice that there is a huge gap between the low-distortion curve (distortion equal to 0.1) and the optimal region. As we will see this gap is filled by using spatial coupling and running MAMP.

Effect of Correlation between the Terminals

In order to investigate the effect of correlation between the two terminals, we have plotted a low distortion contour of the three sources with the same $d(X) = d(Y) = 0.44$ but three different conditional RID 0.248, 0.1820 and 0.0916. Decreasing the conditional RID while fixing the individual entropy, makes the signals in two terminals more correlated. A low distortion curve of the three sources is plotted in Figure 5.10. The plot shows that the required measurement rate is decreasing by increasing the correlation.

5.5.3 Performance with Spatial Coupling

In this section, we simulate the SE equation for the MAMP algorithm. We consider the same source as in Section 5.5.2 where $d(X) = d(Y) = 0.44$ and $d(X|Y) = d(Y|X) = 0.248$. In order to approach the corner point $(d(X), d(Y|X))$, we consider a measurement rate with 10 percent oversampling, i.e., $\rho_x = 1.1d(X)$ and $\rho_y = 1.1d(Y|X)$. The simulation results are shown in Figure 5.11. Similar to the single-

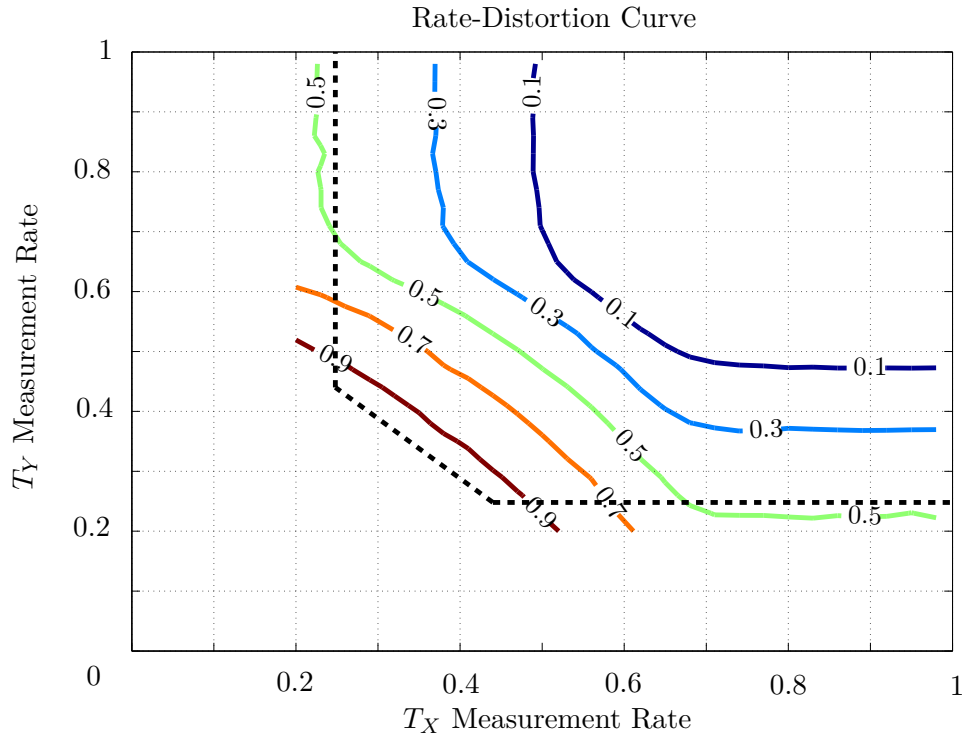


Figure 5.7 – Rate-Distortion region for a linearly correlated Bernoulli-Gaussian source with $d(X) = d(Y) = 0.44$ and $d(X|Y) = d(Y|X) = 0.248$. The dashed lines show the boundaries of the optimal region.




Figure 5.8 – Rate-Distortion Region for the Linearly Correlated Bernoulli-Gaussian Source with $d(X) = d(Y) = 0.44$ and $d(X|Y) = d(Y|X) = 0.1802$. The dashed lines show the boundaries of the optimal region.




Figure 5.9 – Rate-Distortion region for a linearly correlated Bernoulli-Gaussian source with $d(X) = d(Y) = 0.44$ and $d(X|Y) = d(Y|X) = 0.0916$. The dashed lines show the boundary of the optimal region.

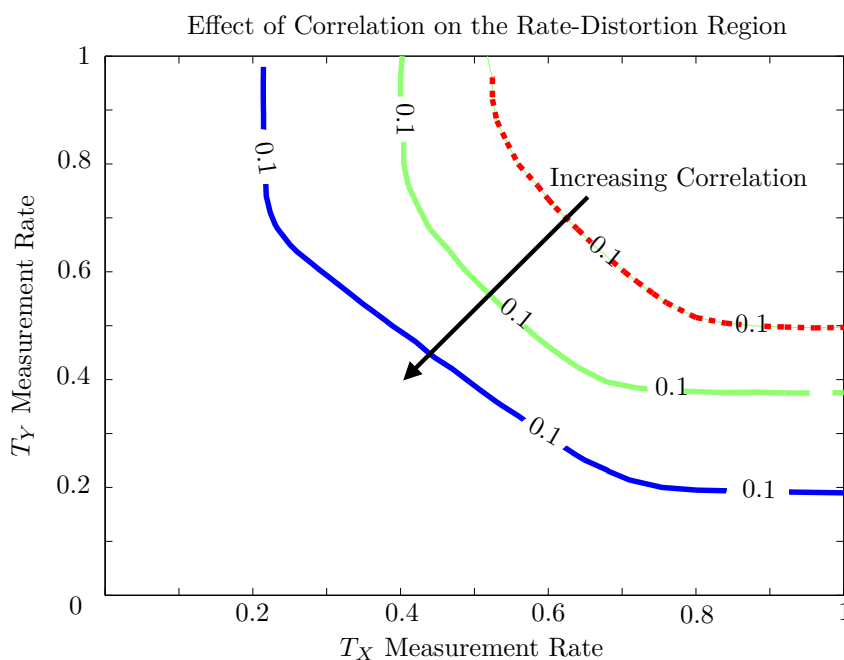


Figure 5.10 – Effect of Correlation on the Measurement Rate Region. The low distortion curve of three different two-terminal sources with the same individual RID is plotted. The required measurement region of the more correlated source is dominated by that of the less correlated one.

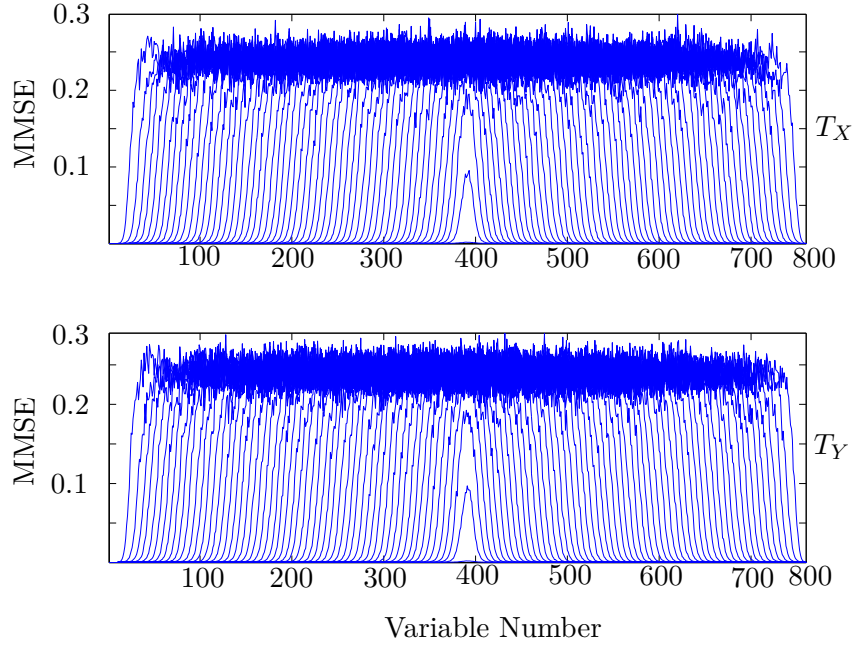


Figure 5.11 – Spatial Coupling Wave for A Linearly Correlated Source with $\rho_x = 1.1d(X)$ and $\rho_y = 1.1d(Y|X)$

terminal case, one can observe a wave-like phenomenon that starts from the boundary variables and proceeds towards the center recovering all the variables gradually. In particular, to create the initial wave at the boundary, one needs to oversample the boundary variables. Figure 5.12 depicts the simulation results for another experiment where ρ_x is kept fixed but ρ_y is reduced. It is observed that, this time spatial coupling wave proceeds to decode the variables in T_X . However the initially generated wave in T_Y stops after a while and can not proceed to recover all the variables in T_Y .

By checking the results for the non-spatially coupled case, one can see that the resulting MSE error decreases gradually by increasing the measurement rate. On the contrary, in the spatially coupled case, either the generated wave proceeds and recovers all the variables or it stops, thus asymptotically, there is a sharp transition in the resulting MSE in terms of measurement rate.

For the same source, we have done the simulations to find boundary of the phase transition. Figure 5.13 depicts the simulation result. It is seen that there is a good match between the simulation and the boundary predicted information theoretically.

5.6 Appendix

5.6.1 Proofs of the Hadamard Construction

For $n \in \mathbb{N}$ and $N = 2^n$, let $\{(X_i, Y_i)\}_{i=1}^N$ be a sequence of random vectors in the space \mathcal{L} , with joint and conditional RID $d(X, Y)$, $d(X|Y)$ and $d(Y|X)$. Let $Z_1^N = H_N X_1^N$ and assume that $W_1^N = H_N Y_1^N$. Let us define two processes I_n and J_n as follows.

$$I_n(i) = d(Z_i|Z_1^{i-1}), J_n(i) = d(W_i|W_1^{i-1}, Z_1^N), i \in [N].$$




Figure 5.12 – Spatial Coupling Wave for A Linearly Correlated Source with $\rho_x = 1.1d(X)$ and $\rho_y < d(Y|X)$




Figure 5.13 – Phase Transition Boundary for MAMP and Comparison with SE Prediction. Dashed curve shows the theoretical boundary of the achievable measurement rate region.

We can label I_n and J_n by a sequence of b_1^n as done for the single-terminal case in Section 4.3.1, and convert them to stochastic processes $I_n = I^{B_1 B_2 \dots B_n}$ and $J_n = J^{B_1 B_2 \dots B_n}$. By this definition, one can prove the following proposition.

Proposition 5.2 (Multi-terminal RID polarization). *($I_n, \mathcal{F}_n, \mathbb{P}$) and ($J_n, \mathcal{F}_n, \mathbb{P}$) are erasure stochastic processes with initial values $d(X)$ and $d(Y|X)$, both polarizing to $\{0, 1\}$.*

Proof. For the initial value, we have $I_0(1) = d(X_1)$ and $J_0(1) = d(Y_1|X_1)$. As $\{(X_i, Y_i)\}_{i=1}^N$ is a memoryless source, similar to the single-terminal case, it is easy to see that I is an erasure process with initial value $d(X_1)$ and it remains to show that J is also an erasure process but with initial value $d(Y_1|X_1)$.

Let \tilde{H}^{i-1} , \tilde{H}^i and \tilde{h}_i denote the first $i-1$ rows, the first i rows, and the i -th row of \tilde{H}_N . As $(X_1, Y_1) \in \mathcal{L}$, there is a sequence of i.i.d. nonsingular random variables E_1^k and two vectors a_1^k and b_1^k such that $X_1 = \sum_{i=1}^k a_i E_i$ and $Y_1 = \sum_{i=1}^k b_i E_i$. As $\{(X_i, Y_i)\}_{i=1}^N$ is memoryless, there is a concatenation of a sequence of i.i.d. copies of E_1^k , $E = [E_1^k(1); E_1^k(2); \dots; E_1^k(N)]$, such that

$$Z_1^N = \tilde{H}_N X_1^N = [\tilde{H}_N \otimes (a_1^k)^t] E, \quad W_1^N = \tilde{H}_N Y_1^N = [\tilde{H}_N \otimes (b_1^k)^t] E,$$

where \otimes denotes the Kronecker product and $(a_1^k)^t, (b_1^k)^t$ are the transpose of the column vectors a_1^k and b_1^k . Let

$$\Gamma = \{\Theta_1, \Theta_2, \dots, \Theta_N\} \tag{5.20}$$

be the random element corresponding to the Θ pattern of $E_1^k(j)$, $j \in [N]$, where $\Theta_j \in \{0, 1\}^k$, $j \in [N]$. Using the rank result developed for the RID, it is easy to see that for every $i \in [N]$, we have the following:

$$J_n(i) = d(W_i | W_1^{i-1}, Z_1^N) = \mathbb{E}\{I([\tilde{H}^{i-1} \otimes (b_1^k)^t; \tilde{H} \otimes (a_1^k)^t; \tilde{h}_i \otimes (b_1^k)^t][C_\Gamma])\}.$$

For $i \in [N]$, let $\mathbf{1}_i(\Theta_1^N) \in \{0, 1\}$ denote the random increase of rank of $[\tilde{H}^{i-1} \otimes (a_1^k)^t]_{C_\Gamma}$ by adding $\tilde{h}_i \otimes (a_1^k)^t$. Now, consider the stage $n+1$, where we are going to combine two copies of \tilde{H}_N to construct the matrix \tilde{H}_{2N} . Then, the row i corresponding to W_i is split into two new rows i^+ and i^- , which correspond to the row number $2i-1$ and the row number $2i$ shown below:

$$\begin{pmatrix} \tilde{H}_N \otimes (a_1^k)^t & , & \tilde{H}_N \otimes (a_1^k)^t \\ \tilde{H}_N \otimes (a_1^k)^t & , & -\tilde{H}_N \otimes (a_1^k)^t \\ \vdots & , & \vdots \\ \tilde{h}_{i-1} \otimes (b_1^k)^t & , & \tilde{h}_{i-1} \otimes (b_1^k)^t \\ \tilde{h}_{i-1} \otimes (b_1^k)^t & , & -\tilde{h}_{i-1} \otimes (b_1^k)^t \\ \tilde{h}_i \otimes (b_1^k)^t & , & \tilde{h}_i \otimes (b_1^k)^t \end{pmatrix}$$

Similar to the single-terminal case, we see that adding $\tilde{h}_i \otimes (b_1^k)^t$ increases the rank of the matrix if it increases the rank of either the first or the second block. In other words,

$$\mathbf{1}_{2i-1}(\Theta_1^{2N}) = \mathbf{1}_i(\Theta_1^N) + \mathbf{1}_i(\Theta_{N+1}^{2N}) - \mathbf{1}_i(\Theta_1^N) \mathbf{1}_i(\Theta_{N+1}^{2N}),$$

where $\mathbf{1}_i(\Theta_1^N), \mathbf{1}_i(\Theta_{N+1}^{2N}) \in \{0, 1\}$ are the corresponding amount of increase of the rank of the first and second block by adding the i -th row. In particular, Θ_1^N and Θ_{N+1}^{2N} are i.i.d. so are $\mathbf{1}_i(\Theta_1^N)$ and $\mathbf{1}_i(\Theta_{N+1}^{2N})$. Taking the expectation value, similar to what did in the single-terminal case, we obtain that

$$J_n(i)^+ = 2J_n(i) - J_n(i)^2. \quad (5.21)$$

Moreover, one can also show that for $i \in [N]$,

$$\frac{J_n(i)^+ + J_n(i)^-}{2} = J_n(i),$$

which together with (5.21), implies that $J_n(i)^- = J_n(i)^2$. Therefore, J is also an erasure process with initial value $d(Y|X)$. Similar to the single-terminal case, one can also show that the shuffled Hadamard matrix \tilde{H}_N can be replaced with H_N and the permutation matrix B_N is not necessary. \square

The next step is to construct the two-terminal measurement ensemble. Let $n \in \mathbb{N}$ and $N = 2^n$. We will construct Φ_N^x by selecting those rows of the Hadamard matrix, H_N , with $I_n(i) \geq \epsilon d(X)$. Similarly, Φ_N^y is constructed by selecting those rows of H_N with $J_n(i) \geq \epsilon d(Y|X)$. It remains to prove that the family $\{\Phi_N^x, \Phi_N^y\}$ labeled with N , a power of 2, and of dimension $m_N^x \times N$ and $m_N^y \times N$ is ϵ -REP with measurement rate $(d(X), d(Y|X))$. By this construction, we can achieve one of the corner points of the dominant face of the rate region. If we switch the role of X and Y we will get the other corner point $(d(X|Y), d(Y))$. One way to obtain any point on the dominant face is to use time sharing for the two family. However, it is also possible to use an explicit construction proposed in [71], which directly gives any point on the dominant face of the measurement rate region without any need to time sharing. We will just prove the achievability for the corner point $(d(X), d(Y|X))$.

Proof of Theorem 5.2. We first show that the family $\{\Phi_N^x, \Phi_N^y\}$ has measurement rate $(d(X), d(Y|X))$. Notice that the processes I_n^x, I_n^y converge almost surely thus, they converge in probability. Specifically, considering the uniform probability assumption and using a similar technique as we used in the single-terminal case, we get the following:

$$\begin{aligned} \limsup_{N \rightarrow \infty} \frac{m_N^x}{N} &= \limsup_{N \rightarrow \infty} \frac{\#\{i \in [N] : I_n^x(i) \geq \epsilon d(X)\}}{N} \\ &= \limsup_{n \rightarrow \infty} \mathbb{P}(I_n^x \geq \epsilon d(X)) = \mathbb{P}(I_\infty^x \geq \epsilon d(X)) = d(X). \end{aligned}$$

Similarly, we can show that $\limsup_{N \rightarrow \infty} \frac{m_N^y}{N} = d(Y|X)$.

It remains to prove that $\{\Phi_N^x, \Phi_N^y\}$ is ϵ -REP. Let $S_X = \{i \in [N] : I_n(i) \geq \epsilon d(X)\}$ and $S_Y = \{i \in [N] : J_n(i) \geq \epsilon d(Y|X)\}$ denote the selected rows to construct $\{\Phi_N^x, \Phi_N^y\}$ and let $Z_1^N = H_N X_1^N$ and $W_1^N = H_N Y_1^N$ be the full measurements for the x and the y terminal. Let $B_i^X = S_X^c \cap [i-1]$ and $B_i^Y = S_Y^c \cap [i-1]$ be the set of

all indices in S_X^c and S_Y^c less than i . We have

$$\begin{aligned}
d(X_1^N, Y_1^N | Z_{S_X}, W_{S_Y}) &= d(Z_1^N, W_1^N | Z_{S_X}, W_{S_Y}) \\
&\leq d(Z_1^N | Z_{S_X}) + d(W_1^N | Z_1^N, W_{S_Y}) \\
&\leq \sum_{i \in S_X^c} d(Z_i | Z_{B_i^X}, Z_{S_X}) + \sum_{i \in S_Y^c} d(W_i | W_{B_i^Y}, W_{S_Y}, Z_1^N) \\
&\leq \sum_{i \in S_X^c} d(Z_i | Z_1^{i-1}) + \sum_{i \in S_Y^c} d(W_i | W_1^{i-1}, Z_1^N) \\
&\leq N\epsilon d(X) + N\epsilon d(Y|X) \\
&= \epsilon N d(X, Y) = \epsilon d(X_1^N, Y_1^N),
\end{aligned}$$

which shows the ϵ -REP property for the ensemble $\{\Phi_N^x, \Phi_N^y\}$. \square

Achievability proof for the discrete case: In the fully discrete case, the construction is very similar to the mixture case with the only difference that instead of using the RID, we will use the entropy. Similar to the single-terminal case, we can prove the following.

Lemma 5.1. $(I_n, \mathcal{F}_n, \mathbb{P})$ and $(J_n, \mathcal{F}_n, \mathbb{P})$ are positive martingale converging to 0 almost surely.

We again construct the family $\{\Phi_N^x, \Phi_N^y\}$ by selecting those rows of H_N with $I_n \geq \epsilon H(X)$ and $J_n \geq \epsilon H(Y|X)$.

Proof of Theorem 5.3. Similar to the single-terminal case, it is easy to show that $\{\Phi_N^x, \Phi_N^y\}$ has measurement rate $(0, 0)$. It remains to prove that $\{\Phi_N^x, \Phi_N^y\}$ is ϵ -REP. Let $S_X = \{i \in [N] : I_n(i) \geq \epsilon H(X)\}$ and $S_Y = \{i \in [N] : J_n(i) \geq \epsilon H(Y|X)\}$ denote the selected rows to construct $\{\Phi_N^x, \Phi_N^y\}$ and let $Z_1^N = H_N X_1^N$ and $W_1^N = H_N Y_1^N$ be the full measurements for the X and the Y terminal. Let $B_i^X = S_X^c \cap [1 : i - 1]$ and $B_i^Y = S_Y^c \cap [1 : i - 1]$ be the set of all indices in S_X^c and S_Y^c less than i . We have the following:

$$\begin{aligned}
H(X_1^N, Y_1^N | Z_{S_X}, W_{S_Y}) &= H(Z_1^N, W_1^N | Z_{S_X}, W_{S_Y}) \\
&\leq H(Z_1^N | Z_{S_X}) + H(W_1^N | Z_1^N, W_{S_Y}) \\
&\leq \sum_{i \in S_X^c} H(Z_i | Z_{B_i^X}, Z_{S_X}) + \sum_{i \in S_Y^c} H(W_i | W_{B_i^Y}, W_{S_Y}, Z_1^N) \\
&\leq \sum_{i \in S_X^c} H(Z_i | Z_1^{i-1}) + \sum_{i \in S_Y^c} H(W_i | W_1^{i-1}, Z_1^N) \\
&\leq N\epsilon H(X) + N\epsilon H(Y|X) \\
&= \epsilon N H(X, Y) = \epsilon H(X_1^N, Y_1^N),
\end{aligned}$$

which shows the ϵ -REP property for the two-terminal ensemble $\{\Phi_N^x, \Phi_N^y\}$. \square

5.6.2 Heuristic Derivation of the Multi-Terminal AMP

In this section, we try to heuristically obtain an approximation to the message passing algorithm given by equations (5.2)-(5.5). Our derivation is similar to the heuristic

derivation of the single-terminal AMP in [21]. Intuitively as the measurement matrices A and B are dense, any two messages emanating from the same check node are only slightly different from each other. The same is true for the messages emanating from a variable node. For example, if one considers messages from check nodes to variable nodes

$$r_{a \rightarrow i}^t = u_a - \sum_{j \in [n] \setminus i} A_{aj} x_{j \rightarrow a}^t = u_a - \sum_{j \in [n]} A_{aj} x_{j \rightarrow a}^t + A_{ai} x_{i \rightarrow a}^t,$$

it is seen that for a fixed $a \in [m_x]$, $r_{a \rightarrow i}^t$ for different values of $i \in [n]$ are different because of the appearance of the last term $A_{ai} x_{i \rightarrow a}^t$ which is of the order $O(\frac{1}{\sqrt{m_x}}) \approx O(\frac{1}{\sqrt{n}})$ as m_x and n are assumed to be proportional. Similarly, considering the messages from variable nodes to check nodes,

$$\begin{aligned} x_{i \rightarrow a}^{t+1} &= \eta_t^x \left(\sum_{b \in [m_x] \setminus a} A_{bi} r_{b \rightarrow i}^t, \sum_{c \in [m_y]} A_{ci} s_{c \rightarrow i}^t \right) \\ &= \eta_t^x \left(\sum_{b \in [m_x]} A_{bi} r_{b \rightarrow i}^t - A_{ai} r_{a \rightarrow i}^t, \sum_{c \in [m_y]} A_{ci} s_{c \rightarrow i}^t \right), \end{aligned}$$

it is observed that for a fixed $i \in [n]$, the difference of messages $x_{i \rightarrow a}^{t+1}$ for different values of $a \in [m_x]$ is again of the order $O(\frac{1}{\sqrt{n}})$. Therefore, one gets

$$\begin{aligned} r_{a \rightarrow i}^t &= r_a^t + \delta r_{a \rightarrow i}^t, \quad s_{c \rightarrow k}^t = s_c^t + \delta s_{c \rightarrow k}^t, \\ x_{i \rightarrow a}^t &= x_i^t + \delta x_{i \rightarrow a}^t, \quad y_{i \rightarrow a}^t = y_i^t + \delta y_{i \rightarrow a}^t, \end{aligned}$$

where the δ terms are of the order $O(\frac{1}{\sqrt{n}})$. Replacing in Equation (5.2) and (5.3), one obtains

$$\begin{aligned} r_a^t + \delta r_{a \rightarrow i}^t &= u_a - \sum_{j \in [n]} A_{aj} (x_j^t + \delta x_{j \rightarrow a}^t) + A_{ai} (x_i^t + \delta x_{i \rightarrow a}^t), \\ s_c^t + \delta s_{c \rightarrow k}^t &= v_c - \sum_{l \in [n]} B_{cl} (y_l^t + \delta y_{l \rightarrow c}^t) + B_{ck} (y_k^t + \delta y_{k \rightarrow c}^t). \end{aligned}$$

The terms $A_{ai} \delta x_{i \rightarrow a}^t$ and $B_{ck} \delta y_{k \rightarrow c}^t$ are of the order $O(\frac{1}{n})$ and negligible asymptotically. Thus, one obtains that

$$r_a^t = u_a - \sum_{j \in [n]} A_{aj} (x_j^t + \delta x_{j \rightarrow a}^t), \quad \delta r_{i \rightarrow a}^t = A_{ai} x_i^t, \quad (5.22)$$

$$s_c^t = v_c - \sum_{l \in [n]} B_{cl} (y_l^t + \delta y_{l \rightarrow c}^t), \quad \delta s_{k \rightarrow c}^t = B_{ck} y_k^t. \quad (5.23)$$

Replacing in Equations (5.4) and (5.5), it results that

$$\begin{aligned} x_i^{t+1} + \delta x_{i \rightarrow a}^{t+1} &= \eta_t^x \left(\sum_{b \in [m_x] \setminus a} A_{bi} (r_b^t + A_{bi} x_i^t), \sum_{d \in [m_y]} B_{di} (s_d^t + B_{di} y_i^t) \right) \\ &= \eta_t^x \left(\sum_{b \in [m_x]} A_{bi} (r_b^t + A_{bi} x_i^t), \sum_{d \in [m_y]} B_{di} (s_d^t + B_{di} y_i^t) \right) \\ &\quad + \partial_1 \eta_t^x (\dots, \dots) A_{ai} (r_a^t + A_{ai} x_i^t). \end{aligned}$$

This implies that

$$x_i^{t+1} = \eta_t^x(x_i^t + \sum_{b \in [m_x]} A_{bi} r_b^t, y_i^t + \sum_{d \in [m_y]} B_{di} s_d^t), \quad (5.24)$$

$$\delta x_{i \rightarrow a}^{t+1} = \partial_1 \eta_t^x(x_i^t + \sum_{b \in [m_x]} A_{bi} r_b^t, y_i^t + \sum_{d \in [m_y]} B_{di} s_d^t) A_{ai} r_a^t, \quad (5.25)$$

where one uses the fact that for any $i \in [n]$, $\sum_{a \in [m_x]} A_{ai}^2 \approx 1$, and $A_{ai} \delta r_{a \rightarrow i}^t = O(\frac{1}{n})$ thus negligible as n tends to infinity. A similar argument holds for the T_Y giving

$$y_k^{t+1} = \eta_t^y(x_k^t + \sum_{b \in [m_x]} A_{bk} r_b^t, y_k^t + \sum_{d \in [m_y]} B_{dk} s_d^t), \quad (5.26)$$

$$\delta y_{k \rightarrow c}^{t+1} = \partial_2 \eta_t^y(x_k^t + \sum_{b \in [m_x]} A_{bk} r_b^t, y_k^t + \sum_{d \in [m_y]} B_{dk} s_d^t) B_{ck} r_c^t. \quad (5.27)$$

Replacing (5.25) in (5.22) and (5.27) in (5.23), and using the approximation $A_{ai}^2 \approx \frac{1}{m_x}$ and $B_{ck}^2 \approx \frac{1}{m_y}$, one obtains that

$$\begin{aligned} r_a^t &= u_a - A_a x^t + \frac{\sum_{j \in [n]} \partial_1 \eta_t^x(x_j^{t-1} + \dots, y_j^{t-1} + \dots)}{m_x} r_a^{t-1} \\ &= u_a - A_a x^t + \frac{\langle \partial_1 \eta_t^x(x^{t-1} + A^* r^{t-1}, y^{t-1} + B^* s^{t-1}) \rangle}{\rho_x} r_a^{t-1}, \end{aligned} \quad (5.28)$$

where A_a denotes the a -th row of the matrix A . Similarly,

$$\begin{aligned} s_c^t &= v_c - B_c y^t + \frac{\sum_{l \in [n]} \partial_2 \eta_t^y(x_l^{t-1} + \dots, y_l^{t-1} + \dots)}{m_y} s_c^{t-1} \\ &= v_c - B_c y^t + \frac{\langle \partial_2 \eta_t^y(x^{t-1} + A^* r^{t-1}, y^{t-1} + B^* s^{t-1}) \rangle}{\rho_y} s_c^{t-1}. \end{aligned} \quad (5.29)$$

Equations (5.24), (5.26), (5.28) and (5.29) give the the MAMP algorithm.

5.6.3 Heuristic Derivation of the State Evolution

To give an intuitive justification (as in [21]) for the validity of SE for the two-terminal AMP in Equations (5.10) and (5.11), consider the following version of the AMP where at each iteration t , the measurement matrices A and B are replaced with independent copies and where we drop the Onsager term in Equations (5.6) and (5.7). In other words, let $\mathbf{u}^t = A(t)\mathbf{x}_0 + w_x$ and $\mathbf{v}^t = B(t)\mathbf{y}_0 + w_y$ be the noisy measurements at iteration t , where w_x and w_y are additive noises consisting of i.i.d. zero mean with variance σ_x^2 and σ_y^2 respectively. The new AMP algorithm can be written as follows

$$\begin{aligned} r^t &= \mathbf{u}^t - A(t)x^t, \quad x^{t+1} = \eta_t^x(A(t)^* r^t + x^t, B(t)^* s^t + y^t), \\ s^t &= \mathbf{v}^t - B(t)y^t, \quad y^{t+1} = \eta_t^y(A(t)^* r^t + x^t, B(t)^* s^t + y^t). \end{aligned}$$

The first equation can be simplified to the following form

$$\begin{aligned} x^{t+1} &= \eta_t^x(\mathbf{x}_0 + A(t)^* w_x + (I - A(t)^* A(t))(x^t - \mathbf{x}_0), \\ &\quad \mathbf{y}_0 + B(t)^* w_y + (I - B(t)^* B(t))(y^t - \mathbf{y}_0)). \end{aligned}$$

Conditioned on w_x , $A(t)^*w_x$ is an n dimensional vector with i.i.d. Gaussian components with zero mean and variance $\frac{\|w_x\|_2^2}{n} \approx \sigma_x^2$. Moreover, in the asymptotic limit as n gets large, by central limit theorem, each row of $I - A(t)^*A(t)$ consists of approximately Gaussian random variables with variance $\frac{n}{m_x} = \frac{1}{\rho_x}$. Similarly, the components of $B(t)^*w_y$ are i.i.d. Gaussian with zero mean and approximate variance σ_y^2 and each row of $I - B(t)^*B(t)$ converges to independent zero mean Gaussian variables with variance $\frac{n}{m_y} = \frac{1}{\rho_y}$. Hence, the components of $A(t)^*w_x + (I - A(t)^*A(t))(x^t - \mathbf{x}_0)$ are approximately Gaussian with variance

$$\tau_x^t = \sigma_x^2 + \frac{1}{\rho_x} \frac{\|x^t - \mathbf{x}_0\|_2^2}{n}. \quad (5.30)$$

At $t = 0$, with the initialization $x^0 = 0$, one obtains that

$$\tau_x^0 = \sigma_x^2 + \frac{1}{\rho_x} \frac{\|\mathbf{x}_0\|_2^2}{n} \rightarrow \sigma_x^2 + \frac{1}{\rho_x} \mathbb{E}(X^2),$$

which is compatible with the SE initialization. A similar derivation gives $\tau_y^0 = \sigma_y^2 + \frac{1}{\rho_y} \mathbb{E}(Y^2)$. Moreover, by induction on t , one can simply check that at iteration $t + 1$,

$$x^{t+1} = \eta_t^x (X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y).$$

Thus, replacing in Equation (5.30) and using a similar argument, one obtains that for the iteration $t + 1$,

$$\frac{\|x^{t+1} - \mathbf{x}_0\|_2^2}{n} \rightarrow \mathbb{E}(X - \eta_t^x (X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y))^2,$$

which implies that at iteration $t + 1$:

$$\tau_x^{t+1} = \sigma_x^2 + \frac{1}{\rho_x} \mathbb{E}(X - \eta_t^x (X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y))^2.$$

A similar argument gives the corresponding equation for τ_y^t :

$$\tau_y^{t+1} = \sigma_y^2 + \frac{1}{\rho_y} \mathbb{E}(Y - \eta_t^y (X + \sqrt{\tau_x^t} Z_x, Y + \sqrt{\tau_y^t} Z_y))^2.$$

5.6.4 MMSE Estimator Linearly Correlated Bernoulli-Gaussian Signals

Suppose Z^k are independent Bernoulli-Gaussian random variables with probability distribution $p_i(z) = (1 - \alpha_i)\delta_0(z) + \alpha_i\mathbf{N}(0, \frac{1}{\alpha_i}, z)$. Let A be a $t \times k$ matrix and let $S = AZ^k$ be a t dimensional linearly correlated signal. Suppose $O = S + \tilde{N}$ is the observation vector, where \tilde{N} is a $t \times 1$ zero mean Gaussian measurement noise with a covariance matrix $\tilde{\Sigma}$. We denote by $\eta_i(x) = \mathbb{E}(S_i|O = x)$ the MMSE estimator of S_i , the i -th component of the signal, given a $t \times 1$ observation vector $O = x$. We will compute $\eta_1(x)$. The other estimators can be computed similarly.

It is easy to check that one can represent Z_i , $i \in [k]$ by $\Theta_i N_i$, where Θ_i^k are independent binary random variables with $\mathbb{P}(\Theta_i = 1) = \alpha_i$ and N^k are independent zero mean Gaussian variables with variance $\frac{1}{\alpha_i}$. Assume that Σ is the covariance matrix of N^k with diagonal elements $\Sigma_{ii} = \frac{1}{\alpha_i}$ and zero elsewhere. Let a_1 denote the

first row of A and assume that for a given binary sequence θ^k and for an arbitrary $n \times k$ matrix B , $B(\theta^k)$ denotes an $n \times k$ matrix whose i -th column is the i -th column of B provided $\theta_i = 1$ and zero otherwise.

Using the conditioning on Θ^k , we have


$$\eta_1(x) = \sum_{\theta^k \in \{0,1\}^k} \mathbb{E}(S_1 | O = x, \Theta^k = \theta^k) \mathbb{P}(\theta^k | O = x).$$

Conditioned on θ^k , $S_1 = a_1(\theta^k)N$ is a zero mean Gaussian with variance $a_1(\theta^k)\Sigma a_1(\theta^k)^*$. The observation vector is also Gaussian with a zero mean and a covariance matrix $A(\theta^k)\Sigma A(\theta^k)^* + \tilde{\Sigma}$, thus the estimation of S_1 is reduced to a Gaussian estimation problem where the estimator is known to be a linear function of observation. Let $\hat{S}_1(\theta^k, x) = a(\theta^k)\Sigma A(\theta^k)(A(\theta^k)\Sigma A(\theta^k)^* + \tilde{\Sigma})^{-1}x$. It is easy to check that $\mathbb{E}(S_1 | O = x, \Theta^k = \theta^k) = \hat{S}_1(\theta^k, x)$. Therefore, one obtains

$$\begin{aligned} \eta_1(x) &= \sum_{\theta^k} \hat{S}_1(\theta^k, x) \mathbb{P}(\theta^k | O = x) = \frac{1}{p_o(x)} \sum_{\theta^k} \hat{S}_1(\theta^k, x) \mathbb{P}(\theta^k) p_o(x | \theta^k) \\ &= \frac{\sum_{\theta^k} \hat{S}_1(\theta^k, x) \mathbb{P}(\theta^k) \mathbf{N}(0, A(\theta^k)\Sigma A(\theta^k)^* + \tilde{\Sigma}, x)}{\sum_{\theta^k} \mathbb{P}(\theta^k) \mathbf{N}(0, A(\theta^k)\Sigma A(\theta^k)^* + \tilde{\Sigma}, x)}, \end{aligned}$$

where $\mathbf{N}(\mu, C, x) = \frac{1}{\sqrt{(2\pi)^n \det(C)}} \exp(-\frac{1}{2}(x - \mu)^* C^{-1}(x - \mu))$ denotes the Gaussian distribution with mean μ and covariance matrix C .

Entropy Power Inequality for Integer-valued Random Variables



In Chapter 2, we observed that in order to prove the absorption phenomenon for integer-valued random variables, it is sufficient to find a lower bound for the gap between the conditional entropy of sum and the individual conditional entropy of a pair of random variables in terms of their individual conditional entropy¹. To be more precise, suppose that (X, Y) are integer-valued random variables with a given conditional entropy $H(X|Y) = c$ for some $c \in \mathbb{R}_+$ and let (X', Y') be an independent copy of (X, Y) . We needed to show that there is a universal function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any pair of integer-valued random variables (X, Y) and its independent copy (X', Y') ,

$$H(X + X'|Y, Y') - H(X|Y) \geq g(H(X|Y)). \quad (\text{A.1})$$

We also required that g be increasing and strictly positive except in the origin where $g(0) = 0$. It is important to emphasize that for the proof of the absorption phenomenon, g must be a universal function not depending on a specific pair (X, Y) but their conditional entropy $H(X|Y)$.

It is interesting to know that finding universal functions like g as in Equation (A.1) that establish universal bounds or inequalities between information measures is vastly studied in information theory, with the first result proposed by Shannon himself generally known as *Entropy Power Inequality* (EPI) [29]. EPI yields lower bounds on the differential entropy of the sum of two independent real-valued random variables in terms of the individual entropies. Versions of the EPI for discrete random variables have been obtained for special families of distributions with the differential entropy replaced by the discrete entropy, but no universal inequality is known (beyond trivial ones). More recently, the sunset theory for the entropy function provides a sharp inequality $H(X + X') - H(X) \geq \frac{1}{2} - o(1)$ when X, X' are i.i.d. with high entropy [72]. We strengthen this result by finding a universal lower bounds which holds for all range of values of $H(X)$ not necessarily large ones. Moreover, we extend the result to non-identically distributed random variables and to conditional entropies.

¹This chapter is the result of collaboration with Emmanuel Abbe and Emre Telatar. I am also grateful to Christophe Vignat for helpful discussions on this problem.

The structure of this appendix is as follows. In Section A.1, we overview a history of the EPI and recent results and extensions. In Section A.2, we state the main results that we obtained and give further intuitions. Section A.3, explains the proof techniques. Finally, in Section A.4, we explain some open problems and state a conjecture under which one can further strengthen the proven lower bounds.

A.1 History and Introduction

For a continuous random variable² \mathbf{X} on \mathbb{R}^n , let $h(\mathbf{X})$ be the differential entropy of \mathbf{X} and let $N(\mathbf{X}) = 2^{\frac{2}{n}h(\mathbf{X})}$ denote the *entropy power* of \mathbf{X} . If \mathbf{Y} is another continuous \mathbb{R}^n -valued random variable independent of \mathbf{X} , the EPI states that

$$N(\mathbf{X} + \mathbf{Y}) \geq N(\mathbf{X}) + N(\mathbf{Y}), \quad (\text{A.2})$$

with equality if and only if \mathbf{X} and \mathbf{Y} are Gaussian with proportional covariance matrices. A weaker statement of the EPI, yet of key use in applications, is the following inequality stated here for $n = 1$,

$$h(X + X') - h(X) \geq \frac{1}{2}, \quad (\text{A.3})$$

where X, X' are i.i.d., and where equality holds if and only if X is Gaussian.

The EPI was first proposed by Shannon [29] who used a variational argument to show that Gaussian random variables \mathbf{X} and \mathbf{Y} with proportional covariance matrices and specified differential entropies constitute a stationary point for $h(\mathbf{X} + \mathbf{Y})$. However, this does not exclude saddle points and local minima. The first rigorous proof of the EPI was given by Stam [73] in 1959, using the De Bruijn's identity which connects the derivative of the entropy with Gaussian perturbation to the Fisher information. This proof was further simplified by Blachman [74]. Another proof was given by Lieb [75] based on an extension of Young's inequality.

While there is a wide variety of entropic inequalities, the EPI is the only general inequality in information theory giving a tight lower bound on the entropy of a sum of independent random variables by means of the individual entropies. It is used as a key ingredient to prove converse results in coding theorems [76–80].

There have been numerous extensions and simplifications of the EPI over the reals [81–91]. There have also been several attempts to obtain discrete versions of the EPI, using Shannon entropy. Of course, it is not clear what is meant by a discrete version of the EPI, since (A.2), (A.3) clearly do not hold verbatim for Shannon entropy.

Several extensions have yet been developed. First, there have been some extensions using finite field additions, for example, the so-called Mrs. Gerber's Lemma (MGL) proved in [92] by Wyner and Ziv for binary alphabets. The MGL was further extended by Witsenhausen [93] to non-binary alphabets, who also provided counter-examples for the case of general alphabets. More recently, [94] obtained EPI and MGL results for abelian groups of order 2^n . Second, concerning discrete random variables and addition over the reals, Harremoës and Vignat [95] proved that the discrete EPI holds for binomial random variables with parameter $\frac{1}{2}$, which later on was generalized by

²All continuous random variables are assumed to have well-defined differential entropies.

Sharma, Das and Muthukrishnan [96]. Yu and Johnson [97] obtained a version of the EPI for discrete random variables using the notion of thinning.

More recently, Tao established in [72] a sumset theory for Shannon entropy, obtaining in particular the sharp inequality

$$H(X + X') - H(X) \geq \frac{1}{2} - o(1),$$

where $o(1)$ vanishes when $H(X)$ tends to infinity. Further results were obtained for the differential entropy in [98].

We are mostly interested in integer-valued random variables with arithmetic over the reals. We show that there exists an increasing function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, such that $g(c) = 0$ if and only if $c = 0$, and

$$H(X + X') - H(X) \geq g(H(X)),$$

for any i.i.d. integer-valued random variables X, X' . Although we have provided an explicit characterization of g , we found that even proving the existence of such a function (without explicit characterization) is equally challenging. We further generalize the result to non-identically distributed random variables and to conditional entropies.

To briefly overview the notations: the set of integers and reals will be denoted by \mathbb{Z} and \mathbb{R} . Similarly, \mathbb{Z}_+ and \mathbb{R}_+ will denote the set of positive integers and positive reals. We will use capital letters for random variables and lower case letters for their realizations (the random variable X can have realization x). The natural logarithm and the logarithm in base 2 will be denoted by \ln and \log_2 respectively, and for $x \in [0, 1]$, $h_2(x) = -x \log_2(x) - (1 - x) \log_2(1 - x)$ will denote the binary entropy function with the convention that $0 \log_2(0) = 0$. The entropy of a discrete random variable X in base 2 (bits) will be denoted by $H(X)$. We will interchangeably use $H(p)$ or $H(X)$, where p is the probability distribution of X . The conditional entropy of a random variable X given another random variable Y will be denoted by $H(X|Y)$. For $a, b \in \mathbb{R}$, we will use $a \vee b$ and $a \wedge b$ for the maximum and minimum of a and b . Also, $a^+ = a \vee 0$ denotes the positive part of a .

A.2 Statement of the Results

In this section, we will give an overview of the results. The first theorem gives a lower bound on the entropy gap of sum of two i.i.d. random variables as a function of their entropies.

Theorem A.1. *There is a function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any two i.i.d. integer-valued random variables X, X' ,*

$$H(X + X') - H(X) \geq g(H(X)). \tag{A.4}$$

Moreover, g is an increasing function, $\lim_{c \rightarrow \infty} g(c) = \frac{1}{8} \log_2(e)$ and $g(c) = 0$ if and only if $c = 0$.

Remark A.1. *The function g in Theorem A.1 is given by*

$$g(c) = \min_{x \in [0,1]} \left\{ (cx - h_2(x)) \vee \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e) \right\}.$$

Remark A.2. As we mentioned in the introduction, a recent result by Tao [72] implies that for i.i.d. integer-valued random variables X, X' of very high entropy, $H(X + X') - H(X) \approx \frac{1}{2}$. In comparison with this result, we get an asymptotic lower bound of $\frac{1}{8} \log_2(e) \approx 0.18$, which is also valid for the general independent case provided that the entropy of both random variables approaches infinity.

The next theorem extends the i.i.d. result to the general independent case.

Theorem A.2. There is a function $g : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ such that for any two independent integer-valued random variables X, X' ,

$$H(X + X') - \frac{H(X) + H(X')}{2} \geq g(H(X), H(X')).$$

Moreover, g is a positive and doubly-increasing³ function, $\lim_{(c,d) \rightarrow (\infty, \infty)} g(c, d) = \frac{1}{8} \log_2(e)$ and $g(c, d) = 0$ if and only if $c = d = 0$.

Remark A.3. One might be tempted to prove the stronger bound

$$H(X + X') - (H(X) \vee H(X')) \geq g(H(X), H(X')), \quad (\text{A.5})$$

for some doubly-increasing function g . However, this fails because, for example, assume that X, X' are random variables uniformly distributed over $\{1, 2, \dots, M\}$ and $\{1, 2, \dots, NM\}$, for some number $N \geq 2$. It is not difficult to show that

$$H(X + X') - (H(X) \vee H(X')) \leq \log_2\left(\frac{N+1}{N}\right),$$

which decreases to 0 with increasing N . Hence, the strong inequality (A.5) does not hold universally over all integer-valued random variables.

The next theorem extends the results in Theorem A.1 to the conditional case.

Theorem A.3. There is a function $\tilde{g} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any two i.i.d. integer-valued pairs of random variables (X, Y) and (X', Y') ,

$$H(X + X' | Y, Y') - H(X | Y) \geq \tilde{g}(H(X | Y)).$$

Moreover, $\tilde{g} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an increasing function and $\tilde{g}(c) = 0$ if and only if $c = 0$.

Remark A.4. The function \tilde{g} is given by

$$\tilde{g}(c) = \min_{\delta \in [0, \frac{1}{2}]} \left\{ (g(c, c) - h_2(\delta)) \vee \delta^2 g(c, c) \right\}, \quad (\text{A.6})$$

where g is as in Theorem A.2.

A.3 Proof Techniques

In this part, we will give an overview and also some intuition about the techniques used to prove the theorems. For clarity of the explanation, we have postponed some of the proofs to Section A.5.

³A function $g : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is doubly-increasing if for any value of one of the arguments, it is an increasing function of the other argument.

A.3.1 EPI for i.i.d. Random Variables

We will start from the EPI for i.i.d. random variables. The main idea of the proof is to find suitable bounds for $H(p \star p) - H(p)$ in two different cases: one case in which p is close to a spiky distribution (a unit mass at a single point) and the other case where p is close to a uniform distribution over a subset of \mathbb{Z} .

Lemma A.1. *Assume that p is a probability distribution over \mathbb{Z} with $H(p) = c$ and let $x = \|p\|_\infty$. Then,*

$$H(p \star p) - c \geq cx - h_2(x).$$

Proof. In Subsection A.5.1. □

Remark A.5. *Notice that Lemma A.1, gives a tight bound for spiky distributions for which $\|p\|_\infty$ is very close to 1, i.e., for $H(p) = c$, one gets $H(p \star p) - c \simeq c$, which is the best one can hope for.*

The next step is to give a bound for non-spiky distributions. The main idea is that in this case, it is possible to decompose the probability distribution p into two different parts p_1, p_2 with disjoint non-interlacing supports such that $p \star p_1$ and $p \star p_2$ are sufficiently far apart in ℓ_1 -distance. We formalize this through the following lemmas.

Lemma A.2. *Assume that p_1, p_2 and p are arbitrary probability distributions over \mathbb{Z} such that p_1 and p_2 have non-overlapping supports and $\|p\|_\infty = x$. Then*

$$\|p \star p_1 - p \star p_2\|_1 \geq 2(2x - 1)^+.$$

Proof. In Subsection A.5.1. □

Lemma A.3. *Let $c > 0$, $0 < \alpha \leq \frac{1}{2}$ and $n \in \mathbb{Z}$. Assume that p is a probability distribution over \mathbb{Z} such that $\alpha \leq p((-\infty, n]) \leq 1 - \alpha$ and $H(p) = c$. Then,*

$$\|p \star p_1 - p \star p_2\|_1 \geq 2\alpha,$$

where $p_1 = \frac{1}{p((-\infty, n])} p|_{(-\infty, n]}$ and $p_2 = \frac{1}{p([n+1, \infty))} p|_{[n+1, \infty)}$ are scaled restrictions of p to $(-\infty, n]$ and $[n+1, \infty)$ respectively.

Proof. In Subsection A.5.1. □

Lemma A.4. *Assuming the hypotheses of Lemma A.3,*

$$H(p \star p) - c \geq \frac{\alpha^2 \|p \star p_1 - p \star p_2\|_1^2}{2} \log_2(e).$$

Proof. In Subsection A.5.1. □

Lemma A.5. *Assume that p is a probability distribution over \mathbb{Z} with $H(p) = c$ and $\|p\|_\infty = x$. Then*

$$H(p \star p) - c \geq \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e).$$

Now that we have the required bounds in the spiky and non-spiky cases, we can combine them to prove Theorem A.1.

Proof of Theorem A.1. Assume that p is a probability distribution over \mathbb{Z} with $H(p) = c$ and $\|p\|_\infty = x$. It is easy to see that $x \geq 2^{-c}$. Using Lemma A.1 and Lemma A.5, it results that $H(p \star p) - c \geq l(c)$, where

$$l(c) = \min_{x \in [2^{-c}, 1]} \left\{ (cx - h_2(x)) \vee \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e) \right\}.$$

We will use a simpler lower bound given by

$$g(c) = \min_{x \in [0, 1]} \left\{ (cx - h_2(x)) \vee \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e) \right\},$$

where obviously $l(c) \geq g(c)$. It is easy to check that $g(c)$ is a continuous function of c . The monotonicity of g follows from monotonicity of $cx - h_2(x)$ with respect to c , for every $x \in [0, 1]$. For strict positivity, note that $(1-x)^2((1-x) \vee (4x-2)^+)^2$ is strictly positive for $x \in [0, 1)$ and it is 0 when $x = 1$, but $\lim_{x \rightarrow 1} cx - h_2(x) = c$. Hence, for $c > 0$, $g(c) > 0$. If $c = 0$ then

$$\begin{aligned} & \{(cx - h_2(x)) \vee \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e)\} \\ &= \frac{(1-x)^2((1-x) \vee (4x-2)^+)^2}{8} \log_2(e), \end{aligned}$$

and its minimum over $[0, 1]$ is 0. For asymptotic behavior, notice that at $x = 0$, $cx - h_2(x) = 0$ and

$$\frac{(1-x)^2((1-x) \vee (4x-2)^+)}{8} \log_2(e) = \frac{1}{8} \log_2(e).$$

Hence, from continuity, it results that $g(c) \leq \frac{1}{8} \log_2(e)$ for any $c \geq 0$. Also for any $0 < \epsilon < \frac{1}{2}$ there exists a c_0 such that for every $c > c_0$ and every x , $\epsilon < x \leq 1$, $cx - h_2(x) \geq \frac{1}{8} \log_2(e)$. Thus for any $\epsilon > 0$ there is a c_0 such that for $c > c_0$, the outer minimum over x in the definition of $g(c)$ is achieved on $[0, \epsilon]$, which is higher than $\frac{(1-\epsilon)^4}{8} \log_2(e)$. This implies that for every $\epsilon > 0$,

$$\frac{1}{8} \log_2(e) \geq \limsup_{c \rightarrow \infty} g(c) \geq \liminf_{c \rightarrow \infty} g(c) \geq \frac{(1-\epsilon)^4}{8} \log_2(e),$$

thus $\lim_{c \rightarrow \infty} g(c) = \frac{1}{8} \log_2(e)$. □

Figure A.1 shows the EPI gap. As expected, the asymptotic value of the gap is $\frac{1}{8} \log_2(e) \approx 0.18$.

A.3.2 EPI for non-i.i.d. random variables

Theorem A.2 is an extension of Theorem A.1 to independent but non-identically distributed random variables. Similar to the i.i.d. case the idea is to distinguish between the spiky and non-spiky distributions.




Figure A.1 – EPI gap (A.4) for i.i.d. integer-valued random variables

Lemma A.6. *Assume that p and q are two probability distributions over \mathbb{Z} with $H(p) = c$ and $H(q) = d$. Suppose that $x = \|p\|_\infty$ and $y = \|q\|_\infty$. Then,*

$$2H(p \star q) - c - d \geq dx - h_2(x) + cy - h_2(y). \quad (\text{A.7})$$

Proof. In Subsection A.5.2. □

When at least one of the distributions is spiky, Lemma A.6 gives a relatively tight bound. Hence, we should try to find a good bound for the non-spiky case.

Lemma A.7. *Let p, q be two probability distributions over \mathbb{Z} . Assume that there are $0 < \alpha, \beta < \frac{1}{2}$ and $m, n \in \mathbb{Z}$ such that $\alpha \leq p((-\infty, m]) \leq 1 - \alpha$ and $\beta \leq q((-\infty, n]) \leq 1 - \beta$. Then*

$$\|q \star p_1 - q \star p_2\|_1 + \|p \star q_1 - p \star q_2\|_1 \geq 2(\alpha + \beta),$$

where $p_1 = \frac{1}{p((-\infty, m])} p|_{(-\infty, m]}$, $p_2 = \frac{1}{p([m+1, \infty))} p|_{[m+1, \infty)}$, $q_1 = \frac{1}{q((-\infty, n])} q|_{(-\infty, n]}$, and $q_2 = \frac{1}{q([n+1, \infty))} q|_{[n+1, \infty)}$.

Proof. In Subsection A.5.2. □

Lemma A.8. *Assume that the hypotheses of Lemma A.7 hold and let $H(p) = c$ and $H(q) = d$. Then*

$$\begin{aligned} H(p \star q) - d &\geq \frac{\alpha^2 \|q \star p_1 - q \star p_2\|_1^2}{2} \log_2(e), \\ H(p \star q) - c &\geq \frac{\beta^2 \|p \star q_1 - p \star q_2\|_1^2}{2} \log_2(e). \end{aligned}$$

Proof. In Subsection A.5.2. □

Lemma A.9. *Let p and q be probability distributions over \mathbb{Z} with $H(p) = c$, $H(q) = d$, $\|p\|_\infty = x$ and $\|q\|_\infty = y$. Then $2H(p \star q) - c - d \geq l(x, y)$, where*

$$l(x, y) = \min_{(a,b) \in T(x,y)} \frac{(1-x)^2 a^2 + (1-y)^2 b^2}{8} \log_2(e),$$

and $T(x, y)$ is a subset of $(a, b) \in \mathbb{R}_+^2$ parameterized by $(x, y) \in [0, 1] \times [0, 1]$ and given by the following inequalities

$$a \geq (4y - 2)^+, b \geq (4x - 2)^+, a + b \geq 2 - x - y.$$

Moreover, $l(x, y)$ is a continuous function of (x, y) , $l(x, y) \geq 0$ and $l(x, y) = 0$ if and only if $(x, y) = (1, 1)$.

Proof. In Subsection A.5.2. □

Proof of Theorem A.2. Let $x = \|p\|_\infty$ and $y = \|q\|_\infty$. It is easy to check that $x \geq 2^{-c}$, $y \geq 2^{-d}$. Using Lemma A.6 and Lemma A.9, we obtain that $H(p \star q) - \frac{c+d}{2} \geq s(c, d)$, where $s(c, d)$ is given by

$$\frac{1}{2} \min_{(x,y) \in R(c,d)} \{(dx - h_2(x) + cy - h_2(y)) \vee l(x, y)\},$$

for $R(c, d) = [2^{-c}, 1] \times [2^{-d}, 1]$. We will use a simpler lower bound given by

$$g(c, d) = \frac{1}{2} \min_{(x,y) \in R} \{(dx - h_2(x) + cy - h_2(y)) \vee l(x, y)\},$$

where $R = [0, 1] \times [0, 1]$. It is easy to see that $g(c, d)$ is a continuous function. It is also a doubly-increasing function of its arguments. To prove the last part, notice that the $l(x, y)$ in the definition of g is strictly positive except for $(x^*, y^*) = (1, 1)$. But $\lim_{(x,y) \rightarrow (1,1)} dx - h_2(x) + cy - h_2(y) = c + d$, which is strictly positive unless $c = d = 0$. Therefore, for $(c, d) \neq (0, 0)$, $g(c, d) > 0$.

The function $dx - h_2(x) + cy - h_2(y)$ is a doubly-increasing function of (c, d) over R , which implies that $g(c, d)$ must be a doubly-increasing function of (c, d) . Also, using an argument similar to what we had in the proof of Theorem A.1, it is possible to show that for high values of c and d , the outer minimum in the definition of g is achieved in a neighborhood of $(0, 0)$, namely, $[0, \epsilon] \times [0, \epsilon]$, where ϵ goes to zero as c, d approach infinity. From the continuity of $l(x, y)$, it can be shown that in this range, the value of $l(x, y)$ is very close to

$$\min_{(a,b): a, b \geq 0, a+b \geq 2} \frac{a^2 + b^2}{8} \log_2(e) = \frac{1}{4} \log_2(e).$$

This implies that $\lim_{(c,d) \rightarrow (\infty, \infty)} g(c, d) = \frac{1}{8} \log_2(e)$. □

A.3.3 Conditional EPI

In this part, we will prove the EPI result for the conditional case, namely, we will find a lower bound for the conditional entropy gap, $H(X + X'|Y, Y') - H(X|Y)$, for i.i.d. integer-valued pairs (X, Y) and (X', Y') assuming that $H(X|Y) = c$, for some

positive number c . Notice that as Y and Y' only appear in the conditioning, we do not lose generality by assuming them to be integer-valued. Let us denote the probability distribution of Y by q , then the conditional entropy gap can be written as follows

$$\sum_{i,j \in \mathbb{Z}} q_i q_j H(p_i \star p_j) - c,$$

where p_i is the conditional distribution of X given $Y = i$.

Notice that we are interested to the infimum of this gap over all possible q, p_i satisfying $\sum_{i \in \mathbb{Z}} q_i H(p_i) = c$. Even if the minimizing q exists, it may not be finitely supported and in general, finding the corresponding gap requires an infinite dimensional constrained optimization.

To cope with this problem, we will show that it is possible to restrict the support size of q to 2 provided that instead of the i.i.d. case we consider the general independent case. Of course, at the end we get a looser bound at the price of simplifying the problem.

To be more specific, let (X, Y) and (X', Y') be independent (not necessarily i.i.d.) integer-valued pairs with $H(X|Y) = H(X'|Y') = c$ and let $t_n(c)$ be the infimum of $H(X + X'|Y, Y') - c$ over all $(X, Y), (X', Y')$ having a conditional entropy equal to c with Y and Y' having a support size at most n . Also, assume that $t_\infty(c)$ is the corresponding infimum when there is no constraint on the support size. We first prove the following lemma.

Lemma A.10. *For every $n \geq 2$, $t_\infty(c) = t_n(c)$.*

Proof. Obviously, $t_n(c) \geq t_\infty(c)$. Moreover, given any $\epsilon > 0$, there is an ϵ -optimal independent pair (X, Y) and (X', Y') such that

$$H(X + X'|Y, Y') - c \leq t_\infty(c) + \epsilon.$$

Let q, q' denote the distribution of Y, Y' and let p_i, p'_j be the conditional distribution of X, X' given $Y = i, Y' = j$. Let

$$V = \{\mathbf{v}_{ij} \in \mathbb{R}^3 : \mathbf{v}_{ij} = (H(p_i \star p'_j), H(p_i), H(p'_j)), i, j \in \mathbb{Z}\}.$$

It is easy to see that

$$\sum_{i,j \in \mathbb{Z}} q_i q'_j \mathbf{v}_{ij} = (H(X + X'|Y, Y'), c, c) = \mathbf{h},$$

which implies that the three dimensional vector $\mathbf{h} = (H(X + X'|Y, Y'), c, c)$ can be written as a convex combination of the vectors $\mathbf{v}_{ij} \in V$ with weights $q_i q'_j$. Let $\mathbf{v}_i = \sum_j q'_j \mathbf{v}_{ij}$. Then, we have $\sum_i q_i \mathbf{v}_i = \mathbf{h}$. Notice that the second component of \mathbf{v}_i is equal to $H(p_i)$. Also, the third component is equal to c independent of i , which implies that there are only two components depending on i in \mathbf{v}_i . Therefore, by Carathéodory theorem, it is possible to write \mathbf{h} as a convex combination of at most three $\mathbf{v}_i, i \in \mathbb{Z}$, which without loss of generality, we can assume to be $\{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2\}$. In other words, there are positive $\gamma_i, i = 0, 1, 2$, $\sum_{i=0}^2 \gamma_i = 1$ and $\mathbf{h} = \sum_{i=0}^2 \gamma_i \mathbf{v}_i$. Also, note that if we change the distribution of Y from q to γ , the resulting $(X, Y), (X', Y')$ is again an ϵ -optimal solution. Now, we claim that we can simplify the problem

further and find a probability triple $\psi = (\psi_0, \psi_1, \psi_2)$ with at most 2 non-zero elements such that $\sum_{i=0}^2 \psi_i H(p_i) = c$ and at the same time

$$\sum_{i=0}^2 \psi_i \mathbf{v}_i^{(1)} \leq \sum_{i=0}^2 \gamma_i \mathbf{v}_i^{(1)} = \sum_{i=0}^2 q_i \mathbf{v}_i^{(1)} = H(X + X'|Y, Y'),$$

where $\mathbf{v}_i^{(1)}$ denotes the first coordinate of the vector \mathbf{v}_i . This implies that if we replace the distribution γ for Y by ψ , which has a support size at most 2, we get a lower $H(X + X'|Y, Y')$. To prove the claim, let us consider the following optimization problem

$$\text{minimize } \sum_{i=0}^2 \psi_i \mathbf{v}_i^{(1)} \text{ s.t. } \begin{cases} \sum_{i=0}^2 \psi_i = 1, \\ \sum_{i=0}^2 \psi_i H(p_i) = c, \\ \psi_i \geq 0. \end{cases}$$

First of all, notice that as $\sum_{i=0}^2 \gamma_i H(p_i) = c$, γ is a feasible point. Therefore, the feasible set is a non-empty subset of the three dimensional probability simplex. Also, as the objective function is linear in ψ , the optimal point must be an extremal point (boundary point) of the feasible set, which implies that there is an optimal solution with at most two non-zero components and this proves the claim.

By symmetry, we can apply the same argument to the probability distribution q' of Y' to get an ϵ -optimal solution in which the support of both q and q' has at most size 2. Hence, this implies that for any $\epsilon > 0$ and any $n \geq 2$, $t_n(c) \leq t_2(c) \leq t_\infty(c) + \epsilon$, thus $t_n(c) = t_\infty(c)$. \square

Lemma A.10 allows us to simplify finding the lower bound. However, we might get a looser bound because we relaxed the condition that (X, Y) and (X', Y') be identically distributed. From now on, we will assume that Y and Y' are binary random variables. We will use the following two lemmas to get a lower bound for the conditional entropy gap.

Lemma A.11. *Let $(X, Y), (X', Y')$ be an independent pair of random variables, $\mathbb{P}(Y = 0) = \alpha$, $\mathbb{P}(Y' = 0) = \beta$ and $H(X|Y) = H(X'|Y') = c$. Then,*

$$H(X + X'|Y, Y') - c \geq g(c, c) - (h_2(\alpha) \wedge h_2(\beta)),$$

where g is the same function as in Theorem A.2.

Proof. In Subsection A.5.3. \square

Lemma A.12. *Assume that all of the conditions of Lemma A.11 hold. Suppose there is a $0 \leq \delta \leq \frac{1}{2}$ such that $\delta < \alpha, \beta < 1 - \delta$. Then,*

$$H(X + X'|Y, Y') - c \geq \delta^2 g(c, c).$$

Proof. In Subsection A.5.3. \square

Proof of Theorem A.3. The proof follows by combining the results obtained in Lemma A.11 and A.12. Let $\delta = \min\{\alpha, 1 - \alpha, \beta, 1 - \beta\}$. Then $0 \leq \delta \leq \frac{1}{2}$ and using Lemma A.12, we get the lower bound $\delta^2 g(c, c)$. Similarly, from Lemma A.11 and

using the fact that $h_2(\alpha) \wedge h_2(\beta) = h_2(\delta)$, we get the lower bound $g(c, c) - h_2(\delta)$. Combining the two, we obtain the desired lower bound

$$\tilde{g}(c) = \min_{\delta \in [0, \frac{1}{2}]} \{(g(c, c) - h_2(\delta)) \vee \delta^2 g(c, c)\}.$$

The monotonicity of \tilde{g} follows from the monotonicity of $g(c, c)$. Also, notice that $\delta^2 g(c, c)$ is strictly positive unless $\delta = 0$ but $\lim_{\delta \rightarrow 0} g(c, c) - h_2(\delta) = g(c, c)$, which is strictly positive if $c > 0$. Therefore, for $c > 0$ we have $\tilde{g}(c) > 0$. This completes the proof. \square

A.4 Open problems

A.4.1 Closure convexity of the entropy set \mathcal{H}

As we saw in the proof of Theorem A.3, the conditional EPI does not directly follow from the unconditional one. In particular, we had to relax the i.i.d. condition in order to get a relatively weak lower bound. In this part, we propose another approach to the problem which uses the closure convexity of the entropy set as we will define in a moment.

Definition A.1. *The entropy set \mathcal{H} is defined as follows*

$$\mathcal{H} = \{(H(p \star q), H(p), H(q)) \in \mathbb{R}_+^3 : p, q \text{ are probability distributions over } \mathbb{Z}\}.$$

Remark A.6. *Notice that multiple (p, q) pairs may be mapped to the same point in \mathcal{H} space. For example, if (p, q) is mapped to a point $\mathbf{v} \in \mathcal{H}$, then any distribution (\tilde{p}, \tilde{q}) in which \tilde{p} and \tilde{q} are shifted versions of p and q is also mapped to \mathbf{v} .*

Remark A.7. *Some of the boundaries of the set \mathcal{H} trivially follow from the properties of the entropy, i.e., for any $\mathbf{v} \in \mathcal{H}$,*

$$\mathbf{v}^{(1)} \geq \mathbf{v}^{(2)}, \mathbf{v}^{(1)} \geq \mathbf{v}^{(3)}, \mathbf{v}^{(1)} \leq \mathbf{v}^{(2)} + \mathbf{v}^{(3)},$$

where $\mathbf{v}^{(i)}$ denotes the i -th coordinate of the vector \mathbf{v} . Also the boundary $\mathbf{v}^{(1)} = \mathbf{v}^{(2)} + \mathbf{v}^{(3)}$ is achievable. To show this, let $\mathbf{v}^{(2)}, \mathbf{v}^{(3)} \in \mathbb{R}_+$ and consider two finite support distributions p and q of support $\{0, 1, \dots, M-1\}$ and $\{0, 1, \dots, N-1\}$ for appropriate M and N such that $H(p) = \mathbf{v}^{(2)}$ and $H(q) = \mathbf{v}^{(3)}$. Now, fix p and define a new distribution \tilde{q} as follows

$$\tilde{q}(i) = \begin{cases} 0 & \frac{i}{M} \notin \mathbb{Z}, \\ q(\frac{i}{M}) & \frac{i}{M} \in \mathbb{Z}. \end{cases}$$

It is not difficult to show that $H(\tilde{q}) = H(q) = \mathbf{v}^{(3)}$ and $H(p \star \tilde{q}) = H(p) + H(\tilde{q}) = \mathbf{v}^{(2)} + \mathbf{v}^{(3)}$.

We propose the following conjecture about the set \mathcal{H} .

Conjecture A.1. *The closure of the set \mathcal{H} is convex.*

Using this conjecture, we can prove the following lemma, which is a stronger version of the conditional EPI.

Theorem A.4. *Assume that Conjecture A.1 holds. Let (X, Y) and (X', Y') be independent pairs of integer-valued random variables with $H(X|Y) = c$, $H(X'|Y') = d$. Then*

$$H(X + X'|Y, Y') - \frac{c + d}{2} \geq g(c, d),$$

where g is the same function as in Theorem A.2.

Proof. Let us assume that the distribution of Y, Y' is q, q' respectively. Also assume that p_i, p'_j is the distribution of X, X' when $Y = i, Y' = j$. Let

$$\mathbf{v}_{ij} = (H(p_i \star p'_j), H(p_i), H(p'_j)), \quad i, j \in \mathbb{Z}.$$

Notice that $\mathbf{v}_{ij} \in \mathcal{H}$. We also have

$$(H(X + X'|Y, Y'), c, d) = \sum_{i, j \in \mathbb{Z}} q_i q'_j \mathbf{v}_{ij},$$

which is a convex combination of the vectors \mathbf{v}_{ij} . By the closure convexity of \mathcal{H} , for any $\epsilon > 0$, it is possible to find an $\mathbf{h} \in \mathcal{H}$ in ϵ -neighborhood of $(H(X + X'|Y, Y'), c, d)$. In other words, for the given $\epsilon > 0$, there are two distributions μ_1, μ_2 over \mathbb{Z} such that

$$\begin{aligned} H(\mu_1 \star \mu_2) - \epsilon &\leq H(X + X'|Y, Y') \leq H(\mu_1 \star \mu_2) + \epsilon, \\ H(\mu_1) - \epsilon &\leq c \leq H(\mu_1) + \epsilon, \\ H(\mu_2) - \epsilon &\leq d \leq H(\mu_2) + \epsilon. \end{aligned}$$

In particular, this implies that

$$\begin{aligned} H(X + X'|Y, Y') - \frac{c + d}{2} &\geq H(\mu_1 \star \mu_2) - \frac{c + d}{2} - \epsilon \\ &\geq H(\mu_1 \star \mu_2) - \frac{H(\mu_1) + H(\mu_2)}{2} - 2\epsilon \\ &\geq g(H(\mu_1), H(\mu_2)) - 2\epsilon \geq g(c - \epsilon, d - \epsilon) - 2\epsilon, \end{aligned}$$

where we used the monotonicity of g with respect to both arguments. As $\epsilon > 0$ is arbitrary and g is a continuous function, $H(X + X'|Y, Y') - \frac{c+d}{2} \geq g(c, d)$. \square

Remark A.8. *In the case that (X, Y) and (X', Y') are i.i.d. pairs with $H(X|Y) = H(X'|Y') = c$, this result reduces to*

$$H(X + X'|Y, Y') - c \geq g(c, c),$$

which is tighter than the bound (A.6) obtained in Theorem A.3.

A.5 Proof of Auxiliary Lemmas

A.5.1 EPI for i.i.d. random variables

Proof of Lemma A.1. Assume that X is an integer-valued random variable with probability distribution p . Let $i \in \mathbb{Z}$ be such that $p(i) = \|p\|_\infty = x$. Let p_i be the probability distribution p shifted by i , i.e., $p_i(k) = p(k + i)$ for every $k \in \mathbb{Z}$. Assume

that $P = p_i$. Note that $H(p \star p) = H(P \star P)$ and $H(P) = H(p) = c$. Hence, without any loss of generality, one can assume that $p(0) = \|p\|_\infty$. Let B be a binary random variable with $\mathbb{P}\{B = 0\} = x = 1 - \mathbb{P}\{B = 1\}$, and let R be a random variable defined by $\mathbb{P}\{R = k\} = p_i(k)/(1 - x)$ for every $k \in \mathbb{Z} \setminus \{0\}$ and $\mathbb{P}\{R = 0\} = 0$. One can check that $X = BR$ for independent B and R . We also have $H(X) = h_2(x) + (1 - x)H(R)$. Let X' be an independent copy of X . Then, we have

$$\begin{aligned} H(p \star p) &= H(BR + X') \geq H(BR + X'|B) = xc + (1 - x)H(X' + R) \\ &\geq xc + (1 - x)H(R) = xc + c - h_2(x). \end{aligned}$$

This yields $H(p \star p) - c \geq xc - h_2(x)$. \square

Proof of Lemma A.2. Let $n_0 \in \mathbb{Z}$ be such that $p(n_0) = \|p\|_\infty = x$. We have the following:

$$\begin{aligned} \|p \star p_1 - p \star p_2\|_1 &= \sum_{i \in \mathbb{Z}} |p \star p_1(i) - p \star p_2(i)| = \sum_{i \in \mathbb{Z}} \left| \sum_{j \in \mathbb{Z}} p(j)(p_1(i - j) - p_2(i - j)) \right| \\ &\geq \sum_{i \in \mathbb{Z}} p(n_0) |p_1(i - n_0) - p_2(i - n_0)| - \sum_{i \in \mathbb{Z}} \sum_{j \neq n_0} p(j) |p_1(i - j) - p_2(i - j)| \\ &= x \|p_1 - p_2\|_1 - (1 - x) \|p_1 - p_2\|_1 = 2(2x - 1), \end{aligned}$$

where we used the fact that p_1 and p_2 have non-overlapping supports thus $\|p_1 - p_2\|_1 = \|p_1\|_1 + \|p_2\|_1 = 2$. Therefore, we get the desired result $\|p \star p_1 - p \star p_2\|_1 \geq 2(2x - 1)^+$. \square

Proof of Lemma A.3. Let $\alpha_1 = p((-\infty, n])$ and $\alpha_2 = p([n + 1, \infty)) = 1 - \alpha_1$. Note that $p = \alpha_1 p_1 + \alpha_2 p_2$. We distinguish two cases $\alpha_1 \leq \frac{1}{2}$ and $\alpha_1 > \frac{1}{2}$. If $\alpha_1 \leq \frac{1}{2}$ then we have

$$\begin{aligned} \|p \star p_1 - p \star p_2\| &= \|\alpha_1 p_1 \star p_1 - (1 - \alpha_1) p_2 \star p_2 + (1 - 2\alpha_1) p_1 \star p_2\|_1 \\ &\geq \|\alpha_1 p_1 \star p_1 - (1 - \alpha_1) p_2 \star p_2\|_1 - (1 - 2\alpha_1) \|p_1 \star p_2\|_1 \\ &= \alpha_1 + (1 - \alpha_1) - (1 - 2\alpha_1) = 2\alpha_1 \geq 2\alpha, \end{aligned}$$

whereas if $\alpha_1 > \frac{1}{2}$, we have

$$\begin{aligned} \|p \star p_1 - p \star p_2\| &= \|\alpha_1 p_1 \star p_1 - (1 - \alpha_1) p_2 \star p_2 + (1 - 2\alpha_1) p_1 \star p_2\|_1 \\ &\geq \|\alpha_1 p_1 \star p_1 - (1 - \alpha_1) p_2 \star p_2\|_1 - (2\alpha_1 - 1) \|p_1 \star p_2\|_1 \\ &= \alpha_1 + (1 - \alpha_1) - (2\alpha_1 - 1) = 2(1 - \alpha_1) \geq 2\alpha, \end{aligned}$$

where we used the triangle inequality, $1 - \alpha_1 \geq \alpha$ and the fact that $p_1 \star p_1$ and $p_2 \star p_2$ have non-overlapping supports, thus the ℓ_1 -norm of the sum is equal to sum of the corresponding ℓ_1 -norms. \square

Proof of Lemma A.4. Let α_1 and α_2 be the same as in the proof of Lemma A.3. Let $\nu_1 = p_1 \star p$, $\nu_2 = p_2 \star p$, and for $x \in [0, 1]$, define $\mu_x = x\nu_1 + (1 - x)\nu_2$ and $f(x) = H(\mu_x)$. We have

$$f'(x) = - \sum (\nu_1(i) - \nu_2(i)) \log_2(\mu_x(i)), \quad f''(x) = - \frac{1}{\ln(2)} \sum \frac{(\nu_1(i) - \nu_2(i))^2}{\mu_x(i)} \leq 0.$$

Therefore, $f(x)$ is a concave function of x . Moreover,

$$f'(0) = D(\nu_1\|\nu_2) + H(\nu_1) - H(\nu_2), \quad f'(1) = -D(\nu_2\|\nu_1) + H(\nu_1) - H(\nu_2).$$

Since p_1 and p_2 have different supports, there are i, j such that $\nu_{1i} = 0, \nu_{2j} > 0$ and $\nu_{1j} > 0, \nu_{2i} = 0$. Hence $D(\nu_1\|\nu_2)$ and $D(\nu_2\|\nu_1)$ are both equal to infinity. In other words, $f'(0) = +\infty, f'(1) = -\infty$.

Hence, the unique maximum of the function f must be between 0 and 1. Assume that for fixed ν_1 and ν_2 , x^* is the maximizer. If $0 < \alpha_1 \leq x^*$ then

$$\alpha_1 f'(\alpha_1) = \sum \alpha_1 (\nu_2(i) - \nu_1(i)) \log_2(\mu_{\alpha_1}(i)) \geq 0,$$

which implies that

$$\begin{aligned} f(\alpha_1) &= - \sum \mu_{\alpha_1}(i) \log_2(\mu_{\alpha_1}(i)) = - \sum (\nu_2(i) + \alpha_1(\nu_1(i) - \nu_2(i))) \log_2(\mu_{\alpha_1}(i)) \\ &\geq - \sum \nu_2(i) \log_2(\mu_{\alpha_1}(i)) = H(\nu_2) + D(\nu_2\|\mu_{\alpha_1}) \\ &\geq H(p) + \frac{1}{2 \ln(2)} \|\nu_2 - \mu_{\alpha_1}\|_1^2 = H(p) + \frac{\alpha_1^2}{2 \ln(2)} \|\nu_1 - \nu_2\|_1^2, \end{aligned}$$

where we used Pinsker's inequality for distributions r and s , $D(r\|s) \geq \frac{1}{2 \ln(2)} \|r - s\|_1^2$. Similarly, we can show that if $x^* \leq \alpha_1 \leq 1$ then

$$f(\alpha_1) \geq H(p) + \frac{(1 - \alpha_1)^2}{2 \ln(2)} \|\nu_1 - \nu_2\|_1^2.$$

As $\alpha \leq \alpha_1 \leq 1 - \alpha$ and $\alpha \leq \frac{1}{2}$ it results that

$$\begin{aligned} H(p \star p) &= H(\alpha_1 p \star p_1 + (1 - \alpha_1) p \star p_2) = f(\alpha_1) \geq H(p) + \frac{\alpha^2}{2 \ln(2)} \|\nu_1 - \nu_2\|_1^2 \\ &\geq c + \frac{\alpha^2}{2 \ln(2)} \|\nu_1 - \nu_2\|_1^2, \end{aligned}$$

which is the desired result. \square

Proof of Lemma A.5. Let $x = \|p\|_\infty$ and $\alpha = \frac{1-x}{2}$. It is easy to show that there is an $n \in \mathbb{Z}$ such that $\alpha \leq p((-\infty, n]) \leq 1 - \alpha$. Also, let p_1 and p_2 , as in Lemma A.3, be the restriction of p to $(-\infty, n]$ and $[n+1, \infty)$. As p_1 and p_2 have disjoint supports, using Lemma A.3 and A.2, it results that $\|p \star p_1 - p \star p_2\|_1 \geq (1-x) \vee (4x-2)^+$. Therefore, using Lemma A.4, we get

$$H(p \star p) - c \geq \frac{(1-x)^2 ((1-x) \vee (4x-2)^+)^2}{8} \log_2(e).$$

This completes the proof. \square

A.5.2 EPI for non-i.i.d. random variables

Proof of Lemma A.6. Let X and Y be two independent random variables with probability distribution p and q . Let $x = \|p\|_\infty$. Similar to the proof of Lemma A.1, it can be shown that there is a binary random variable B , $\mathbb{P}(B=0) = x$ and a random variable R independent of B such that $\tilde{X} = BR$, where \tilde{X} is a suitably

shifted version of X with $\mathbb{P}(\tilde{X} = 0) = x$. Also, $H(X) = h_2(x) + (1-x)H(R)$. Then, we get

$$\begin{aligned} H(p \star q) &= H(X + Y) = H(\tilde{X} + Y) = H(BR + Y) \geq H(BR + Y|B) \\ &\geq \mathbb{P}(B = 0)H(Y) + \mathbb{P}(B = 1)H(R + Y) \geq xd + (1-x)H(R) \\ &= xd + c - h_2(x), \end{aligned}$$

which implies that $H(p \star q) - c \geq xd - h_2(x)$. By symmetry, we also obtain that $H(p \star q) - d \geq yc - h_2(y)$. Combining the two, we get the desired result $2H(p \star q) - c - d \geq dx - h_2(x) + cy - h_2(y)$. \square

Proof of Lemma A.7. Let $\alpha_1 = p((-\infty, m])$, $\alpha_2 = 1 - \alpha_1$, $\beta_1 = q((-\infty, n])$ and $\beta_2 = 1 - \beta_1$. Note that $p = \alpha_1 p_1 + \alpha_2 p_2$ and $q = \beta_1 q_1 + \beta_2 q_2$. Thus we obtain

$$\begin{aligned} \|q \star p_1 - q \star p_2\|_1 + \|p \star q_1 - p \star q_2\|_1 &\geq \|q \star p_1 - q \star p_2 + p \star q_1 - p \star q_2\|_1 \\ &= \|(\alpha_1 + \beta_1)p_1 \star q_1 + (\beta_2 - \alpha_1)p_1 \star q_2 + (\alpha_2 - \beta_1)p_2 \star q_1 - (\alpha_2 + \beta_2)p_2 \star q_2\|_1 \\ &\geq \|(\alpha_1 + \beta_1)p_1 \star q_1 - (\alpha_2 + \beta_2)p_2 \star q_2\|_1 - \|(\beta_2 - \alpha_1)p_1 \star q_2 + (\alpha_2 - \beta_1)p_2 \star q_1\|_1 \\ &\geq \alpha_1 + \beta_1 + \alpha_2 + \beta_2 - |\beta_2 - \alpha_1| - |\alpha_2 - \beta_1| \\ &= 2(1 - |1 - (\alpha_1 + \beta_1)|), \end{aligned}$$

where we used the triangle inequality and the fact that $p_1 \star q_1$ and $p_2 \star q_2$ have non-overlapping supports. We consider two cases: if $\alpha_1 + \beta_1 \leq 1$ then $(1 - |1 - (\alpha_1 + \beta_1)|) = (\alpha_1 + \beta_1) \geq (\alpha + \beta)$. Otherwise, $\alpha_1 + \beta_1 > 1$ and we obtain

$$(1 - |1 - (\alpha_1 + \beta_1)|) = 2 - (\alpha_1 + \beta_1) = \alpha_2 + \beta_2 \geq \alpha + \beta.$$

Therefore, in both cases we get

$$\|q \star p_1 - q \star p_2\|_1 + \|p \star q_1 - p \star q_2\|_1 \geq 2(\alpha + \beta),$$

which is the desired result. \square

Proof of Lemma A.8. Let $\alpha_1 = p((-\infty, m])$, $\alpha_2 = 1 - \alpha_1$, $\nu_1 = p_1 \star q$, $\nu_2 = p_2 \star q$, and for $x \in [0, 1]$, let $\mu_x = x\nu_1 + (1-x)\nu_2$ and $f(x) = H(\mu_x)$. By an argument similar to what we had in the proof of Lemma A.4, we can show that

$$H(p \star q) = f(\alpha_1) \geq d + \frac{\alpha^2}{2 \ln(2)} \|\nu_1 - \nu_2\|_1^2,$$

which implies that

$$H(p \star q) - d \geq \frac{\alpha^2 \|q \star p_1 - q \star p_2\|_1^2}{2} \log_2(e).$$

The other inequality in the lemma follows by symmetry. \square

Proof of Lemma A.9. Let $\alpha = \frac{1-x}{2}$ and $\beta = \frac{1-y}{2}$. It can be checked that α and β satisfy the conditions of Lemma A.7 and A.8. Therefore, using Lemma A.8, we obtain

$$2H(p \star q) - c - d \geq \frac{\alpha^2 a^2 + \beta^2 b^2}{2} \log_2(e) = \frac{(1-x)^2 a^2 + (1-y)^2 b^2}{8} \log_2(e),$$

where $a = \|q \star p_1 - q \star p_2\|_1$ and $b = \|p \star q_1 - p \star q_2\|_1$. Also, from Lemma A.7, we have

$$a + b \geq 2(\alpha + \beta) = 2 - x - y. \quad (\text{A.8})$$

Furthermore, applying Lemma A.2 to the distribution p with $\|p\|_\infty = x$ and q_1, q_2 with disjoint supports, and similarly to q with $\|q\|_\infty = y$ and p_1, p_2 with disjoint supports, we get

$$b \geq (4x - 2)^+, a \geq (4y - 2)^+. \quad (\text{A.9})$$

Therefore, $2H(p \star q) - c - d \geq l(x, y)$, where

$$l(x, y) = \min_{(a,b) \in T(x,y)} \frac{(1-x)^2 a^2 + (1-y)^2 b^2}{8} \log_2(e),$$

and where $T(x, y)$ is defined by the three inequalities derived in (A.8) and (A.9). The continuity of $l(x, y)$ can be easily checked. For the last part of the lemma, notice that if $M = x \vee y < 1$ then it is not difficult to show that

$$l(x, y) \geq \min_{a,b \geq 0, a+b \geq 2-2M} \frac{(1-M)^2(a^2 + b^2)}{8} \log_2(e) \geq \frac{(1-M)^4}{4} \log_2(e) > 0,$$

which is strictly positive. Moreover, if $x \vee y = 1$ but $(x, y) \neq (1, 1)$ then, for example, $y \in [0, 1), x = 1$, which implies that $b \geq 2$. Therefore, we get $l(x, y) \geq \frac{(1-y)^2}{2} \log_2(e)$, which is strictly positive unless $y = 1$. A similar argument applies to $x \in [0, 1), y = 1$. Therefore, over $(x, y) \in [0, 1] \times [0, 1]$, $l(x, y) \geq 0$ and $l(x, y) = 0$ if and only if $(x, y) = (1, 1)$. \square

A.5.3 Conditional EPI

Proof of Lemma A.11. To prove the lemma, notice that we have the constraint $H(X|Y) = H(X'|Y') = c$ and the probability distribution of Y, Y' has a support of size 2. We first prove that it is possible to modify the conditional distribution of the random variables X and X' given Y and Y' in a way that none of the constraints are violated, $H(X + X'|Y, Y')$ remains fixed and simultaneously, $H(Y|X)$ and $H(Y'|X')$ become as small as desired. To show this, let $p_i, p'_j, i, j \in \{0, 1\}$ be the distribution of X, X' conditioned on $Y = i, Y' = j$. Notice that if we shift any p_i, p'_j to the right or to the left by as many steps as we want, the conditional entropies remain unchanged so does $H(X + X'|Y, Y')$. We claim that by suitable shift of distributions, it is possible to make $H(Y|X)$ as small as we want. The same is true for $H(Y'|X')$.

To prove the claim, let $\epsilon > 0$ and assume that A_ϵ and B_ϵ are subsets of \mathbb{Z} of minimal size such that $p_0(A_\epsilon) \geq 1 - \frac{\epsilon}{2}$ and $p_1(B_\epsilon) \geq 1 - \frac{\epsilon}{2}$. In particular, for any $i \in A_\epsilon, j \in B_\epsilon, p_0(i) > 0, p_1(j) > 0$. Moreover,

$$\mathbb{P}(X \in A_\epsilon \cup B_\epsilon) \geq \alpha p_0(A_\epsilon) + (1 - \alpha) p_1(B_\epsilon) \geq 1 - \frac{\epsilon}{2}.$$

For $n \in \mathbb{Z}_+$, let us define $B_\epsilon^{(n)} = \{i + n : i \in B_\epsilon\}$, to be the right shift of B_ϵ by n . Also assume that $p_1^{(n)}$ is the probability distribution shifted to the right by n , namely, for $k \in \mathbb{Z}, p_1^{(n)}(k) = p_1(k - n)$. Specially, this implies that $p_1^{(n)}(B_\epsilon^{(n)}) = p_1(B_\epsilon)$.

Now let us replace p_1 , by $p_1^{(n)}$ and let us denote the resulting random variable by \tilde{X} . This assumption does not change $H(X|Y)$ and $H(X + X'|Y, Y')$. As A_ϵ and B_ϵ are finite sets, there is N_1 such that for all $n > N_1$, the two sets A_ϵ and $B_\epsilon^{(n)}$ are disjoint. For $a \in A_\epsilon$ and $b \in B_\epsilon$, let us compute the conditional distribution of Y given $\tilde{X} = a$ and $\tilde{X} = b + n \in B_\epsilon^{(n)}$. We have

$$\begin{aligned} \mathbb{P}(Y = 0 | \tilde{X} = a) &= \frac{\alpha p_0(a)}{\alpha p_0(a) + (1 - \alpha)p_1(a - n)}, \\ \mathbb{P}(Y = 1 | \tilde{X} = b + n) &= \frac{(1 - \alpha)p_1(b)}{(1 - \alpha)p_1(b) + \alpha p_0(b + n)}. \end{aligned}$$

It is not difficult to see that for all $a \in A_\epsilon$ and all $b \in B_\epsilon$, both of these numbers converge to 1 as n goes to infinity, which implies that both $H(Y|\tilde{X} = a)$ and $H(Y|\tilde{X} = b)$ converge to 0. In particular, there is an N_2 such that for $n > N_2$ these two numbers are less than $\frac{\epsilon}{2}$. Therefore, for $n > \max\{N_1, N_2\}$ we have

$$\begin{aligned} H_n(Y|\tilde{X}) &= \sum_{k \in \mathbb{Z}} p_{\tilde{X}}(k) H(Y|\tilde{X} = k) \leq \sum_{k \in A_\epsilon \cup B_\epsilon^{(n)}} p_{\tilde{X}}(k) \times \frac{\epsilon}{2} + \sum_{k \notin A_\epsilon \cup B_\epsilon^{(n)}} p_{\tilde{X}}(k) \times 1 \\ &= \sum_{k \in A_\epsilon \cup B_\epsilon} p_X(k) \times \frac{\epsilon}{2} + \sum_{k \notin A_\epsilon \cup B_\epsilon} p_X(k) \leq \epsilon, \end{aligned}$$

which proves the claim. Now assume that we have selected $(X, Y), (X', Y')$ such that $H(Y|X), H(Y'|X') < \epsilon$ for some positive small number ϵ . Then we have

$$\begin{aligned} H(X + X'|Y, Y') - c &= H(X + X') - H(X) - I(X + X'; Y, Y') + I(X; Y) \\ &\geq H(X + X') - H(X) - H(Y, Y') + H(Y) - H(Y|X) \\ &\geq H(X + X') - H(X) - H(Y') - \epsilon \\ &\geq g(H(X), H(X')) - h_2(\beta) - \epsilon \\ &\geq g(c, c) - h_2(\beta) - \epsilon, \end{aligned}$$

where we used the independence of Y, Y' , doubly-increasing property of g and the fact that $H(X) \geq H(X|Y) = c$ and similarly $H(X') \geq c$. As this is true for any $\epsilon > 0$, we obtain $H(X + X'|Y, Y') - c \geq g(c, c) - h_2(\beta)$.

By symmetry, we also have $H(X + X'|Y, Y') - c \geq g(c, c) - h_2(\alpha)$. Therefore, we get the desired result $H(X + X'|Y, Y') - c \geq g(c, c) - (h_2(\alpha) \wedge h_2(\beta))$. \square

Proof of Lemma A.12. Assume that the distribution of Y, Y' is q, q' . Also, for $k, l \in \{0, 1\}$, let p_k, p'_l be the conditional distribution of X, X' given $Y = k, Y' = l$. By the assumption, $\delta < \alpha, \beta < 1 - \delta$, we have $q_r, q'_s > \delta$ for any $r, s \in \{0, 1\}$, thus there must be $i, j \in \{0, 1\}$ such that $H(p_i), H(p'_j) \geq c$. Therefore, we have

$$\begin{aligned} H(X + X'|Y, Y') - c &= \sum_{k, l=0}^1 q_k q'_l (H(p_k \star p'_l) - \frac{H(p_k) + H(p'_l)}{2}) \\ &\geq q_i q'_j (H(p_i \star p'_j) - \frac{H(p_i) + H(p'_j)}{2}) \\ &\geq \delta^2 g(H(p_i), H(p'_j)) \geq \delta^2 g(c, c), \end{aligned}$$

where we used the doubly-increasing property of g . \square

Bibliography

- [1] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [2] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [3] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [4] ———, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [5] (2014, Mar.) Inverse problem. [Online]. Available: http://en.wikipedia.org/wiki/Inverse_problem
- [6] J. Pearl, *Causality: models, reasoning and inference*. Cambridge University Press, 2000, vol. 29.
- [7] S. L. Lauritzen, *Graphical models*. Oxford University Press, 1996.
- [8] M. J. Wainwright and M. I. Jordan, “Graphical models, exponential families, and variational inference,” *Foundations and Trends® in Machine Learning*, vol. 1, no. 1-2, pp. 1–305, 2008.
- [9] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, “Model-based compressive sensing,” *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.
- [10] V. Cevher, P. Indyk, C. Hegde, and R. G. Baraniuk, “Recovery of clustered sparse signals from compressive measurements,” DTIC Document, Tech. Rep., 2009.
- [11] V. Cevher, P. Indyk, L. Carin, and R. G. Baraniuk, “Sparse signal recovery and acquisition with graphical models,” *IEEE Signal Processing Magazine*, vol. 27, no. 6, pp. 92–103, 2010.
- [12] V. Cevher, M. F. Duarte, C. Hegde, and R. G. Baraniuk, “Sparse signal recovery using markov random fields,” in *NIPS*, vol. 8, 2008, pp. 257–264.

- [13] S. Jafarpour, W. Xu, B. Hassibi, and R. Calderbank, "Efficient and robust compressed sensing using optimized expander graphs," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 4299–4308, 2009.
- [14] L. Applebaum, S. D. Howard, S. Searle, and R. Calderbank, "Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery," *Applied and Computational Harmonic Analysis*, vol. 26, no. 2, pp. 283–290, 2009.
- [15] S. D. Howard, A. R. Calderbank, and S. J. Searle, "A fast reconstruction algorithm for deterministic compressive sensing using second order reed-muller codes," in *42nd IEEE Annual Conference on Information Sciences and Systems (CISS)*, 2008, pp. 11–15.
- [16] R. Calderbank, S. Howard, and S. Jafarpour, "Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 358–374, 2010.
- [17] A. Amini and F. Marvasti, "Deterministic construction of binary, bipolar, and ternary compressed sensing matrices," *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 2360–2370, 2011.
- [18] A. Agarwal, S. Negahban, and M. J. Wainwright, "Stochastic optimization and sparse statistical recovery: Optimal algorithms for high dimensions," in *NIPS*, 2012, pp. 1547–1555.
- [19] —, "Fast global convergence of gradient methods for high-dimensional statistical recovery," *The Annals of Statistics*, vol. 40, no. 5, pp. 2452–2482, 2012.
- [20] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. motivation and construction," in *Information Theory Workshop (ITW)*, 2010, pp. 1–5.
- [21] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Transactions on Information Theory*, vol. 57, no. 2, pp. 764–785, 2011.
- [22] D. Baron, M. F. Duarte, S. Sarvotham, M. B. Wakin, and R. G. Baraniuk, "An information-theoretic approach to distributed compressed sensing," in *45th Annual Conference on Communication, Control, and Computing (Allerton)*, 2005.
- [23] M. Fornasier and H. Rauhut, "Recovery algorithms for vector-valued data with joint sparsity constraints," *SIAM Journal on Numerical Analysis*, vol. 46, no. 2, pp. 577–613, 2008.
- [24] Multi-sensor and distributed compressive sensing. [Online]. Available: <http://dsp.rice.edu/cs>
- [25] D. Slepian and J. K. Wolf, "Noiseless coding of correlated information sources," *IEEE Transactions on Information Theory*, vol. 19, no. 4, pp. 471–480, 1973.

- [26] A. Goldsmith, *Wireless communications*. Cambridge University Press, 2005.
- [27] E. Arikan, “Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels,” *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3051–3073, 2009.
- [28] —, “Source polarization,” in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2010, pp. 899–903.
- [29] C. E. Shannon, “A mathematical theory of communication,” *Bell Systems Tech. J.*, pp. 379–423, 623–656, July and October 1948.
- [30] M. Karzand and I. Telatar, “Polar codes for q-ary source coding,” in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2010, pp. 909–912.
- [31] R. Durrett, *Probability: theory and examples*. Cambridge University Press, 2010, vol. 3.
- [32] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [33] W. Pratt, J. Kane, and H. C. Andrews, “Hadamard transform image coding,” in *Proceedings of the IEEE*, 1969, pp. 58–68.
- [34] 3GPP TS 25.213 V11.4.0 Release 11, “Spreading and modulation (fdd),” 2013.
- [35] K. J. Horadam, *Hadamard Matrices and Their Applications*. Princeton University Press, 2007.
- [36] S. Haghghatshoar and E. Abbe, “Polarization of the Rényi information dimension for single and multi-terminal analog compression,” *arXiv preprint arXiv:1301.6388*, 2013.
- [37] A. Hedayat and W. Wallis, “Hadamard matrices and their applications,” *The Annals of Statistics*, pp. 1184–1238, 1978.
- [38] M. H. Lee and M. Kaveh, “Fast Hadamard transform based on a simple matrix factorization,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 34, no. 6, pp. 1666–1667, 1986.
- [39] J. R. Johnson and M. Pueschel, “In search of the optimal Walsh-Hadamard transform,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing Proceedings (ICASSP)*, 2000, pp. 3347–3350.
- [40] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss, “Near-optimal sparse fourier representations via sampling,” in *Proceedings of the 34th annual ACM symposium on Theory of computing*, 2002, pp. 152–161.
- [41] A. C. Gilbert, M. J. Strauss, and J. A. Tropp, “A Tutorial on Fast Fourier Sampling,” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 57–66, 2008.
- [42] D. Lawlor, Y. Wang, and A. Christlieb, “Adaptive sub-linear Fourier algorithms,” *arXiv.org*, Jul. 2012.

- [43] H. Hassanieh, P. Indyk, D. Katabi, and E. Price, “Simple and practical algorithm for sparse Fourier transform,” *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1183–1194, 2012.
- [44] —, “Nearly optimal sparse Fourier transform,” *Proceedings of the 44th symposium on Theory of Computing*, pp. 563–578, 2012.
- [45] B. Ghazi, H. Hassanieh, P. Indyk, D. Katabi, E. Price, and L. Shi, “Sample-optimal average-case sparse fourier transform in two dimensions,” *arXiv.org*, Mar. 2013.
- [46] S. Pawar and K. Ramchandran, “A hybrid DFT-LDPC framework for fast, efficient and robust compressive sensing,” in *50th Annual Conference on Communication, Control, and Computing (Allerton)*, 2012, pp. 1943–1950.
- [47] —, “Computing a k -sparse n -length Discrete Fourier Transform using at most $4k$ samples and $O(k \log_2 k)$ complexity,” *arXiv.org*, May 2013.
- [48] T. Richardson and R. L. Urbanke, *Modern coding theory*. Cambridge University Press, 2008.
- [49] N. C. Wormald, “Differential Equations for Random Processes and Random Graphs,” *The Annals of Applied Probability*, vol. 5, no. 4, pp. 1217–1235, Nov. 1995.
- [50] M. G. Luby, M. Mitzenmacher, M. A. Shokrollahi, and D. A. Spielman, “Efficient erasure correcting codes,” *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 569–584, 2001.
- [51] Y. Wu and S. Verdú, “Rényi information dimension: Fundamental limits of almost lossless analog compression,” *IEEE Transactions on Information Theory*, vol. 56, no. 8, pp. 3721–3748, 2010.
- [52] —, “Optimal phase transitions in compressed sensing,” *IEEE Transactions on Information Theory*, vol. 58, no. 10, pp. 6241–6263, 2012.
- [53] D. L. Donoho, A. Javanmard, and A. Montanari, “Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing,” *IEEE Transactions on Information Theory*, vol. 59, no. 11, pp. 7434–7464, 2013.
- [54] S. Haghshatshoar, E. Abbe, and E. Telatar, “Adaptive sensing using deterministic partial hadamard matrices,” in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2012, pp. 1842–1846.
- [55] S. Kudekar and H. D. Pfister, “The effect of spatial coupling on compressive sensing,” in *48th Annual Conference on Communication, Control, and Computing (Allerton)*, 2010, pp. 347–353.
- [56] F. Zhang and H. D. Pfister, “Verification decoding of high-rate LDPC codes with applications in compressed sensing,” *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5042–5058, 2012.

- [57] A. G. Dimakis, R. Smarandache, and P. O. Vontobel, "LDPC codes for compressed sensing," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3093–3114, 2012.
- [58] M. Akçakaya and V. Tarokh, "Shannon-theoretic limits on noisy compressive sampling," *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 492–504, 2010.
- [59] G. Reeves and M. Gastpar, "Sampling bounds for sparse support recovery in the presence of noise," in *IEEE International Symposium on Information Theory (ISIT)*, 2008, pp. 2187–2191.
- [60] S. Sarvotham, D. Baron, and R. G. Baraniuk, "Measurements vs. bits: Compressed sensing meets information theory," in *44th Annual Conference on Communications, Control, and Computing Proceedings (Allerton)*, 2006.
- [61] M. J. Wainwright, "Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting," *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5728–5741, 2009.
- [62] W. Wang, M. J. Wainwright, and K. Ramchandran, "Information-theoretic limits on sparse signal recovery: Dense versus sparse measurement matrices," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2967–2979, 2010.
- [63] D. Guo, D. Baron, and S. Shamai, "A single-letter characterization of optimal noisy compressed sensing," in *47th Annual Conference on Communication, Control, and Computing (Allerton)*, 2009, pp. 52–59.
- [64] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Statistical-physics-based reconstruction in compressed sensing," *Physical Review X*, vol. 2, no. 2, p. 021005, 2012.
- [65] E. Abbe, "Universal source polarization and sparse recovery," in *IEEE Information Theory Workshop (ITW)*, 2010, pp. 1–5.
- [66] A. Rényi, "On the dimension and entropy of probability distributions," *Acta Mathematica Academiae Scientiarum Hungarica*, vol. 10, no. 1-2, pp. 193–215, 1959.
- [67] E. Arikan and I. Telatar, "On the rate of channel polarization," in *IEEE International Symposium on Information Theory (ISIT)*, 2009, pp. 1493–1495.
- [68] J. Barbier, F. Krzakala, and C. Schülke, "Compressed sensing and approximate message passing with spatially-coupled fourier and hadamard matrices," *arXiv preprint arXiv:1312.1740*, 2013.
- [69] Y. Wu and S. Verdú, "MMSE dimension," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 4857–4879, 2011.
- [70] S. Kudekar, T. J. Richardson, and R. L. Urbanke, "Threshold saturation via spatial coupling: Why convolutional LDPC ensembles perform so well over the BEC," *IEEE Transactions on Information Theory*, vol. 57, no. 2, pp. 803–834, 2011.

- [71] E. Arikan, "Polar coding for the slepian-wolf problem based on monotone chain rules," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2012, pp. 566–570.
- [72] T. Tao, "Sumset and inverse sumset theory for Shannon entropy," *Combinatorics, Probability & Computing*, vol. 19, no. 4, pp. 603–639, 2010.
- [73] A. Stam, "Some inequalities satisfied by the quantities of information of Fisher and Shannon," *Information and Control*, vol. 2, no. 2, pp. 101–112, 1959.
- [74] N. M. Blachman, "The convolution inequality for entropy powers," *IEEE Transactions on Information Theory*, vol. 11, no. 2, pp. 267–271, 1965.
- [75] E. H. Lieb *et al.*, "Proof of an entropy conjecture of Wehrl," *Communications in Mathematical Physics*, vol. 62, no. 1, pp. 35–41, 1978.
- [76] P. Bergmans, "Random coding theorem for broadcast channels with degraded components," *IEEE Transactions on Information Theory*, vol. 19, no. 2, pp. 197–207, 1973.
- [77] S. Leung-Yan-Cheong and M. E. Hellman, "The Gaussian wire-tap channel," *IEEE Transactions on Information Theory*, vol. 24, no. 4, pp. 451–456, 1978.
- [78] L. Ozarow, "On a source-coding problem with two channels and three receivers," *Bell System Technical Journal*, vol. 59, no. 10, pp. 1909–1921, 1980.
- [79] Y. Oohama, "The rate-distortion function for the quadratic Gaussian CEO problem," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 1057–1070, 1998.
- [80] H. Weingarten, Y. Steinberg, and S. Shamai, "The capacity region of the Gaussian multiple-input multiple-output broadcast channel," *IEEE Transactions on Information Theory*, vol. 52, no. 9, pp. 3936–3964, 2006.
- [81] S. Verdú and D. Guo, "A simple proof of the entropy-power inequality," *IEEE Transactions on Information Theory*, vol. 52, no. 5, pp. 2165–2166, 2006.
- [82] O. Rioul, "Information theoretic proofs of entropy power inequalities," *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 33–55, 2011.
- [83] M. H. Costa, "A new entropy power inequality," *IEEE Transactions on Information Theory*, vol. 31, no. 6, pp. 751–760, 1985.
- [84] A. Dembo, "Simple proof of the concavity of the entropy power with respect to added Gaussian noise," *IEEE Transactions on Information Theory*, vol. 35, no. 4, pp. 887–888, 1989.
- [85] C. Villani, "A short proof of the concavity of entropy power," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1695–1696, 2000.
- [86] R. Zamir and M. Feder, "A generalization of the entropy power inequality with applications," *IEEE Transactions on Information Theory*, vol. 39, no. 5, pp. 1723–1728, 1993.

- [87] T. Liu and P. Viswanath, "An extremal inequality motivated by multiterminal information-theoretic problems," *IEEE Transactions on Information Theory*, vol. 53, no. 5, pp. 1839–1851, 2007.
- [88] R. Liu, T. Liu, H. V. Poor, and S. Shamai, "A vector generalization of Costa's entropy-power inequality with applications," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1865–1879, 2010.
- [89] S. Artstein, K. Ball, F. Barthe, and A. Naor, "Solution of Shannon's problem on the monotonicity of entropy," *Journal of the American Mathematical Society*, vol. 17, no. 4, pp. 975–982, 2004.
- [90] A. M. Tulino and S. Verdú, "Monotonic decrease of the non-Gaussianness of the sum of independent random variables: A simple proof," *IEEE Transactions on Information Theory*, vol. 52, no. 9, pp. 4295–4297, 2006.
- [91] M. Madiman and A. Barron, "Generalized entropy power inequalities and monotonicity properties of information," *IEEE Transactions on Information Theory*, vol. 53, no. 7, pp. 2317–2329, 2007.
- [92] A. Wyner and J. Ziv, "A theorem on the entropy of certain binary sequences and applications—I," *IEEE Transactions on Information Theory*, vol. 19, no. 6, pp. 769–772, 1973.
- [93] H. S. Witsenhausen, "Entropy inequalities for discrete channels," *IEEE Transactions on Information Theory*, vol. 20, no. 5, pp. 610–616, 1974.
- [94] V. Jog and V. Anantharam, "The Entropy Power Inequality and Mrs. Gerber's Lemma for Abelian Groups of Order 2^n ," *arXiv preprint arXiv:1207.6355*, 2012.
- [95] P. Harremoës and C. Vignat, "An entropy power inequality for the binomial family," *JIPAM. J. Inequal. Pure Appl. Math.*, vol. 4, no. 5, 2003.
- [96] N. Sharma, S. Das, and S. Muthukrishnan, "Entropy power inequality for a family of discrete random variables," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2011, pp. 1945–1949.
- [97] O. Johnson and Y. Yu, "Monotonicity, thinning, and discrete versions of the entropy power inequality," *IEEE Transactions on Information Theory*, vol. 56, no. 11, pp. 5387–5395, 2010.
- [98] I. Kontoyiannis and M. Madiman, "Sumset and inverse sumset inequalities for differential entropy and mutual information," *arXiv preprint arXiv:1206.0489*, 2012.

Saeid Haghghatshoar

Personal Information

Address: EPFL, IC, LCM, INR 030, Station 14, CH-1015 Lausanne
Email: saeid.haghghatshoar@epfl.ch, haghghatshoar@gmail.com
Date of Birth: 23 August 1985
Nationality: Iranian

Research Interests

Compressed Sensing
Information Theory
Statistical Learning Theory
Graphical Models
Dynamical Systems and Chaotic Signal Processing

Education

- 2010-2014 **PhD in Communications Systems**, *EPFL*, Lausanne-Switzerland, Defended on 17 September.
2007–2009 **MSC in Communications Systems**, *Sharif University of Technology*, Tehran-Iran, *GPA – 18.70/20.0*.
2003–2007 **BSC in Electronics**, *Sharif University of Technology*, Tehran-Iran, *GPA – 18.30/20.0*.

Honors and Awards

- 2003 **25th rank** in International University Entrance Exam among more than 300,000 Participants
2007 **1st rank** in International Student Olympiad in Electrical Engineering
2008 Recipient of 2 Year Fellowship for Outstanding Student in Electrical Engineering
2009 Recipient of Research Award for Outstanding Applied MSC Thesis
2012 Recipient of Certificate of Appreciation for the Best Teaching Assistant from EPFL

Publications

Journal Papers:

- “A new entropy power inequality for integer-valued random variables”, S. Haghghatshoar, E. Abbe, E. Telatar, *IEEE Trans. on Information Theory*, vol. 60, pp. 3787–3796, 2014 and available at arXiv:1301.4185 [cs.IT], 2013.
- “A Fast Hadamard Transform for Signals with Sub-linear Sparsity in the Transform Domain”, R. Scheibler, S. Haghghatshoar, M. Vetterli, submitted to *IEEE Trans. on Information Theory* and available at arXiv:1310.1803 [cs.IT], 2013.
- “Polarization of the Rényi Information Dimension for Single and Multi Terminal Analog Compression”, S. Haghghatshoar, E. Abbe, submitted to *IEEE Trans. on Information Theory* and available at arXiv:1301.6388 [cs.IT], 2013.
- “Spatial Sound Localization via Multipath Euclidean Distance Matrix Recovery”, M. J. Taghizadeh, A. Asaei, S. Haghghatshoar, P. N. Garner, and H. Bourlard, submitted to *IEEE Journal of Selected Topics on Signal Processing*, 2014.

Conference Papers:

- “Adaptive sensing using deterministic partial Hadamard matrices”, S. Haghhighatshoar, E. Abbe, E. Telatar, Information Theory Proceedings (ISIT), pp. 1842–1846, 2012.
- “A new entropy power inequality for integer-valued random variables”, S. Haghhighatshoar, E. Abbe, E. Telatar, Information Theory Proceedings (ISIT), pp. 589–593, 2013.
- “Polarization of the Rényi Information Dimension for Single and Multi Terminal Analog Compression”, S. Haghhighatshoar, E. Abbe, Information Theory Proceedings (ISIT), pp. 779–783, 2013.
- “A Fast Hadamard Transform for Signals with Sub-linear Sparsity”, R. Scheibler, S. Haghhighatshoar, M. Vetterli, 51st Annual Allerton Conference on Communication, Control, and Computing, 2013.
- “Multi-terminal Probabilistic Compressed Sensing”, S. Haghhighatshoar, Information Theory Proceedings (ISIT), pp. 221–225, 2014.
- “Robust Microphone Placement for Source Localization from Noisy Distance Measurements”, M. J. Taghizadeh, S. Haghhighatshoar, A. Asaei, P. N. Garner, and H. Boursard, submitted to 40th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2015.

Selected Courses

Information theory, Adaptive filters, Estimation and Detection theory, Routing and congestion control, Queueing networks, Statistical theory, Dynamical systems, Measure theory and integration, Stochastic calculus, Functional analysis, Applied stochastic processes, Statistical physics for computer science, Advanced probability theory, Graphical models, Combinatorial optimization, Randomized algorithms, Stochastic integration, Differential geometry and Riemannian manifolds.

Teaching Experiences

- Principles of Digital Communications, Teaching Assistant, Spring 2010, 2011
- Advanced Digital Communications, Teaching Assistant, Fall 2011
- Information Theory, Teaching Assistant, Fall 2012
- Quantum Information Theory, Teaching Assistant, Fall 2013

Computer Skills

- Operating Systems: Macintosh, Microsoft Windows
- Programming Languages:
 - Proficient in C, C++, Java and Familiar with Python and MySQL
 - Proficient in Matlab and Familiar with Mathematica
- Digital Typography: Proficient in \LaTeX and Microsoft Word

Languages

- **Persian and Azari** Mother Tongue
- **English** Fluent
- **French and German** Beginner