

Sparse Recovery based Grant-free Random Access for Massive Machine-type Communication

vorgelegt von
Alexander Fengler, M.Sc.
ORCID: 0000-0002-8848-2360

an der Fakultät IV - Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
- Dr.-Ing. -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Thomas Magedanz

Gutachter: Prof. Dr. Giuseppe Caire

Gutachter: Prof. Dr. Gianluigi Liva

Gutachter: Prof. Dr. Krishna Narayanan

Tag der wissenschaftlichen Aussprache: 2. März 2021

Berlin 2021

Abstract

This dissertation treats a recently introduced modern random access protocol known as *unsourced* random access (U-RA). This protocol belongs to the family of *grant-free* random access protocols which, in contrast to the established grant-based protocols of current mobile communication standards, do not rely on an initial access phase, in which the active devices identify themselves and await the grant of dedicated transmission resources (e.g. time-frequency blocks) from the base station (BS). When messages are short and the number of active devices is large, which is a central specification of upcoming massive machine-type-communication (mMTC) scenarios, such a grant-based procedure is overly wasteful and may increase the delay to an unacceptable level. In a grant-free scenario the active devices transmit their payload right away without awaiting the grant of dedicated resources. A commonly discussed approach is to assign devices unique fixed identification sequences (called pilots in the following). Active devices then transmit their pilot followed directly by their message. The idea of unsourced random access is to go one step further and, ideally, abolish the identification step, such that active devices directly transmit a message from a common predefined message set. Information-theoretically such a behavior is captured by letting each user employ the same codebook, as opposed to the classical information-theoretic treatment of the multiple access problem where each user is assigned an individual codebook. The assumption of a common codebook allows to treat the random access problem in a way that captures the effect of short messages while still taking into account the physical properties of the channel. In this work I build on a previously introduced coding scheme for the U-RA problem on the AWGN channel termed coded compressed sensing (CCS). I introduce a novel decoding algorithm for the CCS scheme based on approximate message passing (AMP) and give an asymptotic analysis. The analysis shows that it is possible, under optimal decoding, to achieve the fundamental communication limit on the AWGN channel even when arbitrary many users communicate without any coordination. Furthermore, I show that the low-complexity AMP algorithm cannot achieve this limit with a uniform power allocation. Instead, an optimized non-uniform power allocation is necessary to improve the performance of the AMP algorithm beyond a certain limit. I use the developed analysis to derive a method to find such power allocation. Numerical evaluation show that the proposed scheme is efficient and the analysis is precise.

In the second part of the work I present an extension of the U-RA idea to a block-fading AWGN channel with a massive MIMO receiver. I introduce and study a modified CCS scheme that uses a covariance based decoding algorithm. The analysis shows that the sum-spectral-efficiency of such a U-RA system can grow as $\mathcal{O}(L/\log L)$, where L is the length of a coherence block. In the course of the analysis I show that the self Khatri-Rao product of random spherical matrices satisfies the so called restricted isometry property, which is a result of independent interest and find applications in e.g. the analysis of the limits of direction-of-arrival estimation.

Remarkably, the CCS scheme shows to work very well even in a high mobility scenario with short coherence block length $L \approx 100$ and hundreds of concurrent active users. In such a fast fading scenario pilot based random access would completely fail for this number of users. On the other hand, if the coherence block length is large enough, I show that a more conventional scheme based on randomly chosen pilots and maximum-ratio-combining can efficiently implement the U-RA idea in a massive MU-MIMO setting. Finally, I show that the CCS scheme can be deployed in an distributed setting, where multiple BSs can decode the messages of many active users with minimal corporation between the BSs and no required cell management ("cell-free"). Furthermore, I present a novel iterative scheme that allows to jointly detect the activity and the geographic positions of the active users. The practical performance of all the proposed schemes is evaluated by numerical simulations.

Zusammenfassung

In dieser Arbeit behandle ich ein kürzlich vorgestelltes Random-Access Verfahren zur drahtlosen Kommunikation das unter dem Namen “Unsources Random Access” (U-RA) bekannt wurde. Das Verfahren gehört zur Klasse der zuweisungsfreien Random-Access Protokolle, die, im Gegensatz zu den etablierten zuweisungsbasierten Protokollen aktueller mobiler Kommunikationsstandards, nicht auf eine erste Kennenlernphase angewiesen sind, während der die Benutzer sich identifizieren und auf die Zuweisung von festen Kommunikationsressourcen (z.B. Zeit-Frequenzblöcken) durch die Basisstation (BS) warten müssen. So ein zuweisungsbasiertes Protokoll ist besonders ineffizient, wenn die Zahl der Benutzer sehr groß ist, sie aber nur selten aktiv sind und nur kurze Nachrichten zu übermitteln haben. Das ist z.B. bei typischen Szenarien zukünftiger Maschinenkommunikation der Fall, in denen eine Vielzahl von Sensoren in unregelmäßigen Abständen ihren Messungen übermitteln. Die Zuweisungsphase würde in so einem Szenario die Verzögerungen zwischen Sender und Empfänger deutlich erhöhen. In einem zuweisungsfreien Protokoll hingegen senden Benutzer ihre Daten direkt, ohne auf die Zustimmung der BS zu warten. Ein gerne verwendeter Ansatz zur zuweisungsfreien Kommunikation besteht darin, jedem Benutzer eine eindeutige Signatur zuzuweisen. Ein aktiver Benutzer sendet dann erst seine Signatur, direkt gefolgt von der eigentlichen Nachricht. So ein Ansatz ist nicht komplett zuweisungsfrei, da erst jedem Benutzer eine Signatur zugewiesen werden muss. Im U-RA Ansatz werden hingegen alle Benutzer von Anfang an als ununterscheidbar behandelt. In dem Fall muss der Empfänger nur eine Liste der gesendeten Nachrichten rekonstruieren, wobei es prinzipiell nicht möglich ist eine bestimmte Nachricht einem bestimmten Benutzer zuzuordnen. Informationstheoretisch spiegelt sich dieser Ansatz in der Annahme, dass alle Benutzer dasselbe Codebuch haben. Das erlaubt es das Problem auf eine Weise zu behandeln die es ermöglicht die kurze Länge der Nachrichten und gleichzeitig die Art des Kanals mit einzubeziehen. In dieser Arbeit baue ich auf einer bestehenden Methode auf, für U-RA zu kodieren, die als coded-compressed-sensing (CCS) bekannt ist. Ich stelle einen verbesserten Dekodieren für CCS vor, basieren auf dem approximate message passing (AMP) Algorithmus, und analysiere diesen Dekoder. Die Analyse zeigt, dass es unter optimaler Dekodierung möglich ist die fundamentalen Grenzen der Mehrbenutzerkommunikation auf dem realen AWGN Kanal zu erreichen, selbst wenn unendlich Benutzer ohne jede Kooperation senden. Außerdem kann

ich zeigen, unter welchen Bedingungen diese Grenzen auch unter sub-optimaler AMP Dekodierung erreichbar sind. Simulationen zeigen, dass der vorgestellte Dekodierer effizient ist und die theoretische Analyse gute Vorhersagen liefert. Im zweiten Teil der Arbeit stelle ich vor, wie das U-RA Konzept auf einen AWGN Kanal mit Block-Fading und mehreren Empfangsantennen anwendbar ist. Ich stelle eine Modifikation des CCS Schemas vor, die mit einem Kovarianz-basierten Algorithmus dekodiert werden kann. Die Analyse zeigt, dass die gesamte spektrale Effizienz eines solchen Systems in der Größenordnung $\mathcal{O}(L/\log L)$ wachsen kann, wobei L die Länge eines Koheränzblocks ist. Im Rahmen dieser Analysen habe ich gezeigt, dass sogenannte selbst Khatri-Rao product einer zufälligen sphärischen Matrix die restricted isometry property hat. Das ist ein Resultat das von unabhängigem Interesse ist und in Bereichen wie der Bestimmung von Einfallswinkeln in Radaranwendungen benutzt werden kann. Das CCS Verfahren funktioniert sehr gut und erlaubt es selbst in Szenarien mit hoher Mobilität und kurzen Koheränzbloklängen von etwa $L \approx 100$ noch zuverlässige Kommunikationen von einigen hundert Benutzern. Bei solchen kurzen Koheränzböcken und Benutzerzahlen würden klassische Verfahren komplett versagen. Andererseits, für lange Koheränzböcke stelle ich ein alternatives Verfahren vor, das auf der zufälligen Wahl von Signaturen aus einer vorgegebenen Menge basiert und maximum-ratio-combining benutzt um die Nachrichten unterschiedlicher Benutzer zu trennen. Im letzten Teil der Arbeit untersuche ich, wie der U-RA Ansatz mit verteilten Empfangsstationen funktionieren kann. Ich zeige, dass es trotz minimaler Kooperation unter den Empfangsstationen und ohne Zellenmanagement, möglich ist das U-RA Konzept effizient zu implementieren. Außerdem stelle ich einen neuen Algorithmus vor der es den Empfangsstationen erlaubt gleichzeitig die Aktivität der Benutzer und ihre Position zu bestimmen. Die möglichen Leistungen aller vorgestellten Algorithmen werden mit Simulationen belegt.

Acknowledgements

I would like to thank my advisor Prof. Giuseppe Caire for giving me the opportunity of doing a PhD with him even though I did not know a single thing about information theory when I started. Thanks to his support and guidance I learned a lot during the last four years. I would also like to thank Prof. Gianluigi Liva and Prof. Krishna Narayanan for agreeing to review this work and being part of the committee. During my work I was financially supported by the Alexander-von-Humboldt foundation and the DAAD. I would like to thank Peter and Saeid for the many scientific discussions, the endless proofreading of all of my texts and other valuable advice. A special thanks goes to Jana for helping me with a lot of Verwaltungswork and generally organizing all the things. I would like to thank Xiaoshen, Hendrik, Mohammed and Lefteris for being great roommates, Hamdi for regularly disrupting me (in a good way), Mahdi for forwarding me Mozghan's thesis template, Mozghan for her thesis template and Thomas for dealing with all the technical stuff. Generally, I would like to thank all the other colleagues and all the guests who have been part of our group within the last four years, including Osman, Cagkan, Rafael, Kai, Hao, Tianyu, Onur, Yi, Fernando, Ilker, Muah, Andreas, Andreas, Saeid, Saeid, Saeed, Sebastian, Giusi, Lorenzo, Xinping and others for a pleasant and encouraging atmosphere within the chair and for many lunch- and coffeetimes.

Ich möchte auch meinen Eltern Erhard und Tatjana danken, die hart gearbeitet haben um mir ein Studium zu ermöglichen und mich immer unterstützt haben. Ich danke auch ganz besonders meiner Freundin Katayoun, die immer für mich da war.

*Alexander Fengler,
Berlin, Nov 2020*

Contents

| | |
|---|-------------|
| List of Figures | xiii |
| List of Tables | xv |
| 1. Background and Motivation | 1 |
| 1.1. Grant-Free Random Access | 1 |
| 1.2. Unsourced Random Access | 2 |
| 1.3. Sparse Regression Codes | 4 |
| 1.4. Massive MIMO | 5 |
| 1.5. Sparse Bayesian Learning and Joint-Sparse Support Recovery | 6 |
| 1.6. Slow-fading | 9 |
| 1.7. Cell-free MIMO | 9 |
| 1.8. Notation | 11 |
| 1.9. Publications | 12 |
| 2. Multi-User SPARCs for the AWGN Channel | 13 |
| 2.1. Outline and Main Contributions | 13 |
| 2.2. Channel Model | 14 |
| 2.3. Concatenated Coding | 15 |
| 2.4. Inner Channel | 16 |
| 2.4.1. Asymptotic Error Analysis | 17 |
| 2.4.2. The $J \rightarrow \infty$ limit | 23 |
| 2.5. Hard Decision | 25 |
| 2.6. Outer Channel | 26 |
| 2.6.1. Tree code | 30 |
| 2.7. Analysis of the concatenated scheme | 33 |
| 2.8. Optimizing the Power Allocation | 39 |
| 2.9. Considerations for Practical Implementation | 42 |
| 2.10. Finite-length Simulations | 44 |
| 2.11. Summary | 45 |

| | |
|---|------------|
| 3. Block Fading Multi-User-MIMO Channel | 49 |
| 3.1. Outline and Main Contributions | 49 |
| 3.2. Activity Detection | 50 |
| 3.2.1. Signal Model | 50 |
| 3.2.2. Covariance-Based Algorithms | 52 |
| 3.2.3. Empirical Comparison: MF, ML, NNLS and MMV-AMP | 61 |
| 3.3. Massive MIMO Unsourced Random Access | 73 |
| 3.3.1. Unsourced Random Access as AD Problem | 74 |
| 3.3.2. Discussion and Analysis for One-Slot Transmission | 76 |
| 3.3.3. Coding over Multiple Coherence Blocks | 77 |
| 3.3.4. Asymptotic Analysis of the Concatenated Code | 78 |
| 3.3.5. Simulations | 79 |
| 3.3.6. Impact of Large-Scale Fading | 83 |
| 3.4. Massive MIMO Unsourced Random Access - Slow Fading | 85 |
| 3.4.1. Pilot-based Massive MIMO U-RA | 85 |
| 3.4.2. Analysis | 86 |
| 3.4.3. Complexity | 87 |
| 3.4.4. Simulations | 88 |
| 3.4.5. Collisions | 88 |
| 3.5. Going Cell-Free | 91 |
| 3.5.1. Model | 92 |
| 3.5.2. U-RA Simulations | 93 |
| 3.5.3. Joint AD and Position Estimation | 94 |
| 3.6. Summary | 99 |
| 4. Conclusion | 101 |
| Appendices | 103 |
| A. Optimal product distribution | 105 |
| B. $D_{\text{KL}}(\text{Mult} \parallel \text{Bin})$ | 107 |
| C. Exponential L^1 convergence implies exponential pointwise convergence a.e. | 109 |
| D. Proof of Theorem 5 | 111 |
| E. Proof of Theorem 6 | 117 |
| F. Proof of Theorem 11 | 121 |

| | |
|---|-----|
| G. Proof of the RIP, Theorem 12 | 127 |
| H. Proof of the Recovery Guarantee for NNLS, Theorem 13 | 133 |
| I. Analysis of Error of the Sample Covariance Matrix | 137 |
| Bibliography | 141 |

List of Figures

| | |
|--|----|
| 2.1. Detection error trade-off | 26 |
| 2.2. Achievable rates with the tree code | 32 |
| 2.3. Achievable rates with the tree code, function of (α) | 33 |
| 2.4. \mathcal{E}_{in} vs S_{in} optimal | 37 |
| 2.5. \mathcal{E}_{in} vs S_{in} Algorithmic | 38 |
| 2.6. \mathcal{E}_{in} over J | 39 |
| 2.7. Power Allocation Optimization | 43 |
| 2.8. P_e over E_b/N_0 Power Allocation | 44 |
| 2.9. E_b/N_0 over K_a | 46 |
| 3.1. AMP State Evolution | 69 |
| 3.2. Support detection error $K_a = L$ | 70 |
| 3.3. Support detection error $K_a > L$ | 71 |
| 3.4. LSFC estimation error | 71 |
| 3.5. Phase Transition | 72 |
| 3.6. Box Constraint | 73 |
| 3.7. CCS: P_e over M | 80 |
| 3.8. CCS with ML: P_e over E_b/N_0 | 81 |
| 3.9. CCS with ML: E_b/N_0 over K_a | 81 |
| 3.10. CCS with Matched Filter: E_b/N_0 over K_a | 82 |
| 3.11. Comparison MF,AMP,ML | 83 |
| 3.12. Impact of LSF | 84 |
| 3.13. MRC: E_b/N_0 over K_a | 89 |
| 3.14. Collision Recovery AWGN | 90 |
| 3.15. Collision Recovery MRC | 91 |
| 3.16. Square Cell | 93 |
| 3.17. CCS Distributed Decoding: E_b/N_0 over K_a | 94 |
| 3.18. CCS Joint Outer Decoding: E_b/N_0 over K_a | 95 |
| 3.19. Localisation MSE | 97 |
| 3.20. Missed Detection Rate with Joint Positioning | 98 |

| | |
|---|----|
| 3.21. False Alarm Rate with Joint Positioning | 98 |
|---|----|

List of Tables

| | |
|--|----|
| 3.1. Parameters for Figure 3.9 | 82 |
|--|----|

1. Background and Motivation

1.1. Grant-Free Random Access

In the last decades mobile communication has evolved to a point where almost every person on earth owns a mobile phone. The goal for next generation mobile systems is to provide ubiquitous wireless access to a myriad of machine type devices like autonomous vehicles, mobile sensors or smart home devices. Such scenarios have been described with names like internet-of-things (IoT) or massive machine-type-communication (MTC). Current wireless systems are not designed to handle a massive amount of terminals with sporadic activity and short messages. The way that shared access of the communication resources is handled in the current communication standards (e.g. 3G,4G-LTE or 5G NR) [1, 2] is through a *grant-based* procedure: The active devices identify themselves and await the grant of some dedicated transmission resource. When messages are short and the number of active devices is large such a grant-based procedure is overly wasteful and may increase the delay to an unacceptable level. In contrast, in a *grant-free* scenario [3–7] the active devices transmit their payload right away without awaiting the grant of dedicated resources. A prominent way of implementing a grant-free scheme is to assign devices unique fixed identification sequences (pilots) [3–6]. Active devices then transmit their pilot followed directly by their message. The downside of any fixed allocation of pilots to users is that the cost of identification grows with the total number of devices [8], even if most of them remain inactive for a very long time. A novel paradigm, later referred to as *unsourced random-access* (U-RA), gets around this limitation and allows for a communication system that is completely independent of the number of inactive users [9]. The idea of unsourced random access is to, ideally, abolish the identification step, such that active devices directly transmit a message from a common predefined message set.

Information-theoretically such a behavior is captured by letting each user employ the same codebook. In contrast to the massive MTC requirements, the traditional information theoretic treatment of the multiple-access uplink channel is focused on few users K , large blocklength n and coordinated transmission, in the sense that each user is given an individual distinct codebook, and the K users agree on which rate K -tuple inside the capacity region to operate [10–12]. Mathematically, this is reflected by considering the limit

of infinite message- and blocklength while keeping the rate and the number of users fixed. This approach does not sufficiently capture the bursty random arrival of messages in real world multiple access networks [13], which lead to the widespread success of packet based random-access models [14–17]. Such models are based on simplified collision channels [18], which ignore the underlying physical communication channel and are thereby limited in the achievable performance [19, 20]. The assumption of a common codebook in U-RA allows to treat the random access problem in a way that captures the effect of short messages while still taking into account the physical properties of the channel.

1.2. Unsourced Random Access

Polyanskiy formulated U-RA as information-theoretic problem in [9] and presented a way to calculate an upper bound on the error probability of U-RA with a Gaussian random codebook under a finite blocklength constraint [9, Theorem 1]. This result was used to calculate the energy efficiency of random coding, i.e. the least required energy-per-bit to noise power ratio E_b/N_0 such that the error probability is within some tolerance ϵ .

The energy efficiency of the popular slotted Aloha protocol [14] was evaluated by dividing the available blocklength n into V slots. This creates smaller AWGN channels of blocklength V/n , for which the finite blocklength capacity at a fixed error probability can be approximated by the normal approximation [21]. It was found that the energy efficiency of slotted Aloha is very poor compared to the random coding achievability bound, which was mainly due to the rapid rise of the collision probability with the number of active users. To alleviate this problem Polyanskiy introduced the concept of T -fold Aloha, which refers to a hypothetical slotted coding scheme that is able to resolve up to T collision in one Aloha slot. It was shown that the energy efficiency of T -fold Aloha is much better than of regular Aloha, even for small values of T .

An asymptotic converse on the achievable energy efficiency is given by the Shannon limit for the real AWGN channel [10, 11]

$$E_b/N_0 \geq \frac{2^{2\mu} - 1}{2\mu} \quad (1.1)$$

where $\mu = K_a B/n$ is the sum-spectral efficiency, K_a is the number of concurrently transmitting users, B is the length of messages in bits and n is the number of available channel uses.. (1.1) shows we can distinguish between two extremal behaviors in multiple access systems, depending on whether $K_a B \ll n$ or $K_a B > n$. The first regime is called *power-limited* since here $(2^{2\mu} - 1)/(2\mu) \approx \text{const.}$, and therefore the required energy depends only weakly on the number of active users. On the other hand, if $K_a B > n$, any increase in K_a

will exponentially increase the required E_b/N_0 . This is known as the *interference-limited* regime. The Shannon bound is known to be tight for the conventional AWGN-MAC, but it is not necessarily tight for the unsourced case. In Chapter 2 I establish sufficient conditions on the scaling of K_a under which (1.1) is indeed tight in the unsourced case.

Following the initial work of Polyanskiy several practical approaches were suggested which successively reduced the gap to the random coding achievability bound [22–27]. On a high level many of the suggested approaches consist of creating a pool of resources (e.g. timeslots [22, 25, 28], spreading sequences [26, 27]) of which active users pick one at random and use it to communicate. This is reminiscent of the initial access protocol in current communication standards [1, Ch. 14.3], where the pool of resources consists of orthogonal preamble sequences. Also Aloha falls into this broad class of schemes with time slots as orthogonal resources. The main problem with such a strategy is that, as the number of active users grow, the probability of collisions, i.e. multiple users picking the same resource, becomes large. There are two ways to deal with this problem. One is to use a code that is able to resolve a small number of collisions and the other is to make the pool of resources non-orthogonal and large enough such that the probability of collisions becomes negligible. The first way is essentially the concept of T -fold Aloha, introduced by Ordentlich and Polyanskiy [22], and was practically implemented e.g. in [25, 28]. The second approach was chosen e.g. in [23, 26, 27]. Basically all existing approaches to U-RA share the property that the required energy-per-bit starts to grow rapidly as soon as the number of active users that collide on one resource becomes too large. That is because strong codes are necessary to resolve the multiple-access interference of many users. The obvious way to reduce the number of collisions is to make the pool of resources as large as possible. Random coding can be seen as an extreme example of the second approach, where the pool of resources is equal to all possible codewords.

The U-RA problem on the real AWGN channel can be written as a sparse recovery problem of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z} \tag{1.2}$$

where each column of $\mathbf{A} \in \mathbb{R}^{n \times 2^B}$ is one codeword of the common codebook. $\mathbf{x} \in \{0, 1\}^{2^B}$ is a binary vector with entries equal to one at the indices of the transmitted codewords and zero otherwise. $\mathbf{z} \in \mathbb{R}^n$ is the vector of Gaussian noise. Although there are efficient algorithms with polynomial complexity to solve such a sparse recovery problem, the dimension 2^B is extremely large, even for comparably small values like $B = 100$ bits. A divide-and-conquer approach for this type of sparse recovery problems, referred to as Coded-Compressed-Sensing (CCS), was proposed in [24]. The idea is to split each transmission up into S subslots. In each subslot the active users send a column from a common

inner coding matrix, while the symbols across all subslots are chosen from a common *outer tree code*. One of the strong properties of this coding scheme is that the resolving of the multiple-access interference is done in the domain of K_a sparse 2^J dimensional vectors, which simplifies the processing significantly. In the first part of this work I build upon the finding of [24] and its similarity to sparse regression codes (SPARCs) to propose an improved inner approximate message passing (AMP) decoder. I present a numerical method to analyse the error probability of the decoder for finite J and give a closed form result on the achievable rates that holds in the limit of $K_a, J \rightarrow \infty$.

1.3. Sparse Regression Codes

SPARCs were introduced in [29] as a class of channel codes for the point-to-point AWGN channel that can achieve rates up to Shannon capacity under maximum-likelihood decoding. Later, it was shown that SPARCs can achieve capacity under approximate message passing (AMP) decoding with either power allocation [30] or spatial coupling [31]. AMP is an iterative low-complexity algorithm for solving random linear estimation problems or generalized versions thereof [32–34]. A recent survey on SPARCs can be found in [35].

One of the appealing features of the AMP algorithm is that it is possible to analyse its asymptotic error probability, averaged over certain random matrix ensembles, through the so called state evolution (SE) equations [34, 36]. Interestingly the SE equations can also be obtained as the extreme point conditions of the replica symmetric (RS) potential, a $\mathbb{R} \rightarrow \mathbb{R}$ function that was first calculated through the non-rigorous replica method [37, 38]. It was shown that in random linear estimation problems the extreme points of the RS-potential also characterize the symbols-wise posterior distribution of the input elements and therefore also the error probability of several optimal estimators like the symbol-by-symbol maximum-a-posteriori (SBS-MAP) estimator [39, 40]. The difference between the AMP and the SBS-MAP estimate is that the SBS-MAP estimate always corresponds to the global minimum of the RS-potential, while the AMP algorithm gets ‘stuck’ in local minima. The rate below which a local minimum appears was called the *algorithmic* or *belief-propagation* threshold in [31, 39, 41]. It was shown in [41, 42] that, despite the existence of local minima in the RS-potential, the AMP algorithm can still converge to the global minimum when used with spatially coupled matrices. Although the RS-potential was derived by (and named after) the non-rigorous replica method, it was recently proven to hold rigorously [43, 44] for the case of Gaussian iid measurement matrices. The proof of [44] is more general in the sense that it includes the case where the unknown, to be estimated, vector \mathbf{s} consists of blocks of size 2^J and each block is considered to be drawn iid from some distribution on \mathbb{R}^{2^J} .

1.4. Massive MIMO

In the second part of the work I present an extension of the U-RA idea to a block-fading AWGN channel with a massive MIMO receiver [45, 46]. The term massive MIMO was introduced in [47] and describes a variant of multi-user MIMO (MU-MIMO) where a BS with a *large* antenna array consisting of $M \gg K$ antennas serves K users. Massive MIMO differs from conventional MU-MIMO in that M is large enough such that the effect of small scale fading can be averaged out. Specifically, it is possible to achieve large gains with channel-state information (CSI) being available *only* at the receiver [46, 47]. Not having to know the CSI at the transmitter is a huge advantage compared to conventional MU-MIMO, but in an mMTC setting with sporadic activity it can still be very costly to estimate the channel coefficients of all active users. The state-of-the-art approach to grant-free random-access in massive MIMO consist of assigning unique orthogonal pilot sequences to each users. Then, in an initial identification phase, the active users are detected and their channel vectors are estimated. These estimates are subsequently used to estimate the transmitted data symbols. It was shown in [4] that, through the use of random non-orthogonal pilot sequences, it is possible to identify an arbitrary amount of K_a active users out of K_{tot} total users and estimate their channel vectors in the asymptotic regime of $K_a, K_{\text{tot}}, L, M \rightarrow \infty$ when $K_a/L < 1$ and K_a/K_{tot} stay fixed and $L \gg M$. Note, that if $K_a > L$ it is in general impossible to estimate the complete $K_a \times M$ channel matrix reliably from the $L \times M$ measurements because the number of unknowns is larger then the number of constraints.

A novel approach to activity detection (AD) in the massive MIMO setting was presented in [48]. It was shown that through the use of just the covariance information in the measurements it is possible to estimate the Large-Scale-Fading-Coefficients (LSFCs), and therefore also identify the activity, of $K_a > L$, specifically $K_a = \mathcal{O}(L^2/\log^2(K_{\text{tot}}/K_a))$, active users even if it is not possible to estimate the channel vectors of all active users. Similar results were found in the context of direction-of-arrival estimation and will be discussed later. Although this result is remarkable, it is not practicable to use it in the existing grant-free massive MIMO schemes, because they rely on the CSI to estimate the transmitted data symbols. It is possible to embed information in the AD process, e.g. one bit can be transmitted in the AD phase by assigning each user two unique pilot sequences [49]. Clearly such an idea is impractical for the transmission of more than a handful of bits. In contrast, as I will show in Chapter 3, U-RA with a modified CCS scheme allows to combine covariance based AD and data transmission in a way that is scalable to a moderate number of message bits and inherits all the advantages of U-RA. The inner decoding problem in the CCS scheme for the block-fading MIMO channel model

turns out to be mathematically equivalent to a massive MIMO AD problem where L takes the role of the coherence blocklength and K_{tot} takes the role of 2^J where J is the number coded bits each user can transmit per coherence block. In Chapter 3 I present a complete and rigorous proof of the $K_a = \mathcal{O}(L^2/\log^2 K_{\text{tot}})$ scaling result and use it to derive an achievable scaling for the sum-spectral efficiency in the U-RA setting. The results show that the presented scheme can support $K_a = \mathcal{O}(L^2)$ active users and the sum-spectral efficiency can grow as $\mathcal{O}(L)$ up to logarithmic factors with a required energy-per-bit that can be made arbitrary low by increasing M . This can be achieved in a completely non-coherent way, i.e. it is at no point necessary to estimate the channel coefficients of all the active users. In contrast, no coherent scheme can support more active users than the coherence block-length, since in this regime it is not possible to estimate the channel coefficients of all active users. These are important properties to enable easily deployable, low-latency, energy efficient communication in an IoT setting. The scaling result is obtained by analyzing a Non-Negative Least-Squares (NNLS) algorithm applied to the sample covariance information, which was recently considered for LSFC estimation in [50]. The analysis in [50] showed, that with a random choice of pilot sequences the LSFCs of up to $K_a = \mathcal{O}(L^2)$ users could be estimated, but the proof was limited by the assumptions of $K_a = K_{\text{tot}}$, $K_{\text{tot}} \leq L^2$ and $M \rightarrow \infty$. The result in Chapter 3 lifts all these restrictions and shows that K_{tot} may be potentially much larger than L^2 and K_a , where one needs to pay only a poly-logarithmic penalty $\mathcal{O}(\log^2(\frac{K_{\text{tot}}}{K_a}))$ for increasing the total number of users K_{tot} . Furthermore, an improved algorithm for AD based on the Maximum-Likelihood (ML) estimation of the LSFCs of the active users is presented, referred throughout this work as the "ML Algorithm". The resulting likelihood maximization is a non-convex problem, that can be solved (approximately) by iterative componentwise minimization. This yields an iterative scheme based on rank-1 updates whose complexity is comparable to that of NNLS or Multiple-Measurement-Vector AMP (MMV-AMP), the algorithm used in [4,5]. Extensive numerical simulations show that the ML algorithm is superior to NNLS and to MMV-AMP in any regime, and does not suffer from the ill-conditioned behavior of MMV-AMP for the case of large M .

1.5. Sparse Bayesian Learning and Joint-Sparse Support Recovery

The componentwise optimization of the log-likelihood function was developed in [51–54], where the sparse Bayesian learning (SBL) framework was introduced to find the optimal vector of weights in a linear regression problem. In the SBL framework it is assumed that the weight vector follows a Gaussian prior distribution with zero mean and a diagonal

covariance matrix. The entries of the covariance matrix are estimated by maximizing the likelihood of the data. It was observed that the likelihood function maximization typically yields a sparse result, which is a desirable property in statistical learning. The maximum of the likelihood function can be computed iteratively by following the general expectation-maximization (EM) framework [55,56], but it was found that a componentwise optimization leads to faster convergence while still being guaranteed to converge to at least a local maximum of the likelihood function [52], similar to EM. The SBL framework was extended and applied to basis selection [57], compressed sensing [58] and also the MMV problem [59], where the latter was termed M-SBL. Here, the task is to recover $\mathbf{X} \in \mathbb{C}^{K_{\text{tot}} \times M}$ (here we stick to the notation introduced in this paper) from multiple measurements $\mathbf{Y} \in \mathbb{C}^{L \times M}$ of the form

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{Z} \quad (1.3)$$

Following the SBL framework, it is assumed that the rows of \mathbf{X} are distributed according to

$$\mathbf{X}_{:,i} \sim \mathcal{CN}(0, \gamma_i \mathbf{I}_M). \quad (1.4)$$

Note, that in the literature the term M-SBL has often been used ambiguously to refer to the ML estimate of the parameters γ_i as well as to the algorithm used to find this solution, which may be either coordinate-wise optimization or EM. Both of these algorithms lead to similar solutions [59], but here we adopt the componentwise optimization algorithm, since it can be efficiently implemented using rank-1 updates leading to a significant complexity reduction compared to the EM version.

Let \mathbf{X} have K_a non-zero rows. The identifiability limits of the ML solution of M-SBL were analysed in [60–64]. While the early work [60] was restricted to the case $K_a \leq L$, [61] made the distinction between recovering \mathbf{X} , which necessarily requires $K_a \leq L$, and recovering the vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{K_{\text{tot}}})$. It was noticed in [61] that the recovery of $\boldsymbol{\gamma}$ is governed by the properties of the Khatri-Rao product $\mathbf{A} \odot \mathbf{A}$ and it was proven that, with a random choice of \mathbf{A} , up to $K_a = \mathcal{O}(L^2)$ non-zero entries of $\boldsymbol{\gamma}$ can be recovered uniquely if the covariance matrix of \mathbf{Y} is known exactly. The proof of [61] (similar proofs were given independently in [50,64]) is based on the fact that any $2K_a$ columns of $\mathbf{A} \odot \mathbf{A}$ are linearly independent almost surely for K_a up to $\mathcal{O}(L^2)$. This proof deals only with the identifiability though, i.e., it does not apply to a specific recovery algorithm and does not take into account the uncertainty in estimating the covariance matrix of \mathbf{Y} , and therefore gives no clue on the robustness of the recovery with respect to the unavoidable covariance estimation error. It is well known in the compressed sensing literature that stronger conditions are needed to guarantee algorithmic robust recovery [65]. Upper and lower bounds on the performance of the ML solution of M-SBL for the noisy case have been derived in [60] for $K_a \leq L$ and in [63]

for $K_a = \mathcal{O}(L^2)$, but these bounds contained parameters which were exponentially hard to compute for a given matrix \mathbf{A} and so no concrete scaling of K_a could be given. A coherence based argument was given in [61] to analyse the performance of a covariance based LASSO algorithm, but it was only possible to guarantee recovery for up to $K_a = \mathcal{O}(L)$ coefficients. This is a well known limitation of coherence based arguments, known as the "square-root bottleneck" [65].

In Chapter 3 this bottleneck is circumvented by proving the restricted isometry property (RIP) of a properly centered and rescaled version of $\mathbf{A} \odot \mathbf{A}$ for random \mathbf{A} . This allows to prove robust recovery guarantees for both the NNLS algorithm of [50] and a constrained variant of the ML solution, showing that $K_a = \mathcal{O}(L^2)$ coefficients can be recovered. Although the constrained ML yields a combinatorial minimization with exponential complexity and therefore is not useful in practice, it is shown that the scaling law for successful detection of the activity pattern of the constrained ML scheme is the same (up to logarithmic factors in the scaling of M) as what was found for NNLS. Therefore, it is reasonable to believe that the (low-complexity) ML algorithm achieves the same scaling law. An intuitive argument is provided which, at least heuristically, explains why the componentwise optimization can be expected to converge to the global optimum in the considered scaling regime. Note, that an analysis of the constrained ML estimator was recently presented in [66]. However, the results in [66] are based on a RIP result that was first claimed and then withdrawn by the same authors [67]. Hence, the result in Chapter 3, based on a novel RIP result and a few consequent modifications which are duly proven, essentially rigorizes the analysis presented in [66].

The full characterisation of the global (unconstrained) ML solution and the conditions under which the iterative estimate coincides with it remains open. Some progress has recently been made in [68], where it was shown the global optimality of the algorithmic solution can be checked, given \mathbf{A} and the true γ° , by a linear feasibility program. In contrast, the presented recovery guarantees for the NNLS algorithm hold for all K_a -sparse γ and are given in closed form (up to unspecified constants). Based on the asymptotic Gaussianity of ML estimators in general, it was shown in [68] that for large M the distribution of the ML estimation error, for a fixed γ° , can be characterized numerically by the solution of a quadratic program.

The coordinate-wise optimization algorithm for M-SBL was also independently re-discovered and investigated in the context of angle-of-arrival estimation [69–71]. It was noted in [69] that the update equation can be equivalently derived by an iterative weighted least-squares (WLS) approach, which asymptotically minimizes the variance of the estimation. The resulting algorithm has therefore been called *iterative asymptotic sparse minimum variance stochastic ML* (SAMV-SML). Recently, in [72] a similar WLS estimator was derived and

an iterative algorithm was given to find an approximation of the WLS minimizer. The performance reported was very similar to M-SBL, but at a much higher complexity-per-iteration of $\mathcal{O}(L^4 K_{\text{tot}}^2)$, compared to $\mathcal{O}(L^2 K_{\text{tot}})$ for the coordinate-wise optimization with rank-1 updates.

1.6. Slow-fading

Although the presented U-RA scheme with CCS and ML AD shows empirically to be superior to any other MIMO U-RA scheme if run with the same coherence block-length, its complexity grows proportional to $\mathcal{O}(L^2)$ and becomes unfeasibly large when L is in the order of a couple of thousands. Therefore it may be necessary to split up the coherence block-length artificially, which reduces the achievable spectral efficiencies.

In Section 3.4 I introduce a simple new U-RA scheme that works when $L \gg K_a$. If the coherence block-length is long enough the first part of the coherence block-length can be dedicated to transmitting randomly chosen non-orthogonal pilot sequences. In this first phase the BS uses one of the AD algorithms discussed in Section 3.2 to detect the indices of the transmitted pilots and subsequently produces a linear minimum-mean-square-error (LMMSE) estimate of the channel vectors of the active users. The estimated channel vectors are then used in the second phase as temporary identifiers which allow to separate the data streams of the active users via maximum-ratio-combining (MRC). This concept is reminiscent of the more conventional grant-free random-access approach [3–7] where each user is assigned a fixed pilot sequence. The difference here is that pilot collisions, i.e. two or more active users picking the same pilot, are possible. I show that these collisions can be resolved by using a low-rate polar code with a successive-cancellation-list (SCL) decoder [73–75] as single-user error correcting code. The simplicity of the scheme allows to predict its error probability and energy-efficiency analytically with high precision by using methods developed for massive MU-MIMO.

Empirical simulations show that the performance of the presented scheme, despite its simplicity, is comparable to existing approaches designed for large coherence block-lengths.

1.7. Cell-free MIMO

The final part of the thesis treats the problem of combining measurements from several BS. In a conventional massive MU-MIMO system with pilot based uplink a geographical area is divided into cells with one BS at the center of each cell which serves users within the cell only. When all cells share the same set of pilot sequences, users in neighbouring cells may use the same pilot sequence which leads to a coherent corruption of the estimated

channel vectors. This phenomenon is known as pilot contamination and ultimately limits the scalability of massive MU-MIMO systems, since this type of interference cannot be overcome by increasing the number of receive antennas. Also an increase of the cell-density only increases the effect of the pilot contamination. Overcoming the pilot contamination is a major problem in current mobile communication research. One solution is to allocate different sets of pilots to different cells [46], such that neighbouring cells do not share the same set of pilots. This limits the pilot contaminations but reduces the number of available pilots per cell and therefore also the number of supported concurrent users.

The idea of conventional massive MU-MIMO with distributed, instead of collocated, receive antennas which work together cooperatively to serve many users has been studied under the name *cell-free* massive MU-MIMO [76]. And although the distributed systems outperform their centralized counterparts in terms of achievable rates [76], they also come with the increased logistic burden of placing and connecting all the micro BSs. Besides the pilot based uplink treated in [76] still suffers from pilot contamination.

The advantage of the presented CCS U-RA scheme is that it works without pilots and channel estimation and therefore it does not suffer from pilot contamination. Possible collisions of sub-messages within one slot can be efficiently corrected by the outer code. Furthermore, due to the unsourced nature of the scheme it is possible to scale the system by adding more BS into an area without the need for cell management. A simple approach to a multi-BS system that requires almost no cooperation between BSs is to let each BS create a list of received messages individually. Then, a central collector gathers these lists and takes the union to remove duplicate messages.

Finally, I present a more sophisticated approach to AD in the distributed setting which requires the BS to share only the estimated LSFCs. The idea is to use the distant dependent LSFCs to estimate the locations of the active users and use the location estimates to refine the ML estimation of the LSFCs. This process can be iterated until convergence. The accuracy of the position estimation depends mostly on the quality of the assumed propagation model. In this work I assume a commonly used empirical propagation model that has a deterministic pathloss exponent and random log-normal shadowing. It is well known that positioning, based on received signal strengths, has a limited accuracy due to the large effect of the log-normal shadowing [77, 78]. Nonetheless, the described procedure of joint positioning and AD shows a remarkable improvement when geographically close BS are able to exchange their LSFCs estimates. Note, that the log-normal shadowing model is purely empirical. The shadowing component is often a complicated but mostly deterministic function of the surrounding environment. In static scenarios it is possible to create radio maps which describe the pathloss in different locations. Such maps can be generated for a given area by physical measurements, ray-tracing simulations or novel deep-

neural-network based methods [79,80]. If such a map is available it is possible to incorporate it into the likelihood function and improve the accuracy of the position estimation.

1.8. Notation

Scalar constants are represented by non-boldface letters (e.g., x or X), sets by calligraphic letters (e.g., \mathcal{X}), vectors by boldface small letters (e.g., \mathbf{x}), and matrices by boldface capital letters (e.g., \mathbf{X}). The i -th row and the j -th column of a matrix \mathbf{X} are denoted by the row-vector $\mathbf{X}_{i,:}$ and the column-vector $\mathbf{X}_{:,j}$ respectively. A diagonal matrix with elements (s_1, s_2, \dots, s_k) is denoted by $\text{diag}(s_1, \dots, s_k)$. The vectorization operator is denoted by $\text{vec}(\cdot)$. The ℓ_p -norm of a vector \mathbf{x} and the Frobenius norm of a matrix \mathbf{X} are denoted by $\|\mathbf{x}\|_p$ and $\|\mathbf{X}\|_p$ resp. $\|\mathbf{x}\|_0 := |\{i : x_i \neq 0\}|$ denotes the number of non-zero entries of a vector \mathbf{x} . The operator norm of a matrix \mathbf{X} is denoted by $\|\mathbf{X}\|_{op}$. The $k \times k$ identity matrix is represented by \mathbf{I}_k . For an integer $k > 0$, the shorthand notation $[k]$ is used for $\{1, 2, \dots, k\}$. Superscripts $(\cdot)^\top$ and $(\cdot)^H$ are used for transpose and Hermitian transpose. \odot denotes the elementwise product of vectors or matrices of the same size. $\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{x}^H \mathbf{y}$ denotes the Euclidean scalar product between two vectors. Universal constants are defined as numbers, which are independent of all system parameters. Such constants are typically denoted by c, C, c', c_0, c_1 etc., and different universal constants may be denoted by the same letter. $\log(x)$ denotes the natural logarithm of x .

1.9. Publications

The contents of Chapter 2 and Chapter 3 correspond to the preprints [81] and [82] resp., which have been submitted for publication but are still being peer-reviewed at the point of writing. The following conference publications and preprint were written in the course of the dissertation.

- A. Fengler, P. Jung, and G. Caire, “SPARCs and AMP for Unsourced Random Access,” *IEEE Int. Symp. Inf. Theory Proc.*, pp. 2843–2847, July 2019
- A. Fengler, S. Haghighatshoar, P. Jung, and G. Caire, “Grant-Free Massive Random Access With a Massive MIMO Receiver,” *Asilomar Conf. Signals Syst. Comput.*, Nov. 2019
- A. Fengler, P. Jung, and G. Caire, “Unsourced Multiuser Sparse Regression Codes achieve the Symmetric MAC Capacity,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 3001–3006, June 2020
- A. Fengler and P. Jung, “On the Restricted Isometry Property of Centered Self Khatri-Rao Products,” *arXiv:1905.09245*, May 2019
- [submitted] A. Fengler, S. Haghighatshoar, P. Jung, and G. Caire, “Non-Bayesian Activity Detection, Large-Scale Fading Coefficient Estimation, and Unsourced Random Access with a Massive MIMO Receiver,” *ArXiv191011266 Cs Math*, Aug. 2020
- [submitted] A. Fengler, P. Jung, and G. Caire, “SPARCs for Unsourced Random Access,” *ArXiv190106234 Cs Math*, June 2020

2. Multi-User SPARCs for the AWGN Channel

2.1. Outline and Main Contributions

This chapter is organized as follows. Section 2.2 formalizes the U-RA setting on the real AWGN channel without fading. In Section 2.3 the concatenated coding scheme is defined and the decomposition into inner and outer channel is introduced. Section 2.4 defines the inner AMP decoder and the optimal, but uncomputable, inner MAP decoder and analyses their asymptotic error probabilities. Section 2.5 analyses the quantization step, which is necessary for a binary-input outer decoder. Section 2.6 formulates the outer channel and gives converse and achievability results. Section 2.7 analyses the concatenated code. Section 2.8 gives an algorithm to optimize the power allocation and Section 2.9 introduces a low-complexity approximation of the suggested AMP algorithm. In Section 2.10 we give finite-length simulations and compare them to the analytical results. The main contribution in this chapter are as follows

- The concept of sparse regression codes is extended to the unsourced random access setting by making use of the tree code of [24].
- For the resulting inner-outer concatenated coding scheme, we introduce a matching outer channel model and analyse the achievable rates on this outer channel and compare it to existing practical solutions.
- We propose a modified approximate message passing algorithm as an inner decoder and analyse its asymptotic error probability through its SE. We use the connections between SE and the rigorously confirmed RS formula to find the error probability of a hypothetical MAP decoder.
- We find that the error probability of the inner decoder admits a simple closed form in the limit of $K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$ for some $\alpha > 1$. The limit was also considered in [24], motivated by the fact that J is the number of bits required to encode the identity of each of up to K_{tot} users if $K_a = K_{\text{tot}}^{1/\alpha}$. We show that the

per-user error probability of the concatenated scheme vanishes in the limit of large blocklength and infinitely many users, if the sum-rate is smaller than the symmetric Shannon capacity $0.5 \log_2(1 + K_a \text{SNR})$. This shows that an unsourced random access scheme can, even with no coordination between users, achieve the same symmetric rates as a non-unsourced scheme.

- Using the results from the asymptotic analysis we identify parameter regions where the AMP decoder can achieve the same error probability as the MAP decoder. In parameter regions where there is a gap between the achievable error probability of the AMP and the MAP decoder we propose a method for finding an optimal power allocation that is able to improve the performance of the AMP decoder significantly.
- Finite-length simulations show the efficiency of the proposed coding scheme and the accuracy of the analytical predictions.

2.2. Channel Model

Let K_a denote the number of active users, n the number of available channel uses and $B = nR$ the size of a message in bits. The spectral efficiency is given by $\mu = K_a R$. The channel model used is

$$\mathbf{y} = \sum_{i=1}^{K_{\text{tot}}} q_i \mathbf{x}_i + \mathbf{z}, \quad (2.1)$$

where each $\mathbf{x}_i \in \mathcal{C} \subset \mathbb{R}^n$ is taken from a common codebook \mathcal{C} and $q_i \in \{0, 1\}$ are binary variables indicating whether a user is active. The number of active users is denoted as $K_a = \sum_{i=1}^{K_{\text{tot}}} q_i$. The codewords are assumed to be normalized $\|\mathbf{x}_i\|_2^2 = nP$ for a given energy-per-symbol P and the noise vector \mathbf{z} is Gaussian iid $z_i \sim \mathcal{N}(0, N_0/2)$, such that $\text{SNR} = 2P/N_0$ denotes the real per-user SNR. All the active users pick one of the 2^B codewords from \mathcal{C} , based on their message $W_k \in [1 : 2^B]$ and transmit it.¹ The decoder of the system produces a list $g(\mathbf{y})$ of at most K_a messages. An error is declared if one of the transmitted messages is missing in the output list $g(\mathbf{y})$ and we define the per-user probability of error as:

$$P_e = \frac{1}{K_a} \sum_{k=1}^{K_a} \mathbb{P}(W_k \notin g(\mathbf{y})). \quad (2.2)$$

Note that the error is independent of the user identities in general and especially independent of the inactive users. The performance of the system is measured in terms of the

¹Throughout this work, as in [9] and in [87], we assume that users are synchronized.

required $E_b/N_0 := P/(RN_0)$ for a target P_e and the described coding construction is called *reliable* if $P_e \rightarrow 0$ as $n \rightarrow \infty$.

2.3. Concatenated Coding

In this work we focus on a special type of codebook, where each transmitted codeword is created in the following way: First, the B -bit message W_k of user k is mapped to an LJ -bit codeword from some common *outer* codebook. Then each of the J -bit sub-sequences is mapped to an index $i_k(l) \in [1 : 2^J]$ for $l = [1 : L]$ and $k = [1 : K_a]$. The inner codebook is based on a set of L coding matrices $\mathbf{A}_l \in \mathbb{R}^{n \times 2^J}$. Let $\mathbf{a}_i^{(l)}$ with $i = [1 : 2^J]$ denote the columns of \mathbf{A}_l . The inner codeword of user k corresponding to the sequence of indices $i_k(1), \dots, i_k(L)$ is then created as

$$\mathbf{x}_k = \sum_{l=1}^L \sqrt{P_l} \mathbf{a}_{i_k(l)}^{(l)}. \quad (2.3)$$

The columns of \mathbf{A}_l are assumed to be scaled such that $\|\mathbf{a}_i^{(l)}\|_2^2 = 1$ and the power coefficients P_l are chosen such that $\|\mathbf{x}_k\|_2^2 \leq nP$. The above encoding model can be written in matrix form as

$$\mathbf{y} = \sum_{k=1}^{K_a} \mathbf{A} \mathbf{m}_k + \mathbf{z} = \mathbf{A} \left(\sum_{k=1}^{K_a} \mathbf{m}_k \right) + \mathbf{z}. \quad (2.4)$$

where $\mathbf{A} = (\mathbf{A}_1 | \dots | \mathbf{A}_L)$ and $\mathbf{m}_k \in \mathbb{R}^{L2^J}$ is a non-negative vector satisfying $m_{k,(l-1)2^J+i_k(l)} = \sqrt{P_l}$ and zero otherwise, for all $l = [1 : L]$. Let $\boldsymbol{\theta} = \sum_{k=1}^{K_a} \mathbf{m}_k$ and let \mathbf{s} denote the integer part of $\boldsymbol{\theta}$, i.e. $\boldsymbol{\theta} = (\sqrt{P_1} \mathbf{s}^1 | \dots | \sqrt{P_L} \mathbf{s}^L)^\top$. The linear structure allows to write the channel (2.4) as a concatenation of the inner point-to-point channel $\boldsymbol{\theta} \rightarrow \mathbf{A} \boldsymbol{\theta} + \mathbf{z}$ and the outer binary input adder MAC $(\mathbf{m}_1, \dots, \mathbf{m}_{K_a}) \rightarrow \mathbf{s}$. We will refer to those as the *inner* and *outer channel*, the corresponding encoder and decoder will be referred to as *inner* and *outer encoder/decoder* and the aggregated system of inner and outer encoder/decoder as the *concatenated system*. The per-user inner rate in terms of bits/c.u. is given by $R_{\text{in}} := LJ/n$ and the outer rate is given by $R_{\text{out}} = B/LJ$. For the analysis we assume that all entries of the coding matrices \mathbf{A}_l are Gaussian iid and that the outer encoded indices $i_k(l)$ are distributed uniformly and independently in $[1 : 2^J]$. Furthermore, we assume that the power coefficients are uniformly $P_l \equiv nP/L$, such that the power constraint $\mathbb{E}[\|\mathbf{x}_k\|_2^2] = nP$ is fulfilled on average. This latter assumption will be relaxed in Section 2.8, where we will consider a non-uniform power allocation.

2.4. Inner Channel

In this section we focus on the inner decoding problem of recovering \mathbf{s} from

$$\mathbf{y} = \mathbf{A}\boldsymbol{\theta} + \mathbf{z} = \sqrt{\hat{P}}\mathbf{A}\mathbf{s} + \mathbf{z} \quad (2.5)$$

where $\hat{P} = nP/L$. Let $k_i \in [0 : K_a]$ for $i \in [1 : 2^J]$ be non-negative integers. The probability of observing a specific \mathbf{s}^l is given by:

$$p\left(\mathbf{s}^l = (k_1, \dots, k_{2^J})^\top\right) = 2^{-K_a J} \frac{K_a!}{k_1! \cdots k_{2^J}!} \quad (2.6)$$

if $\sum_{i=1}^{2^J} k_i = K_a$ and zero otherwise. This is a multinomial distribution with uniform event probabilities. The marginals of such a distribution are known to be Binomial, i.e.:

$$p_k := \mathbb{P}(s_i^l = k) = \binom{K_a}{k} 2^{-K_a J} (1 - 2^{-J})^{K_a - k} \quad (2.7)$$

and specifically, the probability of observing a zero is:

$$p_0 := \mathbb{P}(s_i^l = 0) = (1 - 2^{-J})^{K_a}. \quad (2.8)$$

We define two estimators for \mathbf{s} . The first is a variant of the approximate message passing (AMP) algorithm, which we will refer to as AMP-estimator. An estimate of \mathbf{s} is obtained by iterating the following equations:

$$\begin{aligned} \boldsymbol{\theta}^{t+1} &= f_t(\mathbf{A}^\top \mathbf{z}^t + \boldsymbol{\theta}^t) \\ \mathbf{z}^{t+1} &= \mathbf{y} - \mathbf{A}\boldsymbol{\theta}^{t+1} + \frac{2^J L}{n} \mathbf{z}^t \langle f'_t(\mathbf{A}^\top \mathbf{z}^t + \boldsymbol{\theta}^t) \rangle \end{aligned} \quad (2.9)$$

where the functions $f_t : \mathbb{R}^{2^J L} \rightarrow \mathbb{R}^{2^J L}$ are defined componentwise $f_t(\mathbf{x}) = (f_{t,1}(x_1), \dots, f_{t,2^J L}(x_{2^J L}))^\top$ and each component is given by

$$f_{t,i}(x) = \frac{\sqrt{\hat{P}}}{Z(x)} \sum_{k=0}^{K_a} p_k k \exp\left(\frac{1}{2\tau_t^2} (x - k\sqrt{\hat{P}})^2\right) \quad (2.10)$$

with $\tau_t^2 = \|\mathbf{z}^t\|_2^2/n$, p_0 as in (2.8) and

$$Z(x) = \sum_{k=0}^{K_a} p_k \exp\left(\frac{1}{2\tau_t^2} (x - k\sqrt{\hat{P}})^2\right). \quad (2.11)$$

$\langle \mathbf{x} \rangle = (\sum_{i=1}^N x_i)/N$ in (2.9) denotes the average of a vector, f'_t denotes the component-wise derivative of f_t and $\mathbf{s}^0 = \mathbf{0}$ is chosen as the initial value. After the equations (2.9) are iterated for some fixed amount of iterations T_{\max} , a final estimate of \mathbf{s} is obtained by quantizing $\boldsymbol{\theta}^{T_{\max}}$ to the nearest integer multiple of $\sqrt{\hat{P}}$ and dividing by $\sqrt{\hat{P}}$.

The second estimator that we analyse is the symbol-by-symbol maximum-a-posteriori (SBS-MAP) estimator of \mathbf{s}

$$\hat{s}_i = \arg \max_{s \in [0:K_a]} \mathbb{P}(s_i = s | \mathbf{y}, \mathbf{A}), \quad (2.12)$$

which minimizes the SBS error probability $\mathbb{P}(\hat{s}_i \neq s_i)$ but is unfeasible to compute in practice. Let

$$P_e^{\text{MAP}} = \frac{1}{2^J} \sum_{j=1}^{2^J} \mathbb{P}(\hat{s}_j \neq s_j) \quad (2.13)$$

denote the error rate of the SBS-MAP estimator and P_e^{AMP} the error rate of the AMP estimator.

2.4.1. Asymptotic Error Analysis

The error analysis is based on the self-averaging property of the random linear recovery problem (2.5) in the asymptotic limit $L, n \rightarrow \infty$ with a fixed J and fixed R_{in} . That is, although \mathbf{A}, \mathbf{s} and \mathbf{z} are random variables, the error probability of both mentioned estimators converges sharply to its average value. The convergence behavior is fully characterized by the external parameters $J, R_{\text{in}}, \text{SNR}$ and K_a . Let $S_{\text{in}} = K_a R_{\text{in}}$ denote the inner sum-rate and $\mathcal{E}_{\text{in}} = P/(R_{\text{in}} N_0)$ the energy-per-coded-bit in the inner channel. The main theoretical result about the inner channel is to show that the $(S_{\text{in}}, \mathcal{E}_{\text{in}})$ plane consists of three regions, divided by two curves, see Figure 2.4 and Figure 2.5. The two curves are called the *optimal threshold* $\mathcal{E}_{\text{in}}^{\text{opt}}(S_{\text{in}})$ and the *algorithmic threshold* $\mathcal{E}_{\text{in}}^{\text{alg}}(S_{\text{in}})$ (or equivalently $S_{\text{in}}^{\text{opt}}(\mathcal{E}_{\text{in}})$ and $S_{\text{in}}^{\text{alg}}(\mathcal{E}_{\text{in}})$). The operative meaning of these thresholds is given by the following theorem.

Theorem 1. *Let $L, n \rightarrow \infty$ for fixed J and fixed R_{in} . Let $\epsilon > 0$ be some fixed error rate in the inner channel. If S_{in} and \mathcal{E}_{in} are such that $\mathcal{E}_{\text{in}} > \mathcal{E}_{\text{in}}^{\text{opt}}(S_{\text{in}})$, then $P_e^{\text{MAP}} < \epsilon$ and if $\mathcal{E}_{\text{in}} > \mathcal{E}_{\text{in}}^{\text{alg}}(S_{\text{in}})$ then $P_e^{\text{AMP}} < \epsilon$.*

□

Since the SBS-MAP estimator minimizes the error probability for all parameters it holds that $\mathcal{E}_{\text{in}}^{\text{opt}} \leq \mathcal{E}_{\text{in}}^{\text{alg}}$. In the remainder of the section we establish the existence of the thresholds and calculate them.

The following Theorem was first discovered heuristically, using the non-rigorous replica method, in [37] for binary iid signals \mathbf{s} and generalized to arbitrary iid \mathbf{s} in [33, 88]. It states

that asymptotically the error probability of various estimators, including the SBS-MAP estimator (2.12) [40], can be described by the error probability in a *decoupled* Gaussian channel model with an effective SNR that is obtained as the minimizer of a certain *potential function*. Named after the method of discovery, this function is termed the *replica symmetric* (RS) potential. This heuristic discovery was later confirmed rigorously for iid signals in [43, 44]. The proof in [44, Ch. 4] is based on an adaptive interpolation approach and holds for the case of block iid signals as in the model (2.6):

Theorem 2. *In the limit $n, L \rightarrow \infty$ for fixed J and fixed $R_{in} = JL/n$ the error probability P_e^{MAP} of the SBS-MAP estimator (2.12) converges to the error probability of the SBS-MAP estimator of \mathbf{s}^l in the Gaussian vector channel*

$$\mathbf{r}^l = \sqrt{\eta \hat{P}} \mathbf{s}^l + \mathbf{z}^l \quad (2.14)$$

where $\mathbf{s}^l \in \mathbb{R}^{2^J}$ is distributed according to $p(\mathbf{s}^l)$, specified in (2.6), the distribution of a single section of \mathbf{s} and $\mathbf{z}^l \in \mathbb{R}^{2^J}$ is Gaussian iid with $\mathcal{N}(0, 1)$ components, independent of \mathbf{s}^l . The factor $\eta \geq 0$ is given as the global minimizer of

$$i^{RS}(\eta) = I_{2^J}(\eta \hat{P}) + \frac{2^J}{2\beta} [(\eta - 1) \log_2(e) - \log_2(\eta)] \quad (2.15)$$

where $I_{2^J}(\eta \hat{P})$ denotes the mutual information between \mathbf{r}^l and \mathbf{s}^l in the Gaussian vector channel (2.14) and $\beta = 2^J L/n = R_{in} 2^J / J$ is the aspect ratio of the matrix \mathbf{A} in (2.5). \square

Proof. See [44, Ch. 4]. \square

It is known that the asymptotic estimation error of the AMP algorithm can be analyzed by the so called *state evolution* (SE) equations [34]. Using the SE result a statement similar to Theorem 2 can be made about the AMP algorithm (2.9):

Theorem 3. *In the limit $n, L \rightarrow \infty$ for fixed J and fixed R_{in} the MSE of the AMP estimate (2.9) converges to the MSE of estimating s in the scalar Gaussian channel*

$$r = \sqrt{\eta \hat{P}} s + z \quad (2.16)$$

where s is distributed according to the binomial distribution $p(s = k)$, specified in (2.7), the marginal distribution of a single section of \mathbf{s} and $z \sim \mathcal{N}(0, 1)$ is Gaussian iid, independent of s . The factor $\eta \geq 0$ is given as the smallest local minimizer of

$$i_{AMP}^{RS}(\eta) = 2^J I(\eta \hat{P}) + \frac{2^J}{2\beta} [(\eta - 1) \log_2(e) - \log_2(\eta)] \quad (2.17)$$

where $I(\eta\hat{P})$ denotes the mutual information between r and s in the scalar Gaussian channel (2.16). \square

Proof. The theorem is merely a restatement of the SE result in [34]. To see this, set the derivative of (2.17) with respect to η to zero. The derivative can be obtained by using the I-MMSE theorem [89], which states that:

$$\frac{1}{\log_2(e)} \frac{d}{d\eta} I(\eta\hat{P}) = \frac{\hat{P}}{2} \text{mmse}(\eta\hat{P}) \quad (2.18)$$

This gives the following condition for a local minimum:

$$\eta^{-1} = 1 + \beta\hat{P}\text{mmse}(\eta\hat{P}) \quad (2.19)$$

where $\text{mmse}(\eta\hat{P})$ is the minimum-mean-square error (MMSE) in estimating s given an observation r that is jointly distributed with s according to the scalar Gaussian channel (2.16). It is well known that the MMSE is achieved by the posterior-mean estimator (PME) $\mathbb{E}[s|y]$. With the substitution $\tau^2 = \eta^{-1}$ it is apparent that (2.19) resembles the fixpoint condition for the SE of AMP [34, Eq. 1.4] if the PME is used as a componentwise denoising function. It is apparent from the definition in (2.10) that f_t was chosen precisely as the PME of $\sqrt{\hat{P}}s$ in a scalar Gaussian channel of the form (2.16) with $\eta = 1/\tau_t^2$. \square

Note, that the assumed distribution on \mathbf{s} is not iid. However, the AMP algorithm in (2.9) uses a separable denoiser. In this case the SE result of [34] does not require \mathbf{s} to be iid, but only that its empirical marginal distributions converge to some limit, which is required for the calculation of the SE. In fact, the presented AMP algorithm (2.9) is not optimal since it does not make full use of the distribution of \mathbf{s} , i.e. it ignores the correlation among different components of \mathbf{s} within a section. The optimal AMP algorithm would use the PME of \mathbf{s} in the Gaussian vector channel (2.14) as a denoiser, which, however, is unfeasible for large values of J and K_a . Nonetheless, the two potential functions (2.15) and (2.17) are strongly related and in the following we show that the componentwise PME in (2.10) is the best componentwise approximation of the vector PME and furthermore, in the typical sparse setting, i.e. $K_a \ll 2^J$, the difference between the minima of the two potential functions (2.15) and (2.17) is negligibly small and therefore the global minimizer of the scalar potential (2.17) can be used to approximate the global minimizer of (2.15). This means that the error probability of both, the SBS-MAP estimator, and the presented AMP estimator can be characterized by the local and global minimizers of a single scalar potential function. This idea is made precise in the following Theorem:

Theorem 4. Let η_{opt} be the global minimizer of (2.15) and let $\tilde{\eta}_{opt}$ be the global minimizer of (2.17) then $\eta_{opt} > \tilde{\eta}_{opt}$ and

$$\eta_{opt} - \tilde{\eta}_{opt} = \mathcal{O}\left(\frac{\log K_a}{2^{(1-\delta)J}}\right) \quad (2.20)$$

for some $0 < \delta < 1$ □

Proof. Analog to (2.19) the optimality condition for η_{opt} can be found by setting the derivative of (2.15) to zero and using the I-MMSE theorem for a Gaussian vector channel. This gives the condition

$$\eta^{-1} = 1 + \beta \hat{P} \frac{\text{mmse}_{2^J}(\eta \hat{P})}{2^J} \quad (2.21)$$

where $\text{mmse}_{2^J}(\eta \hat{P})$ is the MMSE of estimating \mathbf{s}^l in the Gaussian vector channel (2.14). The mismatched MSE function for an arbitrary probability distribution $q : [0 : K_a]^{2^J} \rightarrow [0, 1]$ is defined as

$$\text{mse}_q(t) = \mathbb{E} \|\mathbf{s} - \hat{\mathbf{s}}_q(\sqrt{t}\mathbf{s} + \mathbf{Z}, t)\|_2^2 \quad (2.22)$$

with

$$\hat{\mathbf{s}}_q(\mathbf{r}, t) = \sum_{\mathbf{s} \in [0:K_a]^{2^J}} \mathbf{s} \frac{\exp(-\|\mathbf{r} - \sqrt{t}\mathbf{s}\|_2^2/2)q(\mathbf{s})}{\sum_{\mathbf{s}'} \exp(-\|\mathbf{r} - \sqrt{t}\mathbf{s}'\|_2^2/2)q(\mathbf{s}')} \quad (2.23)$$

The expression in (2.22) is the MSE of a (mismatched) PME in a Gaussian vector channel of the form (2.14) with respect to some prior distribution $q(\mathbf{s})$, which may differ from the true prior $p_{\mathbf{s}}(\mathbf{s})$. It is clear that from the minimality of the MMSE function that $\text{mmse}_{2^J}(t) \leq \text{mse}_q(t)$ for all t with equality if and only if $q(\mathbf{s}) = p_{\mathbf{s}}(\mathbf{s})$ almost everywhere. This means, that calculating the fixed-point of (2.21), with $\text{mmse}_{2^J}(t)$ replaced by $\text{mse}_q(t)$ for any choice of $q(\mathbf{s})$ gives an upper bound on η_{opt} . The L^1 -distance between the functions $\text{mmse}_{2^J}(t)$ and $\text{mse}_q(t)$ is quantified by the following result from [90]

$$\frac{1}{2} \|\text{mse}_q - \text{mmse}_{2^J}\|_{L^1} = \frac{1}{2} \int_0^\infty [\text{mse}_q(t) - \text{mmse}_{2^J}(t)] dt = D(p_{\mathbf{s}} \parallel q). \quad (2.24)$$

where $D(p_{\mathbf{s}} \parallel q)$ denotes the KL-divergence between the distributions $p_{\mathbf{s}}$ and q . We focus on product distributions of the form $q(\mathbf{s}) = \prod_{i=1}^{2^J} q_i(s_i)$, since for such distributions the vector MSE function (2.22) becomes the sum of scalar MSE functions, which are easy to calculate. Moreover, a simple calculation in Appendix A shows that the L^1 distance (2.24) is minimized by the product distribution whose factors q_i match the marginals of $p_{\mathbf{s}}$. We prove in Appendix B that the KL-divergence between the multinomial distribution $p_{\mathbf{s}}$ and

the product distribution of its marginals satisfies

$$D \left(p_{\mathbf{s}} \left\| \prod_{i=1}^{2^J} p_i(s_i) \right\| \right) = \mathcal{O}(\log K_a) \quad (2.25)$$

Since all marginals of $p_{\mathbf{s}}$ are identical and given by the binomial distribution the mismatched MSE function of $q(\mathbf{s}) = \prod p_i(s_i)$ takes the form

$$\text{mse}_q(t) = 2^J \text{mmse}(t) \quad (2.26)$$

where $\text{mmse}(t)$ is the MMSE function in the scalar Gaussian channel (2.16) that appears in (2.19). So it follows from (2.24) and (2.25) that

$$\left\| \text{mmse} - \frac{\text{mmse}_{2^J}}{2^J} \right\|_{L^1} = \mathcal{O} \left(\frac{\log K_a}{2^J} \right) \quad (2.27)$$

Therefore, with $J \rightarrow \infty$, the difference between the per-component vector MMSE function in (2.21) and the scalar MMSE function in (2.19) converges exponentially fast in L^1 norm to zero. We show in Appendix C that exponentially fast convergence in L^1 -norm implies exponentially fast pointwise uniform convergence almost everywhere. Since MMSE functions of Gaussian channels are smooth [91, Proposition 7] the 'almost' can be dropped. We have established that the difference of the right hand sides of the fixpoint equations (2.19) and (2.21) converges pointwise to zero. Due to the smoothness of the MMSE functions the convergence carries over to the solutions of (2.19) and (2.21) which implies (2.20). \square

Theorem 4 allows to calculate the minima of the potential (2.15) by solving the scalar fixpoint equation (2.19) numerically. It guarantees that the error will be small, since typically 2^J is much larger than K_a . Nonetheless, the MMSE function in equation (2.19) can be further simplified if J grows large. That is because the coefficients p_k defined in (2.7), which can be expressed as

$$p_k = p_0 \frac{\binom{K_a}{k}}{(2^J - 1)^k} \quad (2.28)$$

decay exponentially with kJ . This suggests that for large J all p_k with $k \geq 2$ can be dropped:

Theorem 5. *Let $\text{mmse}_{OR}(t)$ be the MMSE function of estimating the binary variable $s_{OR} \in \{0, 1\}$ in the scalar Gaussian channel*

$$r = \sqrt{t} s_{OR} + z \quad (2.29)$$

2. Multi-User SPARCs for the AWGN Channel

where $p(s_{OR} = 0) = p_0$ and $p(s_{OR} = 1) = 1 - p_0$ and $z \sim \mathcal{N}(0, 1)$ independent of s_{OR} . Then $\text{mmse}_{OR}(t) \geq \text{mmse}(t)$ for all $t > 0$ and

$$\text{mmse}_{OR}(t) - \text{mmse}(t) = \mathcal{O}\left(\frac{K_a^2}{2^{2J}}\right) \quad (2.30)$$

□

Proof. See Appendix D. □

We choose the nomenclature OR in mmse_{OR} because the distribution of s_{OR} arises as the marginal distribution of \mathbf{s} if the OR-sum of the \mathbf{m}_k

$$\mathbf{s} = \bigvee_{k=1}^{K_a} \mathbf{m}_k \quad (2.31)$$

is used instead of conventional real summation. Furthermore, let $I_{OR}(t)$ be the input-output mutual information in the channel (2.29) and

$$i_{J,OR}^{\text{RS}}(\eta) = I_{OR}(\eta\hat{P}) + \frac{2^J}{2\beta}[(\eta - 1)\log_2(e) - \log_2(\eta)] \quad (2.32)$$

the corresponding RS-potential. A consequence of Theorems 4 and 5 is that (2.32) can be used to find the global minimizer of (2.15) and the local minimizer of (2.17).

Corollary 1. *Let η_{opt} be the global minimizer of (2.15), let $\tilde{\eta}_{opt}^{OR}$ be the global minimizer of (2.32), let η_{alg} be the smallest local minimizer of (2.17) and let $\tilde{\eta}_{alg}^{OR}$ be the smallest local minimizer of (2.32), then $\eta_{opt} > \tilde{\eta}_{opt}^{OR}$ and*

$$\eta_{opt} - \tilde{\eta}_{opt}^{OR} = \mathcal{O}\left(\frac{\log K_a}{2^{(1-\delta)J}}\right) \quad (2.33)$$

Furthermore, if $\eta_{opt} \neq \eta_{alg}$ then $\eta_{alg} > \tilde{\eta}_{alg}^{OR}$ and:

$$\eta_{alg} - \tilde{\eta}_{alg}^{OR} = \mathcal{O}\left(\frac{\log K_a}{2^{(1-\delta)J}}\right) \quad (2.34)$$

for some $0 < \delta < 1$. □

Note that we have only shown that the difference of the MMSE functions in the channels (2.14) and (2.29) converges to zero as $J \rightarrow \infty$. This shows that those functions converge to the same limiting function. The derivation of this limiting function itself is the subject of the following Section 2.4.2.

2.4.2. The $J \rightarrow \infty$ limit

The problem with the numerical evaluation of (2.19), even with the simplified mmse_{OR} function, is that 2^{-J} is very small and the scalar potential (2.17) is hard to calculate even for moderately large J due to machine precision errors. So even though (2.30) guarantees that the common $J \rightarrow \infty$ limit of the two MMSE functions in the channels (2.29) and (2.14) exists, it is not obvious how to calculate it numerically. To solve these problems we calculate the limit of (2.32) analytically in a regime where both K_a and J go to infinity with a fixed ratio $\alpha := J/\log_2 K_a$ for some $\alpha > 1$. The parameter α determines the sparsity in the vector \mathbf{s} , e.g. $K_a = 300$ and $J = 15$ gives $\alpha \sim 1.82$. In this limit $K_a/2^J = K_a^{1-\alpha} \rightarrow 0$, i.e. the sparsity in \mathbf{s} goes to zero and the error term in Theorem 5 vanishes. We find that a non-trivial limit of the MMSE function exists in the *energy-efficient* regime, i.e. for $R_{\text{in}}, P \rightarrow 0$ with fixed sum-rate $S_{\text{in}} = K_a R_{\text{in}}$ and fixed energy-per-coded-bit $\mathcal{E}_{\text{in}} = \text{SNR}/(2R_{\text{in}})$. It was already established in the proof of Theorem 4 and 5 that the derivatives of the potential functions (2.15) and (2.32) converge uniformly to a common limit almost everywhere. By a standard result about uniform convergence, e.g. [92, Theorem 7.17], it is enough to calculate the pointwise limit $i_{\infty}^{\text{RS}}(\eta) := \lim_{J \rightarrow \infty} i_{J, \text{OR}}^{\text{RS}}(\eta)$ to guarantee that the derivatives of $i_{J, \text{OR}}^{\text{RS}}(\eta)$ will converge to the derivative of $i_{\infty}^{\text{RS}}(\eta)$ at all continuity points of $i_{\infty}^{\text{RS}}(\eta)$.

Theorem 6. *In the limit $K_a, J \rightarrow \infty$, $R_{\text{in}}, \text{SNR} \rightarrow 0$ with fixed \mathcal{E}_{in} , S_{in} and $J = \alpha \log_2 K_a$ for some $\alpha > 1$ the pointwise limit of the RS-potential (2.32) is given by (up to additive or multiplicative terms that are independent of η and therefore do not influence the critical points of $i^{\text{RS}}(\eta)$):*

$$i_{\infty}^{\text{RS}}(\eta) := \lim_{J \rightarrow \infty} i_{J, \text{OR}}^{\text{RS}}(\eta) = \eta S_{\text{in}} \mathcal{E}_{\text{in}} [1 - \theta(\eta - \bar{\eta})] + \frac{S_{\text{in}}}{\log_2 e} \left(1 - \frac{1}{\alpha} \right) \theta(\eta - \bar{\eta}) + \frac{1}{2} [(\eta - 1) - \ln \eta] \quad (2.35)$$

where

$$\theta(x) := \begin{cases} 1, & \text{if } x > 0 \\ \frac{1}{2}, & \text{if } x = 0 \\ 0, & \text{if } x < 0 \end{cases} \quad (2.36)$$

and

$$\bar{\eta} = \frac{1 - \frac{1}{\alpha}}{\mathcal{E}_{\text{in}} \log_2 e} \quad (2.37)$$

□

Proof. See Appendix E. □

The critical points of $i_\infty^{\text{RS}}(\eta)$ can then be calculated analytically, resulting in simple conditions:

Theorem 7. $\eta^* = 1$ is a global minimizer of $i_\infty^{\text{RS}}(\eta)$, if and only if

$$S_{\text{in}} \left(1 - \frac{1}{\alpha}\right) < \frac{1}{2} \log_2(1 + 2S_{\text{in}}\mathcal{E}_{\text{in}}) \quad (2.38)$$

and $\eta_{\text{loc}}^* = (1 + 2S_{\text{in}}\mathcal{E}_{\text{in}})^{-1}$ is a local minimizer of $i_\infty^{\text{RS}}(\eta)$ if and only if

$$2S_{\text{in}} \geq \log_2 e \left(1 - \frac{1}{\alpha}\right)^{-1} - \frac{1}{\mathcal{E}_{\text{in}}} \quad (2.39)$$

□

Proof. According to Theorem 6 the derivative of $i_\infty^{\text{RS}}(\eta)$ in (2.35) is given by

$$\frac{\partial i_\infty^{\text{RS}}}{\partial \eta}(\eta) = S_{\text{in}}\mathcal{E}_{\text{in}}[1 - \theta(\eta - \bar{\eta})] + \frac{1}{2} \left(1 - \frac{1}{\eta}\right) \quad (2.40)$$

for $\eta \neq \bar{\eta}$. The critical points of the derivative are

$$\eta_0^* = (1 + 2S_{\text{in}}\mathcal{E}_{\text{in}})^{-1} \quad (2.41)$$

and

$$\eta_1^* = 1. \quad (2.42)$$

The first point η_0^* is critical if and only if $\eta_0^* < \bar{\eta}$, which, after rearranging, gives precisely condition (2.39). Also note, that the second derivative of i_∞^{RS} is $(4\eta)^{-2}$, so it is non-negative for all $\eta > 0$. Therefore the critical points are indeed minima. A local maximum may appear only at $\eta = \bar{\eta}$ where i_∞^{RS} is not differentiable. The values of i_∞^{RS} at the minimal points are

$$\begin{aligned} i_\infty^{\text{RS}}(\eta_0^*) &= \frac{S_{\text{in}}\mathcal{E}_{\text{in}}}{1 + 2S_{\text{in}}\mathcal{E}_{\text{in}}} + \frac{1}{2} \left[\frac{-2S_{\text{in}}\mathcal{E}_{\text{in}}}{1 + 2S_{\text{in}}\mathcal{E}_{\text{in}}} + \ln(1 + 2S_{\text{in}}\mathcal{E}_{\text{in}}) \right] \\ &= \frac{\log_2(1 + 2S_{\text{in}}\mathcal{E}_{\text{in}})}{2 \log_2 e} \end{aligned} \quad (2.43)$$

if $\eta_0^* < \bar{\eta}$, and

$$i_\infty^{\text{RS}}(\eta_1^*) = \frac{S_{\text{in}}}{\log_2 e} \left(1 - \frac{1}{\alpha}\right) \quad (2.44)$$

It is apparent that $i_{\infty}^{\text{RS}}(\eta_1^*)$ is the global minimum if and only if condition (2.38) is fulfilled. It was implicitly used here that $\bar{\eta} \leq 1$, which is justified because condition (2.38) implies $\bar{\eta} \leq 1$. This can be seen by using the inequality $\ln(1+x) \leq x$ for $x > 0$ with (2.38). \square

2.5. Hard Decision

Theorem 2 shows that the asymptotic error probability of the SBS-MAP detector can be calculated from the vector Gaussian channel (2.14). Since the detection is performed symbol-wise and all marginals of \mathbf{s}^l are identical the problem is reduced to MAP detection in the scalar Gaussian channel (2.16), which by Theorem 3 also describes the error distribution of the AMP-estimator although with a different effective channel strength. The previous section described in depth how the effective channel strengths can be obtained. In this section we consider the problem of deciding between $s = 0$ and $s \geq 1$ from a Gaussian observation r , specified by (2.16), with a fixed channel strength $\eta\hat{P}$. We focus only on the support information for two reasons. The first reason is that all known practical outer codes only make use of the support information. The second reason is that in the typical setting where $K_a \ll 2^J$ collisions are so rare that the error-per-component of treating each component as if collisions are impossible is negligibly small. Let $\hat{\rho}$ be an estimate of $\mathbb{1}(s \geq 1)$ given a observation of s in Gaussian noise according to (2.16). We define two types of errors, the probability of missed detections (Type I errors)

$$p_{\text{md}} = p(\hat{\rho} = 0 | s \geq 1) \quad (2.45)$$

and the probability of false alarms (Type II errors)

$$p_{\text{fa}} = p(\hat{\rho} = 1 | s = 0). \quad (2.46)$$

According to Neyman-Pearson the optimal trade-off between the two types of errors is achieved by choosing $\hat{\rho} = 1$ whenever

$$\frac{p(s \geq 1 | r)}{p(s = 0 | r)} \geq \theta, \quad (2.47)$$

where θ is some appropriately chosen threshold. If s takes on only binary values with $p(s = 1) = 1 - p_0$ and $p(s > 1) = 0$, as in the OR-approximation introduced in Theorem 5, a straightforward calculation shows that by varying θ the trade-off between p_{md} and p_{fa} follows the curve defined by the equation

$$Q^{-1}(p_{\text{md}}) + Q^{-1}(p_{\text{fa}}) = \sqrt{\eta\hat{P}} \quad (2.48)$$

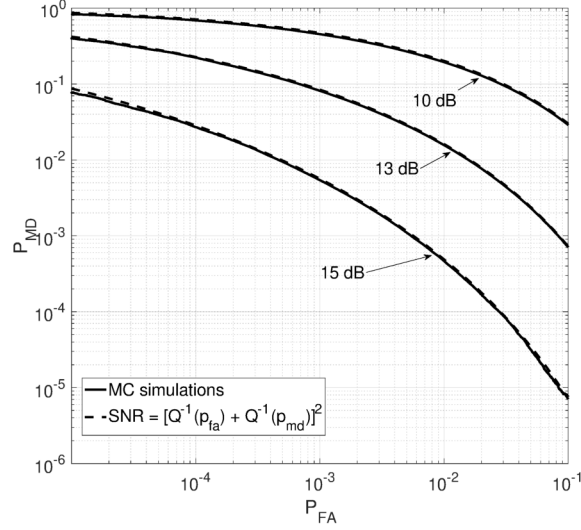


Figure 2.1.: Detection error trade-off for $K_a = 300$ and $J = 12$. The solid curves are created by Monte-Carlo simulations of the channel (2.16) with subsequent likelihood-ratio test, where the threshold θ is varied. The dashed lines are calculated by (2.48).

where $Q(x)$ denotes the Q -function. In Figure 2.1 we plot this curve for $K_a = 300$, $J = 12$ and various values of $\eta\hat{P}$ together with the curves obtained from the precise evaluation of the Neyman-Pearson error probabilities without the OR-approximation. It is apparent that the difference between those two is barely recognizable. The choice $\theta = 1$ gives the SBS-MAP estimator, which minimizes the total error probability

$$p(\hat{\rho} \neq \rho) = p_{fa}p_0 + p_{md}(1 - p_0). \quad (2.49)$$

Nonetheless, it is useful to vary the threshold θ to balance false alarms and missed detections in a way that is adapted to the outer decoder.

2.6. Outer Channel

Assume that the inner decoder (SBS-MAP or AMP) has recovered the support of $\mathbf{s} = \sum_{k=1}^{K_a} \mathbf{m}_k$ such that the symbol-wise error probabilities are given by p_{md} and p_{fa} . It was shown in the previous two sections how p_{md} and p_{fa} can be found in the asymptotic limit.

The support of \mathbf{s} is given by the OR-sum of the messages:

$$\boldsymbol{\rho} = \text{supp}(\mathbf{s}) = \bigvee_{k=1}^{K_a} \mathbf{m}_k \quad (2.50)$$

Let $\hat{\boldsymbol{\rho}}$ denote the estimated support vector. It can be interpreted as the output of L uses of a vector OR-MAC with symbol-wise, asymmetric noise. Let the lists of active indices in the l -th section be

$$\mathcal{S}_l = \left\{ i \in [1 : 2^J] : \hat{\rho}_i^l = 1 \right\}. \quad (2.51)$$

Given L such lists, the outer decoder is tasked to recover the list of transmitted messages up to permutation. An outer code is a subset $\mathcal{C} \subset [1 : 2^J]^L$ of $|\mathcal{C}| = 2^{JLR_{\text{out}}}$ codewords. An outer codeword can be equivalently represented as either a tuple (c_1, \dots, c_L) of L indices in $[1 : 2^J]$ or as a binary vector in \mathbb{R}^{L2^J} with a single one in each section of size 2^J .

Classical code constructions for the OR-MAC, like [93, 94], have been focussed on zero-error decoding, which does not allow for finite per-user-rates as $K_a \rightarrow \infty$, see e.g. [95] for a recent survey. Capacity bounds for the OR-MAC under the given input constraint have been derived in [96] and [97], where it was called the ‘‘T-user M-frequency noiseless MAC without intensity information’’ or ‘‘A-channel’’. An asynchronous version of this channel was studied in [98]. Note, that the capacity bounds in the literature are combinatorial and hard to evaluate numerically for large numbers of K_a and 2^J . In the following we will show that, in the typical case of $K_a \ll 2^J$, a simple upper bound on the achievable rates based on the componentwise entropy is already tight.

From the channel coding theorem for discrete memoryless channels [12] it is known that a code with per-user-rate R_{out} [bits/coded bits] and an arbitrary small error probability exists if and only if

$$R_{\text{out}} < \frac{I(\mathbf{m}_1, \dots, \mathbf{m}_{K_a}; \hat{\boldsymbol{\rho}})}{JK_a}. \quad (2.52)$$

The coding theorem assumes that each user has his own codebook, so the resulting rate constraint as an upper bound on the achievable rates of an outer code with the same codebook constraint. The mutual information is:

$$I(\mathbf{m}_1, \dots, \mathbf{m}_{K_a}; \hat{\mathbf{s}}) = H(\hat{\mathbf{s}}) - H(\hat{\mathbf{s}} | \mathbf{m}_1, \dots, \mathbf{m}_{K_a}) \quad (2.53)$$

where

$$H(\hat{\boldsymbol{\rho}} | \mathbf{m}_1, \dots, \mathbf{m}_{K_a}) = 2^J (p_0 \mathcal{H}_2(p_{\text{fa}}) + (1 - p_0) \mathcal{H}_2(p_{\text{md}})) \quad (2.54)$$

and $\mathcal{H}_2(\cdot)$ denotes the binary entropy function. The output entropy $H(\hat{\mathbf{s}})$ for general asymmetric noise is hard to compute. A simple upper bound on the entropy of the 2^J -ary

vector OR-channel can be obtained by the sum of the marginal entropies of 2^J independent binary-input binary-output channels. If the coded messages are uniformly distributed, i.e. $\mathbb{P}(\mathbf{m}_k^l = \mathbf{e}_j) = 2^{-J}$ for all $j = [1 : 2^J]$, then for all i

$$p(\rho_i = 1) = 1 - p_0 \quad (2.55)$$

with p_0 given in (2.8) and

$$p(\hat{\rho}_i = 1) = (1 - p_0)(1 - p_{\text{md}}) + p_0 p_{\text{fa}} \quad (2.56)$$

Therefore, after reordering:

$$H(\hat{\mathbf{s}}) \leq 2^J (\mathcal{H}_2((1 - p_0)(1 - p_{\text{fa}} - p_{\text{md}}) + p_{\text{fa}})) \quad (2.57)$$

Technically, this is only an upper bound, but we find numerically that it is very tight and furthermore, in the next section we will show that in the noiseless case it is actually achievable by an explicit outer code in the already familiar limit $K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$. To find the limit of (2.57) is assumed that $p_{\text{fa}} \leq c K_a / 2^J = c K_a^{1-\alpha}$ for some constant $c > 0$, i.e. the ratio of false positives to true positives remains at most constant as $K_a, J \rightarrow \infty$. An equivalent condition is

$$\lim_{K_a \rightarrow \infty} \frac{\log_2 p_{\text{fa}}}{\log_2 K_a} \leq 1 - \alpha \quad (2.58)$$

If this is not fulfilled the false positives dominate the entropy terms in the mutual information and the achievable rates go to zero. For small arguments, the binary entropy function becomes

$$\mathcal{H}_2(p) \approx p(1 - \log_2 p) \quad (2.59)$$

and $(1 - p_0) \approx K_a / 2^J$. With this, a straightforward calculation shows that

$$\lim_{K_a, J \rightarrow \infty} \frac{I(\mathbf{m}_1, \dots, \mathbf{m}_{K_a}; \hat{\mathbf{s}})}{J K_a} \leq (1 - p_{\text{md}}) \left(1 - \frac{1}{\alpha}\right) \quad (2.60)$$

For now we assume for simplicity that the inner decoder works error free, i.e. $p_{\text{fa}} = p_{\text{md}} = 0$. Interestingly, the bound $1 - \alpha^{-1}$ is achievable by a random code with a *cover decoder*, a construct often used in group testing literature. Given $\text{OR}(\mathcal{L})$, the OR-combination of \mathcal{L} , a list of K_a codewords, the cover decoder goes through the whole codebook and produces a list of codewords that are covered by $\text{OR}(\mathcal{L})$. By construction the cover decoder will find all codewords in \mathcal{L} and the error probability is governed by the number of false positives n_{fa} . I assume that if the decoder finds more than K_a codewords, it discards exceeding

codewords at random until the list contains only K_a codewords. Therefore the per-user error probability of the cover decoder is given as $P_e = n_{fa}/(K_a + n_{fa})$. We write $\mathbf{c}_1 \subset \mathbf{c}_2$ if a binary vector \mathbf{c}_1 is covered by a binary vector \mathbf{c}_2 , that is if for all i with $c_{1,i} = 1$ also $c_{2,i} = 1$.

Theorem 8. *Let \mathcal{C} be an outer codebook of size $2^{LJR_{out}}$, where the position of each codeword in each section is chosen uniformly at random. Then the error probability of the cover decoder vanishes in the limit $L, K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$ for some $\alpha > 1$ if*

$$R_{out} < 1 - \frac{1}{\alpha} \quad (2.61)$$

□

Proof. Let \mathcal{L} be a list of K_a arbitrary codewords from \mathcal{C} . Then

$$P(n_{fa} \geq 1) = P\left(\bigcup_{\mathbf{c} \notin \mathcal{L}} \{\mathbf{c} \subset \text{OR}(\mathcal{L})\}\right) \quad (2.62)$$

$$\leq 2^{LJR_{out}} \max_{\mathbf{c} \notin \mathcal{L}} P(\mathbf{c} \subset \text{OR}(\mathcal{L})) \quad (2.63)$$

$$= 2^{LJR_{out}} \max_{\mathbf{c} \notin \mathcal{L}} \prod_{l=1}^L P(\mathbf{c}^l \subset \text{OR}(\mathcal{L}^l)) \quad (2.64)$$

$$\approx 2^{LJR_{out}} \left(\frac{K_a}{2^J}\right)^L \quad (2.65)$$

$$= 2^{LJR_{out} + L(1-\alpha) \log_2 K_a} \quad (2.66)$$

$$= 2^{LJ(R_{out} - (1-\alpha^{-1}))} \quad (2.67)$$

In the second and third line we have used the union bound and the independence of the sections. In the fourth line $\text{OR}(\mathcal{L}^l)$ denotes the OR-combination of the l -th section of the codewords in \mathcal{L} , we have used that the probability that a random number from $[1 : 2^J]$ is contained in a fixed set of size K_a is given by $1 - (1 - 2^{-J})^{K_a}$, which is tightly approximated by $K_a/2^J$ in the considered limit. It is apparent that the error probability vanishes for any L and $J \rightarrow \infty$ if condition (2.61) is fulfilled. □

Remark 1. *The proof of Theorem 8 can easily be extended to include false positives. For that introduce modified lists $\tilde{\text{OR}}(\mathcal{L}^l)$, which in addition to the list of transmitted symbols in section l also contain n_{fa} random erroneous entries. If $n_{fa} = cK_a$ for some constant $c > 0$ the result of Theorem 8 is unchanged. Since $p_{fa} = n_{fa}/2^J$, this condition is equivalent to (2.58).*

One can also derive a finite length upper bound on the achievable outer rates with the cover decoder in a more direct, combinatorial way.

Theorem 9. *Any outer code that can guarantee error-free recovery under cover decoding for K_a users with L -sections in the limit $K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$ has to satisfy:*

$$R_{\text{out}} \leq 1 - \frac{1}{\alpha} + \frac{1}{\alpha L} \quad (2.68)$$

□

Proof. We first show that any error free code has to satisfy

$$\frac{2^{LJR_{\text{out}}}}{K_a} \leq \frac{2^{JL}}{K_a^L} \quad (2.69)$$

To see this, assume an outer code that is error-free, i.e. for any list of K_a codewords the OR-combination of these codewords does not cover any other codeword that is not in the list. Then any two non-intersecting lists, \mathcal{L}_1 and \mathcal{L}_2 , of K_a codewords create two non-intersecting lists of K_a^L possible sensewords. To see that they are non-intersecting, note, that due to the error-free property none of the codewords in \mathcal{L}_2 is covered by the OR-combination of all codewords from \mathcal{L}_1 , which is denoted by $\text{OR}(\mathcal{L}_1)$. This means that each codeword from \mathcal{L}_2 differs from $\text{OR}(\mathcal{L}_1)$ in at least one position. But this also means that any OR-combination of codewords from \mathcal{L}_2 differs from $\text{OR}(\mathcal{L}_1)$ in at least one position. Now divide the set of all $2^{LJR_{\text{out}}}$ codewords into distinct lists of length K_a , then each of these lists corresponds to a distinct list of K_a^L sensewords, whose total number has to be limited by the size of the space:

$$\frac{2^{LJR_{\text{out}}}}{K_a} K_a^L \leq 2^{JL} \quad (2.70)$$

This is precisely the statement of (2.69). Using the scaling condition $2^J = K_a^\alpha$ and taking the limit $K_a, J \rightarrow \infty$ gives the statement of the theorem. □

For $L = 1$ it holds that $R_{\text{out}} \leq 1$, which can obviously be achieved, since for a single section no outer code is necessary.

2.6.1. Tree code

The first practical decoder for the outer OR-MAC with the sectionized structure has been presented in [24]. It works as follows: The B -bit message is divided into blocks of size b_1, b_2, \dots, b_L such that $\sum_l b_l = B$ and such that $b_1 = J$ and $b_l < J$ for all $l = 2, \dots, L$. Each subblock $l = 2, 3, \dots, L$ is augmented to size J by appending $\pi_l = J - b_l$ parity bits,

obtained using pseudo-random linear combinations of the information bits of the previous blocks $l' < l$. Therefore, there is a one-to-one association between the set of all sequences of coded blocks and the paths of a tree of depth L . The pseudo-random parity-check equations generating the parity bits are identical for all users, i.e., each user makes use exactly of the same outer *tree code*. For more details on the outer coding scheme, please refer to [24].

Let \mathcal{S}_l , $l = 1, \dots, L$ be the list of active indices in the l -th section, defined in (2.51). Since the sections contain parity bits with parity profile $\{0, \pi_2, \dots, \pi_L\}$, not all message sequences in $\mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_L$ are possible. The role of the outer decoder is to identify all possible message sequences, i.e., those corresponding to paths in the tree of the outer tree code [24]. The output list \mathcal{L} is initialized as an empty list. Starting from $l = 1$ and proceeding in order, the decoder converts the integer indices \mathcal{S}_l back to their binary representation, separates data and parity bits, computes the parity checks for all the combinations with messages from the list \mathcal{L} and extends only the paths in the tree which fulfill the parity checks. A precise analysis of the error probability in various asymptotic regimes is given in [24]. Specifically, the analysis shows that the error probability of the outer code goes to zero in the familiar limit $K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$ and some $\alpha > 1$ ² if the total number of parity bits $P = \sum_{l=2}^L \pi_l$ is chosen as ([24, Theorem 5 and 6])

1. $P = (L + \delta - 1) \log_2 K_a$ for some constant $\delta > 0$ if all the parity bits are allocated in the last slots.
2. $P = c(L - 1) \log_2 K_a$ for some constant $c > 1$ if the parity bits are allocated evenly at the end of each subslot except for the first.

In the first case the complexity scales like $\mathcal{O}(K_a^{R_{\text{out}}L} \log K_a)$, since there is no pruning in the first $R_{\text{out}}L$ subslots, while in the second case the complexity scales linearly with L like $\mathcal{O}(LK_a \log K_a)$. The corresponding outer rates are

$$\begin{aligned}
 R_{\text{out}} &= B/(B + P) \\
 &= 1 - P/(B + P) \\
 &= 1 - P/(LJ) \\
 &= 1 - \frac{L + \delta - 1}{L\alpha} \\
 &= 1 - \frac{1}{\alpha} + \frac{1}{L} \frac{\delta - 1}{\alpha}
 \end{aligned} \tag{2.71}$$

²We deviate slightly from the notation in [24], where the scaling parameter α' is defined by $B = \alpha' \log_2 K_a$ and the number of subslots is considered to be constant. It is apparent that those definitions are connected by $\alpha' = LR_{\text{out}}\alpha$.

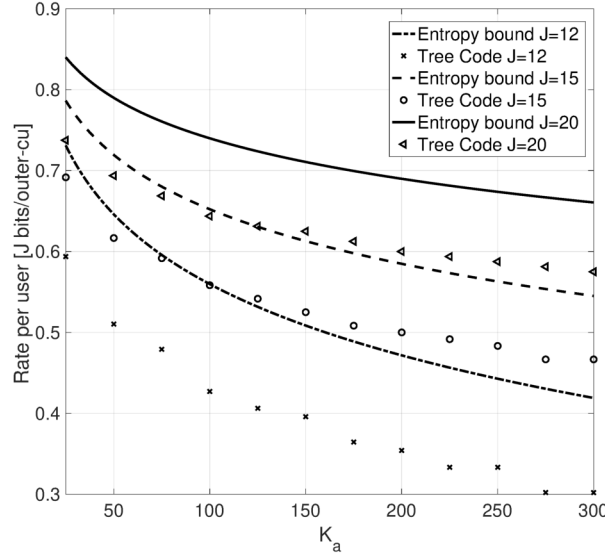


Figure 2.2.: Achievable rates with the tree code with $L = 8$ and $B = 100$ bits together with the upper bound (2.57).

for the case of all parity bits in the last sections and

$$\begin{aligned}
 R_{\text{out}} &= 1 - \frac{c(L-1)}{L\alpha} \\
 &= 1 - \frac{c}{\alpha} - \frac{c}{L\alpha}
 \end{aligned} \tag{2.72}$$

for the case of equally distributed parity bits. In the limit $L \rightarrow \infty$ the achievable rates are therefore $R_{\text{out}} = 1 - 1/\alpha$ and $R_{\text{out}} = 1 - c/\alpha$ respectively, which coincides with the asymptotic upper bound (2.60) for $p_{\text{md}} = 0$ (up to a constant for the second case).

In Figure 2.2 and Figure 2.3 we compare empirical simulations with the developed theory. For the empirical results we fix $B = 100$ bits and $L = 8$. Furthermore, for various values of J and K_a we increase the number of parity bits until a per-user error probability $p_e < 0.05$ is reached. In practice we use a mixture of the two types of parity profile described above. We choose the last section to be only parity bits, while the remaining parity bits are distributed uniformly over the sections $2, \dots, L-1$. The entropy bounds are calculated according to (2.57) with $p_{\text{fa}} = p_{\text{md}} = 0$. In Figure 2.3 we plot the results as a function of $\alpha = J/\log_2 K_a$. The achievable rates of the tree code with the used parity profile turn out to be well described by the formula $R_{\text{out}} = 1 - \alpha^{-1}$ for different J and K_a . Furthermore, the line $R_{\text{out}} = 1 - \alpha^{-1}$ is approached from above by the upper bound (2.57) as $J \rightarrow \infty$.

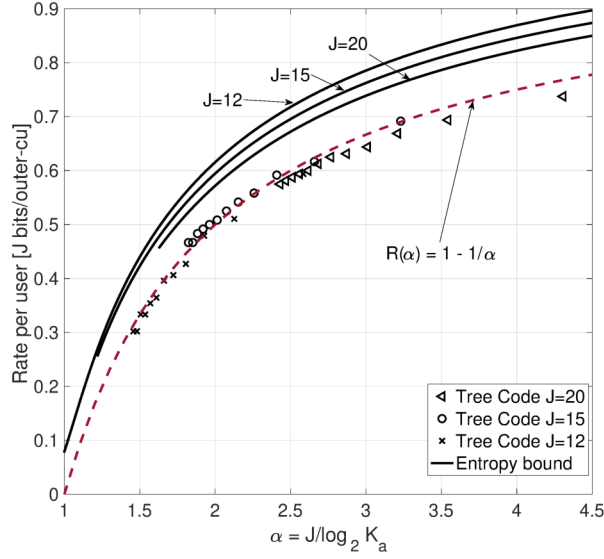


Figure 2.3.: Achievable rates with the tree code with $L = 8$ and $B = 100$ bits together with the upper bound (2.57) as a function of α .

2.7. Analysis of the concatenated scheme

According to Theorem 7, reformulated in terms of the parameters of the concatenated code, the sum-rate is given as

$$S = K_a R = K_a R_{\text{in}} R_{\text{out}} = S_{\text{in}} R_{\text{out}} \quad (2.73)$$

and similarly

$$\frac{E_b}{N_0} = \frac{P}{R N_0} = \frac{\mathcal{E}_{\text{in}}}{R_{\text{out}}}. \quad (2.74)$$

As we have shown in the previous section, the best achievable outer rate is $R_{\text{out}} = 1 - \alpha^{-1}$, which turns out to coincide with the factor appearing in Theorem 7. Since the channel strength in the inner channel is given by $\eta \hat{P}$ and $\hat{P} = n \text{SNR} / L = J \text{SNR} / R_{\text{in}} = 2J \mathcal{E}_{\text{in}}$, the channel strength goes to infinity with J and the error probabilities p_{fa} and p_{md} in the outer channel vanish according to (2.48). An important condition to get the asymptotic limit $R_{\text{out}} = (1 - p_{\text{md}})(1 - \alpha^{-1})$ in (2.60) is (2.58), i.e. that the ratio of false positives to true positives remains at most constant. This condition requires that the channel strength in the inner channel has to grow fast enough to ensure that the probability of false alarms vanishes faster than the sparsity. In the following Theorem we show that the scaling condition for the error probability puts a constraint on the factor η .

Theorem 10. *Let $\eta\hat{P}$ be the channel strength in the scalar Gaussian channel (2.16). In the limit $K_a, J \rightarrow \infty$ with $J = \alpha \log_2 K_a$ for some $\alpha > 1$ the condition (2.58) is fulfilled if and only if*

$$\eta \geq \bar{\eta} \quad (2.75)$$

where $\bar{\eta}$ is given in (2.37). \square

Proof. We show only the direction

$$\eta \geq \bar{\eta} \implies \lim_{K_a \rightarrow \infty} \frac{\log_2 p_{\text{fa}}}{\log_2 K_a} \leq 1 - \alpha \quad (2.76)$$

the reverse implication can be shown similarly.

One can choose the points p_{md} and p_{fa} on the curve defined by (2.48) in a way that $Q^{-1}(p_{\text{md}}) < \epsilon$ for some constant $\epsilon > 0$. Therefore, for $\hat{P} = 2J\mathcal{E}_{\text{in}}$

$$Q^{-1}(p_{\text{fa}}) \geq \sqrt{\eta\hat{P}} - \epsilon = \sqrt{2\eta J\mathcal{E}_{\text{in}}} - \epsilon \quad (2.77)$$

so

$$p_{\text{fa}} \leq Q(\sqrt{2\eta J\mathcal{E}_{\text{in}}} - \epsilon) \quad (2.78)$$

and

$$\lim_{K_a \rightarrow \infty} \frac{\log_2 p_{\text{fa}}}{\log_2 K_a} \leq \lim_{K_a \rightarrow \infty} \frac{1}{\log_2(e)} \left(\frac{-\eta J\mathcal{E}_{\text{in}}}{\log_2 K_a} + \frac{\mathcal{O}(\sqrt{J})}{\log_2 K_a} \right) \quad (2.79)$$

$$= -\frac{\eta\alpha\mathcal{E}_{\text{in}}}{\log_2(e)} \quad (2.80)$$

where the first line follows from the standard bound on the Q-function

$$Q(x) \leq (2\pi)^{-1/2} \exp(-x^2/2)/x. \quad (2.81)$$

By reordering, it holds that (2.80) $< 1 - \alpha$ if

$$\eta > \frac{1 - \alpha^{-1}}{\log_2(e)\mathcal{E}_{\text{in}}} = \bar{\eta} \quad (2.82)$$

\square

The consequences of Theorem 10 for the concatenated code are summarized in the following Corollary.

Corollary 2. *Let $n, L, J, K_a \rightarrow \infty$ and $R, \text{SNR} \rightarrow 0$ with fixed $E_b/N_0 = \text{SNR}/(2R)$, $S = K_a R$ and $J = \alpha \log_2 K_a$ for any $\alpha > 1$. In this limit there is an outer code such that*

the concatenated code described in Section 2.3 with a random Gaussian codebook can be decoded with the SBS-MAP estimator as inner decoder and $P_e \rightarrow 0$ if and only if

$$S < \frac{1}{2} \log_2(1 + K_a \text{SNR}) \quad (2.83)$$

If the AMP algorithm (2.9) is used as inner decoder, reliable decoding is possible if and only if

$$S < \frac{1}{2} (\log_2 e - (E_b/N_0)^{-1}) \quad (2.84)$$

□

Proof. The statement follows immediately from Theorems 7 and 10 together with the relation $\eta_{\text{loc}}^* < \bar{\eta} \leq 1$, which is discussed in the proof of Theorem 7. □

Remark 2. In the case $K_a = 1$ no outer code is necessary, so $R_{\text{in}} = R$ and furthermore $S_{\text{in}} = R$ and $2S_{\text{in}}\mathcal{E}_{\text{in}} = \text{SNR}$. Hence, if $K_a = 1$ is fixed and $J \rightarrow \infty$, which corresponds to $\alpha \rightarrow \infty$, then Corollary 2 recovers the statements of [29, 42], i.e. that SPARCs are reliable at rates up to the Shannon capacity $0.5 \log_2(1 + \text{SNR})$ under optimal decoding. Also the algorithmic threshold (2.84) coincides with the result of [31]. In that sense Theorem 7 and Corollary 2 are an extension of [31] and show that SPARCs can achieve the optimal rate limit in the unsourced random access scenario. However, notice that the concept of the proof technique is simpler, since Theorem 5 is used, which states that not only the sections are described by a decoupled channel model, but in the limit $J \rightarrow \infty$ also the individual components. So the result of Theorem 7 can be derived from the fixpoints of a simple scalar-to-scalar function.

Remark 3. In general, most classical multiple-access variants on the AWGN, where all the users are assumed to have their own codebook, can be represented as sparse recovery problems like (2.4). For that, let $K_a = 1$ and identify the number of section with the number of users. The matrices $\mathbf{A}_1, \dots, \mathbf{A}_L$ are then the codebooks of the individual users and P_l are the transmit power coefficients of different users:

- Fixed L in the limit $J, n \rightarrow \infty$ describes the classical AWGN Adder-MAC from [12], where each user has his own codebook.
- $L, J, n \rightarrow \infty$, where only a fraction of the sections are non-zero describes the many-access channel treated in [8]
- J fixed and $L, n \rightarrow \infty$ describes specific version of the many-access MAC treated in [9, 99]

It is interesting, that in the first case Theorem 7 gives the correct result, after letting $\alpha \rightarrow \infty$, $K_a = 1$ and $L = K$. The case of $J, n \rightarrow \infty$ at finite L is not directly covered though by the presented analysis framework. Nonetheless, the empirical results (e.g. Figure 2.5) show a good agreement with the SE predictions even for small L . Especially the case $L = 1$ is interesting since it resembles the U-RA formulation with random coding, as it was already noted in [99].

In Figure 2.4 and Figure 2.5 we visualize the results of Theorems 2, 3 and 7. For that fix $\alpha = 2$. For various values of J we set K_a such that $J = \alpha \log_2 K_a$. For each value of R_{in} then η_{opt} and η_{alg} are calculated using the approximations of Theorem 4 and Corollary 1. This process is repeated with increasing SNR until $\eta_{\text{opt}} \hat{P}$ and $\eta_{\text{alg}} \hat{P}$ resp. reach a value of $(Q^{-1}(p_{\text{md}}) + Q^{-1}(p_{\text{fa}}))^2$ where the error probabilities are chosen as $p_{\text{md}} = 0.05/L$, with $L = 8$, and $p_{\text{fa}} = 0.01K_a/2^J$. These are the solid lines in Figure 2.4 and Figure 2.5. The dashed lines are the threshold lines from Theorem 7. Additionally, Figure 2.5 shows empirical results, obtained by Monte-Carlo simulations with $L = 8$, where the inner channel with the AMP decoder is simulated with increasing SNR until the error probabilities satisfy $p_{\text{md}} < 0.05/L$ and $p_{\text{md}} < 0.01K_a/2^J$, matching the values above. There are several interesting effects. The asymptotic trade-off curve $S_{\text{in}}(1 - \alpha^{-1}) = 0.5 \log_2(1 + 2S_{\text{in}}\mathcal{E}_{\text{in}})$ is approached from below by the curves for finite J . Also the finite length curves exhibit a region for small S_{in} , where \mathcal{E}_{in} stays almost constant up to some value of S_{in} , and then it starts to grow linearly. Such a behavior was also observed in e.g. [99] in the context of finite-blocklength multiple access on the AWGN channel. This constant regions becomes smaller with increasing J and disappears completely in the asymptotic limit. Also, there is a region of S_{in} in which the algorithmic curve stays almost constant and matches the optimal curve. That is the region where there is only one unique minimum in the RS-potential. The empirical simulations in Figure 2.4 confirm the qualitative behavior of the calculated curves. The required energy stays constant over a large region of S_{in} until some point, where it start to grow rapidly. Note, that Theorem 2 and 3 assume infinite L , but nevertheless the theoretical results match the empirical simulations with $L = 8$ very precisely. According to Theorem 7 in the limit $J \rightarrow \infty$ the required energy grows to infinity as S_{in} approaches $\log_2(e)/[2(1 - 1/\alpha)]$. This value is increased for finite J . In general, one can observe that the asymptotic algorithmic limit is very slowly approached from *above*. That has the interesting consequence, that for each S_{in} there is a value J^* below which the required energy decreases with J and above which the required energy starts to increase again (see Figure 2.6 where this is visualized for $S_{\text{in}} = 2$ with the same parameters as above).

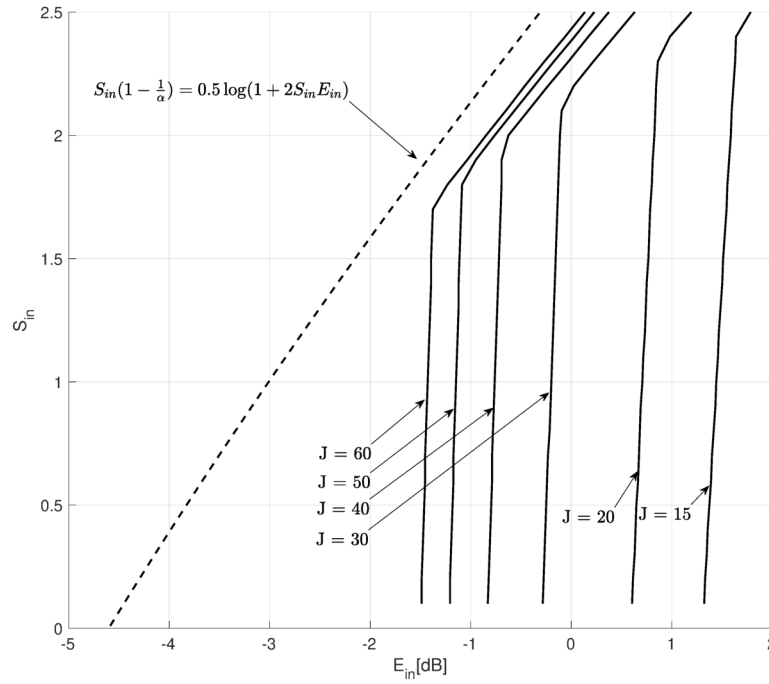


Figure 2.4.: Required \mathcal{E}_{in} to reach specific target error probabilities under optimal decoding and $J = \alpha \log_2 K_a$ for $\alpha = 2$ according to Theorems 2 and 4. The dashed line is the asymptotic limit according to Theorem 7.

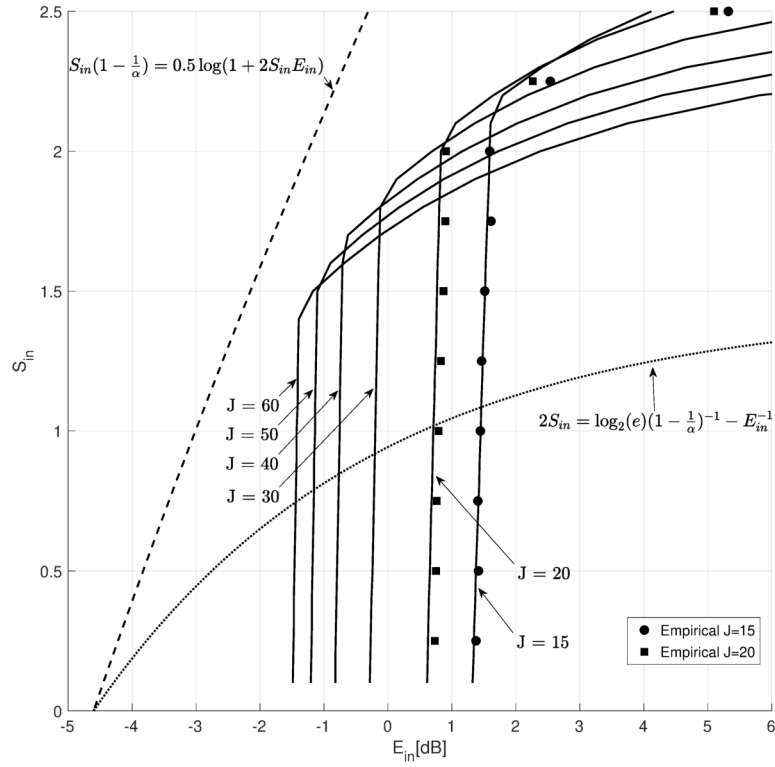


Figure 2.5.: Required \mathcal{E}_{in} to reach specific target error probabilities with AMP decoding and $J = \alpha \log_2 K_a$ for $\alpha = 2$ according to Theorem 3. The empirical simulations were conducted with $L = 8$. The dashed lines are the asymptotic limits according to Theorem 7.

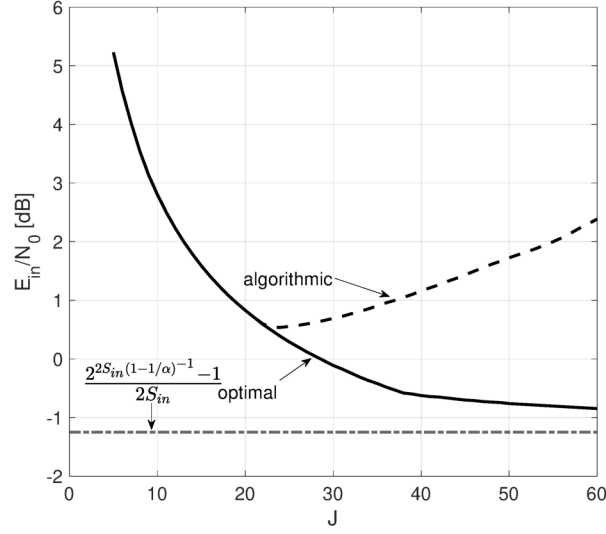


Figure 2.6.: Required \mathcal{E}_{in} to reach specific target error probabilities with $\alpha = 2$ and $S_{\text{in}} = 2$ according to Theorems 2,3 and 7.

2.8. Optimizing the Power Allocation

The foregoing asymptotic analysis has important implications for the code design. It was empirically observed that there is a critical number of users at which the required energy-per-bit under inner AMP decoding increases sharply and that this critical number gets smaller as J grows larger. According to the analytic insight developed in the previous section this behavior is to be attributed to the sub-optimality of the decoder, see Figure 2.5, and it is expected that the required energy will keep decreasing further with J if an optimal decoder is used.

For single user sparse regression codes it is possible to get rid of the local minimum in the potential function through a non-uniform power allocation [30] or through spatial coupling [31, 100]. More generally, it was shown in [101, 102] that when a set of identical copies of a recursive equation with a scalar potential function are coupled in a special way, the potential of the coupled system has only one minimum, which coincides with the global minimum of the scalar potential function, even if the uncoupled potential has a local, non-global minimum. This phenomenon was also termed *threshold saturation* in the context of belief-propagation decoding [103].

We prefer the power allocation approach, because it is easy to obtain an optimized power allocation for an AMP algorithm with separable denoising functions and a given input distribution. In the following, we present a linear programming algorithm to optimize

the power allocation for the AMP algorithm that follows very closely the optimization procedure of [104], which was developed in the context of CDMA.

The denoising function in (2.10) was specific to a uniform power allocation $P_l = P/L$ for all l . For a generic power allocation one can replace the componentwise denoising functions $f_{t,i}$ with ones that depend on the section index:

$$f_{t,i}^l(x) = \frac{\sqrt{\hat{P}_l}}{Z(x)} \sum_{k=0}^{K_a} p_k k \exp \left(\frac{1}{2\tau_t^2} \left(x - k\sqrt{\hat{P}_l} \right)^2 \right) \quad (2.85)$$

where $\hat{P}_l = nP_l/L$ and

$$Z(x) = \sum_{k=0}^{K_a} p_k \exp \left(\frac{1}{2\tau_t^2} \left(x - k\sqrt{\hat{P}_l} \right)^2 \right) \quad (2.86)$$

To analyse the error probability of this modified AMP algorithm in the asymptotic limit $L \rightarrow \infty$, we assume that the powers P_l take values only in a finite set $\{\Pi_1, \dots, \Pi_I\}$ and that the ratio of sections which use Π_i is given by $\alpha_i = |\{l : P_l = \Pi_i\}|/L$ satisfying $\sum_{i=1}^I \alpha_i = 1$ and $\sum_{i=1}^I \alpha_i \Pi_i = P$. Assume that these ratios stay constant as $L \rightarrow \infty$. According to the generalized SE in [36] the recursive equation which describes the behavior of the modified AMP algorithm is given by

$$\tau_{t+1}^2 = \sigma_w^2 + \lim_{L \rightarrow \infty} \frac{\beta}{L2^J} \mathbb{E} [\|\eta_t(\boldsymbol{\theta} + \tau_t \mathbf{Z}) - \boldsymbol{\theta}\|_2^2] \quad (2.87)$$

where $\sigma_w^2 = P^{-1}$,

$$\boldsymbol{\theta} = (\boldsymbol{\theta}^1 | \dots | \boldsymbol{\theta}^L) = \left(\sqrt{\hat{P}_1} \mathbf{s}^1 \mid \dots \mid \sqrt{\hat{P}_L} \mathbf{s}^L \right)^\top \quad (2.88)$$

is a rescaled version of \mathbf{s} and η_t is componentwise given by $\eta_{t,i}^l = \sqrt{\hat{P}_l} f_{t,i}^l$. Since $f_t^l = (f_{t,1}^l, \dots, f_{t,2^J}^l)$ is chosen to be separable, the expected value in (2.87) decouples as follows

$$\mathbb{E} [\|\eta_t(\boldsymbol{\theta} + \tau_t \mathbf{Z}) - \boldsymbol{\theta}\|_2^2] = \sum_{l=1}^L \sum_{j=1}^{2^J} \mathbb{E} \left[\left\| \sqrt{P_l} f_{t,j}^l \left(\sqrt{P_l} s_j^l + \tau_t Z \right) - \sqrt{P_l} s_j^l \right\|_2^2 \right] \quad (2.89)$$

As L goes to infinity the sum over l converges to its mean for each j

$$\lim_{L \rightarrow \infty} \frac{1}{L} \sum_{l=1}^L \mathbb{E} \left[\left\| \sqrt{P_l} f_{t,j}^l \left(\sqrt{P_l} s_j^l + \tau_t Z \right) - \sqrt{P_l} s_j^l \right\|_2^2 \right] = \sum_{i=1}^I \alpha_i \hat{\Pi}_i \mathbb{E} \left[\left\| f_{t,i}^l \left(\sqrt{\Pi_i} s + \tau_t Z \right) - s \right\|_2^2 \right] \quad (2.90)$$

where now s is a random variable distributed according to the marginal empirical distribution of \mathbf{s} . This holds for each component j and each j has the same marginal distribution, so the sum over j becomes redundant.

Remark 4. *In this calculation the sums over j and l are interchangeable. If 2^J is large enough the sum over j already converges to its mean value and the sum over l becomes redundant. This explains heuristically why one can observe a good correspondence between the state evolution and the empirical performance even for small L . Technically the exponential scaling regime with $\beta \rightarrow \infty$ is not covered by the result of [36].*

Note, that the denoising functions $f_{t,i}^l$ were chosen precisely as the PME of s in a scalar Gaussian channel like (2.16) with power $\sqrt{\hat{P}_l}$. So they minimize the MSE in (2.90) and substituting $\tau^2 = \eta^{-1}$ gives the fixpoint condition

$$\eta^{-1} = 1 + \beta \sum_{i=1}^I \alpha_i \hat{\Pi}_i \text{mmse}(\eta \hat{\Pi}_i) \quad (2.91)$$

where $\hat{\Pi}_i = n\Pi_i/L = J\Pi_i/R_{\text{in}}$. The function $\text{mmse}(t)$ is precisely the same as in (2.19). The right hand side of (2.91) is a linear combination of rescaled versions of the original MMSE function in (2.19). The condition that (2.91) has no local minima besides the global minimum around $\eta = 1$ can be formulated as follows:

$$\begin{aligned} & \underset{\boldsymbol{\alpha}}{\text{minimize}} && \sum_{i=1}^I \alpha_i \Pi_i \\ & \text{subject to} && 1 + \beta \sum_{i=1}^I \alpha_i \hat{\Pi}_i \text{mmse}(\eta \hat{\Pi}_i) < \eta^{-1} - \epsilon, \quad \forall \eta \in [0, 1 - \delta] \\ & && \sum_{i=1}^I \alpha_i = 1 \\ & && \alpha_i \geq 0 \end{aligned} \quad (2.92)$$

where $\epsilon, \delta > 0$ are appropriately chosen slack variables. The optimization problem (2.92) is a linear program and therefore easily solvable. The discrete set of Π_i is chosen as follows. For fixed K_a, J, R_{in} set a target inner channel strength. Then Theorem 4 is used to determine the smallest power P_{opt} such that for η_{opt} , the global minimizer of (2.17), it holds that $\eta_{\text{opt}} \hat{P}_{\text{opt}}$ exceeds the target inner channel strength. This P_{opt} serves as lower bound on the set of Π_i . The upper bound is chosen arbitrary, e.g. $5P_{\text{opt}}$. The Π_i are then chosen as a uniform discretisation of the interval $[P_{\text{opt}}, 5P_{\text{opt}}]$. In Figure 2.7 we visualize this for $K_a = 300, J = 20, R_{\text{in}} = 0.0061$, where the slack parameters are $\epsilon = 0.01, \delta = 0.1$.

The solution of (2.92) gives an optimal power distribution that puts weight on only two values, $\Pi_1 = P_{\text{opt}}$ and another value P_* in a ratio of $\alpha_1 \sim 0.81$ to $\alpha_2 \sim 0.19$. The total average power $\alpha_1 P_{\text{opt}} + \alpha_2 P_*$ is about 0.5dB smaller than P_{alg} , the power at which (2.17) has no local minimizers. Therefore, letting one fifth of the sections use a higher power allows to let the other four fifths of the section use the optimal power without having a local convergence point. In Figure 2.7 a) we plot

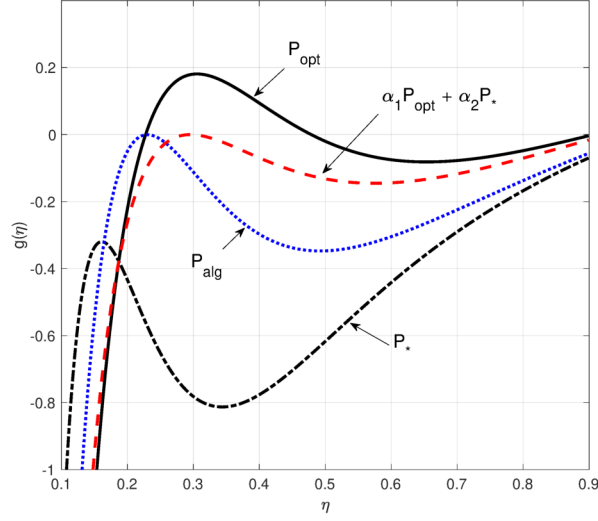
$$g(\eta) = 1 + \beta \left(\alpha_1 \hat{P}_{\text{opt}} \text{mmse}(\eta \hat{P}_{\text{opt}}) + \alpha_2 \hat{P}_* \text{mmse}(\eta \hat{P}_*) \right) - \eta^{-1} \quad (2.93)$$

and its counterparts without power allocation. Figure 2.7 b) shows the integral of $g(\eta)$, which resemble the RS-potential (2.17) with a non-uniform power allocation. The efficiency of the power allocation is demonstrated in Figure 2.8 with finite-length simulations. We choose $L = 8, J = 20$ and use the outer tree code with 0 parity bits in the first section, 20 parity bits in the last, and alternating between 8 and 9 parity bits in the remaining sections. This leaves a total of 89 data bits. We choose those numbers to stay below the typical number of 100 bits that is commonly used in IoT scenarios. The blocklength is chosen as $n = 26229$, which results in $R_{\text{in}} = 0.0061$ and a per-user spectral efficiency of $R_{\text{in}} R_{\text{out}} = 0.0034$. To distribute the power as close as possible to the optimized power allocation obtained above let two sections have a power roughly twice as high as the remaining six sections. It shows that by using this power allocation a gain of about 0.5-1 dB is achievable, which matches the theoretical prediction. With the same parameters as above but $K_a = 200$ it turns out that $P_{\text{alg}} = P_{\text{opt}}$. So the desired inner channel strength of 15dB can be obtained by the AMP algorithm with a flat power allocation. In that case a non-uniform power allocation may be detrimental, because it could introduce unwanted fixpoints into (2.17). Indeed, simulations confirm that the 2-level power allocation that was effective for $K_a = 300$ actually worsens the performance for $K_a \leq 250$. This means that the power allocation has to be tailored carefully to the expected parameters and it only improves the performance if there is a gap between P_{alg} and P_{opt} .

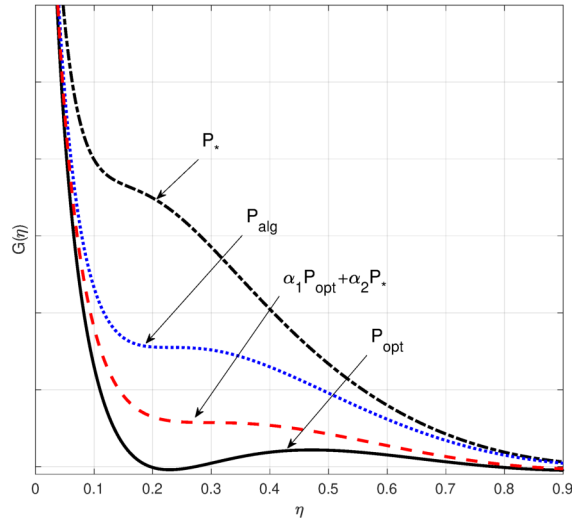
2.9. Considerations for Practical Implementation

The denoising function in the AMP algorithm, (2.10) or its counterpart for non-uniform powers in (2.85), can be simplified if $K_a \ll 2^J$. Then the probabilities p_k are very small for $k \geq 2$ and by neglecting them we get the modified denoising function

$$f_{t,i}^{\text{OR}}(x) = \sqrt{\hat{P}} \left(1 + \frac{p_0}{1 - p_0} \exp \left(\frac{\hat{P} - 2\sqrt{\hat{P}}x}{2\tau_t^2} \right) \right)^{-1} \quad (2.94)$$



(a)



(b)

Figure 2.7.: Visualization of the solution of the optimization problem for $K_a = 300$, $R_{\text{in}} = 0.0061$, $J = 20$. The solution puts weights only on two powers P_{opt} and $P_* \sim 1.9P_{\text{opt}}$ with ratios $\alpha_1 = 0.81$ and $\alpha_2 = 0.19$. The total power in the power allocated system is $\mathcal{E}_{\text{in}} = 1.6\text{dB}$ while the algorithmic threshold is $\mathcal{E}_{\text{in,alg}} = 2.1\text{dB}$. So the power allocation in this case gives a gain of 0.5dB .

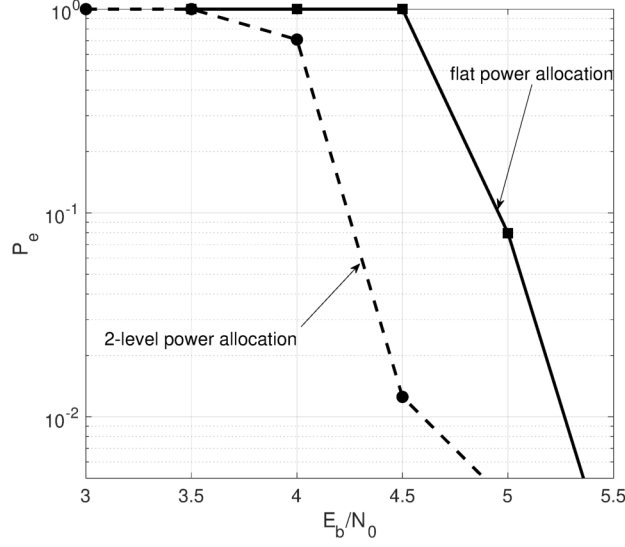


Figure 2.8.: $J = 20, L = 8, B = 89$ bits, $n = 26229, R_{\text{out}} = 0.55$, which leads to $R_{\text{in}} = 0.0061$ and $\mu = 0.0034$. P_e is the per-user-probability of error.

which we denote with the suffix OR because it is the PME in the channel (2.29). In the parameter range of the simulations one can see almost no difference between the full version (2.10) and the truncated OR-estimator (2.94). For moderate values of 2^J compared to K_a , one can improve on (2.94) slightly by taking into account p_2 :

$$f_{t,i}^{\text{OR}+}(x) = \sqrt{\hat{P}} \left(\frac{p_1 \exp\left(-\frac{(x-\sqrt{\hat{P}})^2}{2\tau_t^2}\right) + 2(1-p_0-p_1) \exp\left(-\frac{(x-2\sqrt{\hat{P}})^2}{2\tau_t^2}\right)}{p_0 \exp\left(-\frac{x^2}{2\tau_t^2}\right) + p_1 \exp\left(-\frac{(x-\sqrt{\hat{P}})^2}{2\tau_t^2}\right) + (1-p_0-p_1) \exp\left(-\frac{(x-2\sqrt{\hat{P}})^2}{2\tau_t^2}\right)} \right) \quad (2.95)$$

If necessary, further terms can be included in the same way.

2.10. Finite-length Simulations

In Figure 2.9 the required E_b/N_0 to achieve a per-user probability $P_e < 0.05$ is shown. For the empirical curves with $J = 15$ we use $B = 100$ bits and $n = 30000$ real symbols. For $J = 20$ we use $B = 89$ bits and $n = 26226$ real symbols. This results in a per-user spectral efficiency of $\mu = 0.0033$ and $\mu = 0.0034$, which is the typical value used in comparable works [9, 23–26]. As an inner decoder we use the AMP algorithm (2.9) with (2.95) as denoiser. After the inner decoder converged, in each section, the $K_a + \Delta$ largest entries are declared

as active and added to the list \mathcal{S}_l for the outer tree code. Let $\Delta = 50$. For $J = 15$ we use $L = 16$ and the parity bits for the outer code are chosen as: $\boldsymbol{\pi} = (0, 7, 8, 8, 9, \dots, 9, 13, 14)$. For $J = 20$ we use $L = 8$ and a parity bit distribution $\boldsymbol{\pi} = (0, 9, 8, 9, 8, 9, 8, 20)$. The outer rates are $R_{\text{out}} = 0.4167$ and $R_{\text{out}} = 0.5563$ respectively. For the SE curves we first estimate the required effective inner channel strength by setting the error rates in the outer channel to $p_{\text{md}} = P_e/L$ and $p_{\text{fa}} = \Delta/2^J$. The required inner channel strength is then calculated by (2.48). The potential (2.17) is used to estimate the power to achieve the required inner channel strength. For the curves with power allocation we use the method of Section 2.8 to find the optimal power allocation for $K_a = 300$. In both cases, $J = 15$ and $J = 20$, the required power decreases for $K_a = 300$ but increases for all other values of K_a . So the power allocation has to be adapted to the expected number of users. The empirical values match the theoretically estimated SE curves very well, which confirms the precision of the asymptotic analysis, even though the number of sections is very small.

The obtained value of 4.3 dB for $K_a = 300$ is at the point of writing 0.7 dB better than the best reported value of 5 dB, which was achieved in [105]. For smaller values of K_a other coding schemes have achieved better results, the best of which at the point of writing are [26] and [25], but both of those schemes have shown a rapid increase in required energy as K_a grows large. In [105] an enhanced version of the discussed concatenated coding scheme was presented, where another outer code was used that enabled the passing of soft decoding information between the AMP decoder and the outer decoder, resulting in an turbo-like iterative decoding scheme, alternating between inner and outer decoder. With this type of decoding the required power for $K_a \leq 250$ is reduced significantly, but for $K_a = 300$ the required power is still around 5dB.

2.11. Summary

In this chapter we have introduced a concatenated coding construction that extends the concept of sparse regression codes to the unsourced random access scenario. In this construction an inner code is used as an efficient single user channel code for the AWGN channel and an outer code is used to resolve the multiple access interference. The structural similarity to the coupled compressed sensing scheme allows to use the tree code presented in [24] as an outer code. We use the AMP algorithm as inner decoder, for which we have introduced a low-complexity approximation to the Bayesian optimal denoiser. We introduced a decomposition of the channel into an inner and an outer channel to analyse the asymptotic limits under optimal decoding. Furthermore, we calculated the asymptotically required energy-per-bit of the inner AMP decoder and compared it to the optimal decoder. Finite-length simulations show that the calculated results describe the actual

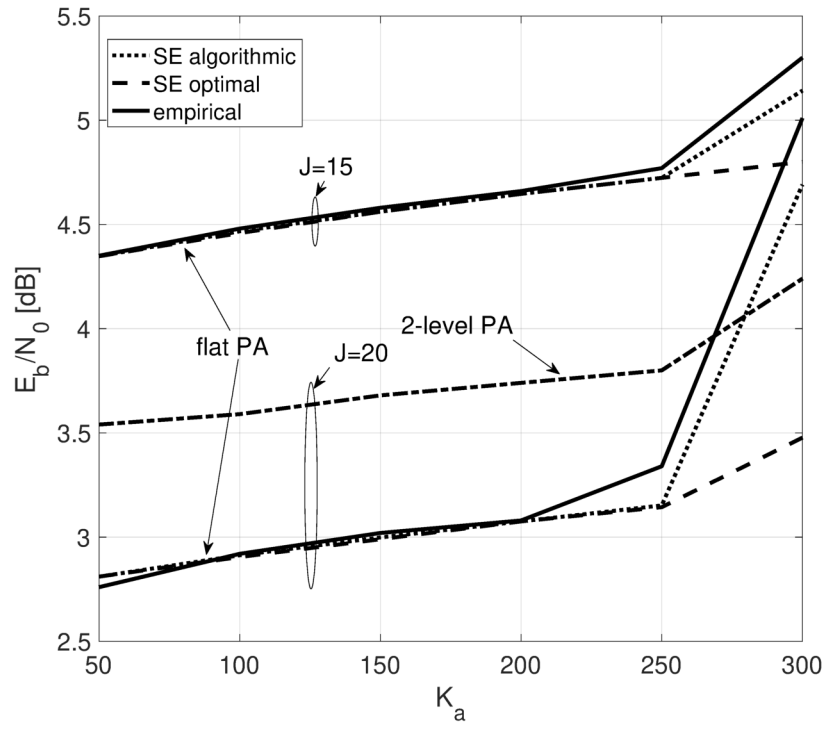


Figure 2.9.: Required E_b/N_0 to achieve a per-user error probability $P_e < 0.05$. The $J = 15$ curve uses $B = 100$, $n = 30000$, and the $J = 20$ curve uses $B = 89$, $n = 26229$.

required energy-per-bit for a fixed per-user error probability very precisely. We find that as $J \rightarrow \infty$, where 2^J is the size of the outer alphabet, the achievable sum-rates of the concatenated code converge to the Shannon limit, even if the number of active users K_a grows to infinity simultaneously, but much faster than J . This is in stark contrast to the typical information theoretic limit, where the size of the message is considered to be much larger than the number of active users. Therefore, it is noteworthy that even under short message length (compared to the number of users) and no coordination between users the Shannon limit can be achieved.

Unfortunately, also the difference in required energy between the AMP decoder and the optimal decoder grows rapidly with J once J surpasses a certain value that depends on the rate. When there is a difference in performance between the AMP and the optimal inner decoder, a non-uniform power allocation can be used to improve the performance of the AMP decoder. We present an linear programming algorithm to find an optimized power allocation. Although for small sum-spectral efficiencies existing U-RA coding schemes like [25,26] are more energy-efficient than the presented scheme, or a sum-spectral efficiency of 1 bit/c.u. and $K_a = 300$ users the presented approach improves on existing ones by almost 1 dB. The good performance at high spectral efficiencies and the availability of a precise analysis make the presented coding scheme stand out among the existing U-RA approaches and therefore an interesting candidate for massive MTC.

The extension of the presented coding scheme and analysis to more general channel models incorporating fading, asynchronicity or multiple receivers seems to be in reach and promising research directions. Furthermore, the presented analysis can be seen as a basis to analyse the more complex turbo-like decoder of [105].

3. Block Fading Multi-User-MIMO Channel

3.1. Outline and Main Contributions

This chapter is organized as follows. Section 3.2 introduces the AD problem on the block-fading MU-MIMO AWGN channel and presents two variants of covariance-based AD algorithms. Furthermore, a novel RIP result is given which allows for a scaling analysis of the presented covariance-based algorithms. The remainder of Section 3.2 describes two alternative AD algorithms and compares them empirically to the covariance-based methods. Section 3.3 introduces the U-RA on the block-fading AWGN channel with a MIMO receiver, presents the CCS algorithm for this channel and establishes the connection to AD. A scaling analysis of the CCS algorithm is given, based on the results from Section 3.2 and finally, the performance of the CCS algorithm is evaluated numerically. In Section 3.4 I present a simple alternative U-RA approach that works well when the coherence block length is large compared to the number of active users. The last Section 3.5 introduces U-RA in a distributed setting and compares different decoding strategies. The main contribution in this part are as follows

- Two efficient iterative algorithms for covariance based AD are introduced and analysed.
- The analysis shows that it is possible to detect the activity of up to

$$K_a = \mathcal{O}(L^2 / \log^2(K_{\text{tot}}/K_a)) \quad (3.1)$$

users out of K_{tot} total users. This is the first time, that such a result has been proven for a concrete algorithm with polynomial complexity. A thorough discussion of related works was given in the introduction.

- In order to prove this scaling result a novel RIP result for Khatri-Rao product matrices is introduced, which is of independent interest, e.g. in angle-of-arrival estimation or machine learning.

- The concept of U-RA is extended to the block fading AWGN channel with multiple receive antennas.
- The CCS algorithm is adapted to work with the covariance-based ML algorithm as inner decoder. This provides a valid alternative for grant-free random access on a block-fading channel with a multi-antenna receiver.
- An analysis of the concatenated scheme shows that sum-spectral efficiencies in the order of $L/\log L$ are achievable, where L is the coherence block-length.
- Finite-length simulations show that the achievable sum-spectral efficiencies are indeed much higher than in current mobile communication standards.
- For large coherence block-lengths another U-RA scheme, based on the pilot transmission and channel estimation, is introduced and analysed.
- It is shown that U-RA in a setting with distributed BSs can circumvent many of the downsides of existing proposed grant-free random access solutions like pilot contamination and the necessity for power control.
- A novel algorithm is introduced that improves the performance of the ML algorithm in the distributed setting by simultaneous AD and position estimation.

3.2. Activity Detection

3.2.1. Signal Model

In this chapter a block-fading wireless channel between each user and the BS is considered where the channel coefficients remain constant over a coherence block consisting of L signal dimensions in the time-frequency domain [45], and change from block to block according to some stationary and ergodic fading process. In this section the problem of AD in a single coherence block is considered. AD refers to the process of identifying the set of K_a active users out of K_{tot} total potential users. For this purpose each user is given a user-specific and a priori known pilot sequence. The problem of AD arises e.g. in several proposed grant-free random access schemes [3–7]. Let the pilot sequence of user k be denoted as $\mathbf{a}_k = (a_{k,1}, \dots, a_{k,L})^\top \in \mathbb{C}^L$. If user k is active, it transmits the components of \mathbf{a}_k in the AD slot of L signal dimensions. Denoting by \mathbf{h}_k the M -dimensional channel vector (small-scale fading coefficients) of user k to the M antennas at the BS, we can write the

received signal at the BS over the AD slot as

$$\mathbf{y}[i] = \sum_{k=1}^{K_{\text{tot}}} b_k \sqrt{P_k g_k} a_{k,i} \mathbf{h}_k + \mathbf{z}[i], \quad i \in [L], \quad (3.2)$$

where $[L] := \{1, \dots, L\}$, $g_k \in \mathbb{R}_+$ denotes the LSFC (channel strength) of user k , $b_k \in \{0, 1\}$ is a binary variable with $b_k = 1$ for active and $b_k = 0$ for inactive users, P_k is the transmit power per signal dimensions of user k and $\mathbf{z}[i] \sim \mathcal{CN}(0, N_0 \mathbf{I}_M)$ denotes the additive white Gaussian noise (AWGN) at the i -th signal dimension. For simplicity we assume $P_k \equiv P$ for all k .

Denoting by $\mathbf{Y} = [\mathbf{y}[1], \dots, \mathbf{y}[L]]^\top$ the $L \times M$ received signal over L signal dimensions and M BS antennas, we can write (3.2) more compactly as

$$\mathbf{Y} = \mathbf{A} \mathbf{\Gamma}^{\frac{1}{2}} \mathbf{H} + \mathbf{Z}, \quad (3.3)$$

where $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_{K_{\text{tot}}}]$ denotes the $L \times K_{\text{tot}}$ matrix of pilot sequences and $\mathbf{\Gamma} = P \mathbf{B} \mathbf{G}$ with \mathbf{G} being a $K_{\text{tot}} \times K_{\text{tot}}$ diagonal matrix consisting of the LSFCs $(g_1, \dots, g_{K_{\text{tot}}})^\top$ and \mathbf{B} being a $K_{\text{tot}} \times K_{\text{tot}}$ diagonal matrix consisting of the binary activity patterns $(b_1, \dots, b_{K_{\text{tot}}})^\top$ of the users. $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_{K_{\text{tot}}}]^\top$ denotes the $K_{\text{tot}} \times M$ matrix containing the M -dimensional normalized channel vectors of the users.

In line with the classical massive MIMO setting [46], we assume for simplicity an independent Rayleigh fading model, such that the channel vectors $\{\mathbf{h}_k : k \in \mathcal{K}_{\text{tot}}\}$ are independent from each other and are spatially white (i.e., uncorrelated along the antennas), that is, $\mathbf{h}_k \sim \mathcal{CN}(0, \mathbf{I}_M)$.

The user pilots are normalized to unit energy per symbol, i.e., $\|\mathbf{a}_k\|_2^2 = L$. Then, the average received SNR of a generic user over L pilot dimensions is given by

$$\text{snr}_k = \frac{\|\mathbf{a}_k\|_2^2 \gamma_k \mathbb{E}[\|\mathbf{h}_k\|_2^2]}{\mathbb{E}[\|\mathbf{Z}\|_F^2]} = \frac{L \gamma_k M}{L M N_0} = \frac{\gamma_k}{N_0}, \quad (3.4)$$

where $\gamma_k = P b_k g_k$ ($b_k = 1$ for active users) is the k -th diagonal element of $\mathbf{\Gamma}$. The vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{K_{\text{tot}}})^\top$ or equivalently the diagonal matrix $\mathbf{\Gamma} = \text{diag}(\boldsymbol{\gamma})$ is called the “active LSFC pattern” of the users. $\mathcal{K}_a \subseteq \{1, \dots, K_{\text{tot}}\}$ shall denote the subset of active users in the current AD slot, with size $|\mathcal{K}_a| = K_a$. Thus, $\boldsymbol{\gamma}$ is a non-negative sparse vector with only K_a nonzero elements. In the presence of LSFCs the definition of active users is slightly relaxed to contain only the users with sufficiently strong received signals $\mathcal{K}_a(\nu) := \{k \in \{1, \dots, K_{\text{tot}}\} : \gamma_k > \nu N_0\}$, for a pre-specified threshold $\nu > 0$, from the noisy observations as in (3.3). As a side goal, we wish also to estimate the LSFCs γ_k of the active users (at least those above the threshold). This information may be useful in practice to

accomplish tasks such as user-BS association, user scheduling, and possibly other high-level network optimization tasks where the knowledge of the user channel strength is relevant, like positioning.

Since the channel vectors are assumed to be spatially white and Gaussian, the columns of \mathbf{Y} in (3.3) are i.i.d. Gaussian vectors with $\mathbf{Y}_{:,i} \sim \mathcal{CN}(0, \mathbf{\Sigma}_{\mathbf{y}})$ where

$$\mathbf{\Sigma}_{\mathbf{y}} = \mathbf{A}\mathbf{\Gamma}\mathbf{A}^H + N_0\mathbf{I}_L = \sum_{k=1}^{K_{\text{tot}}} \gamma_k \mathbf{a}_k \mathbf{a}_k^H + N_0\mathbf{I}_L \quad (3.5)$$

denotes the covariance matrix, which is common among all the columns $\mathbf{Y}_{:,i}$, $i \in [M]$. The empirical covariance matrix of the observations \mathbf{Y} in (3.3) shall be defined as

$$\hat{\mathbf{\Sigma}}_{\mathbf{y}} = \frac{1}{M} \mathbf{Y}\mathbf{Y}^H = \frac{1}{M} \sum_{i=1}^M \mathbf{Y}_{:,i} \mathbf{Y}_{:,i}^H. \quad (3.6)$$

3.2.2. Covariance-Based Algorithms

Maximum Likelihood Estimation

The Gaussianity of the channel vectors allows to express the negative log-likelihood cost function of the conditional distribution $p(\mathbf{Y}|\boldsymbol{\gamma})$ in a particular simple form

$$f(\boldsymbol{\gamma}) := -\frac{1}{M} \log p(\mathbf{Y}|\boldsymbol{\gamma}) \stackrel{(a)}{=} -\frac{1}{M} \sum_{i=1}^M \log p(\mathbf{Y}_{:,i}|\boldsymbol{\gamma}) \quad (3.7)$$

$$\propto \log |\mathbf{A}\mathbf{\Gamma}\mathbf{A}^H + N_0\mathbf{I}_L| + \text{trace} \left(\left(\mathbf{A}\mathbf{\Gamma}\mathbf{A}^H + N_0\mathbf{I}_L \right)^{-1} \hat{\mathbf{\Sigma}}_{\mathbf{y}} \right), \quad (3.8)$$

where (a) follows from the fact that the columns of \mathbf{Y} are i.i.d. (due to the spatially white user channel vectors), and $\hat{\mathbf{\Sigma}}_{\mathbf{y}}$ denotes the sample covariance matrix of the columns of \mathbf{Y} as in (3.6). For spatially white channel vectors considered here, $\hat{\mathbf{\Sigma}}_{\mathbf{y}} \rightarrow \mathbf{\Sigma}_{\mathbf{y}}$ as the number of antennas $M \rightarrow \infty$. It is apparent that the log-likelihood function of $p(\mathbf{Y}|\boldsymbol{\gamma})$ depends on \mathbf{Y} only through the covariance matrix $\hat{\mathbf{\Sigma}}_{\mathbf{y}}$. Therefore, $\hat{\mathbf{\Sigma}}_{\mathbf{y}}$ is a *sufficient statistic* for the estimation of $\boldsymbol{\gamma}$ or any function thereof. Especially in a Massive MIMO scenario, where $M > L$, the use of the covariance matrix $\hat{\mathbf{\Sigma}}_{\mathbf{y}} \in \mathbb{C}^{L \times L}$ instead of the raw measurements $\mathbf{Y} \in \mathbb{C}^{M \times L}$ results in a significant dimensionality reduction. Assuming the number of active users K_a is known, let the *constrained ML estimator* of $\boldsymbol{\gamma}$ be defined as

$$\boldsymbol{\gamma}_{\text{c-ML}}^* = \arg \min_{\boldsymbol{\gamma} \in \Theta_{K_a}^+} f(\boldsymbol{\gamma}). \quad (3.9)$$

where the constraint set $\Theta_{K_a}^+ = \{\boldsymbol{\gamma} \in \mathbb{R}_+^{K_{\text{tot}}} : \|\boldsymbol{\gamma}\|_0 \leq K_a\}$ is the (non-convex) set of non-negative K_a -sparse vectors. There are two problems with this estimator: 1) K_a is generally not known a priori, and 2) the minimization in (3.9) is combinatorial and has exponential complexity in K_{tot} , which can be very large. Therefore, this ML estimator has no practical value. Nevertheless, its performance yields a useful bound to the performance of the following *relaxed ML estimator* of $\boldsymbol{\gamma}$ given by

$$\boldsymbol{\gamma}_{\text{r-ML}}^* = \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}_+^{K_{\text{tot}}}} f(\boldsymbol{\gamma}). \quad (3.10)$$

It is not difficult to check that $f(\boldsymbol{\gamma})$ in (3.8) is the sum of a concave function and a convex function, so also the problem in (3.10) is not convex in general. Besides, the estimator in (3.10) does not require any prior knowledge of K_a .

In the following, for the sake of analysis, we shall denote the true vector of LSFCs as \mathbf{g}° and the true activity pattern as \mathbf{b}° . Next, we consider the performance of the constrained ML estimator (3.9). The idea of the proof is based on [66], which was relying on a RIP result [67] which was then withdrawn because the proof had a flaw. In Appendix F we give a complete and streamlined proof for the case, where the true vector of LSFCs \mathbf{g}° is known at the receiver and all entries satisfy $g_k^\circ \in [g_{\min}, g_{\max}]$. Therefore, the goal consist of estimating the activity pattern \mathbf{b}° and the active LSFC pattern is eventually given by $\boldsymbol{\gamma}_{\text{c-ML}}^* = \mathbf{b}^* \odot \mathbf{g}^\circ$, where \mathbf{b}^* is the estimate of \mathbf{b}° . We hasten to say that our proof technique extends easily also to the case where \mathbf{g}° is unknown, provided that the per-component upper and lower bounds g_{\min} and g_{\max} are known, using the arguments of [66]. We have omitted this general case for the sake of brevity, since it requires a few more technicalities which can be found in [66].

For the case at hand, we define the constrained ML estimator of the activity pattern $\mathbf{b}^\circ \in \{0, 1\}^{K_{\text{tot}}}$ as

$$\mathbf{b}^* := \arg \min_{\mathbf{b} \in \Theta_{K_a}} f(\mathbf{b} \odot \mathbf{g}^\circ), \quad (3.11)$$

with $f(\cdot)$ as defined in (3.7) and $\Theta_{K_a} = \{\mathbf{b} \in \{0, 1\}^{K_{\text{tot}}} : \sum b_k = K_a\}$, the set of binary K_a -sparse vectors. We have the following result:

Theorem 11. *Let the LSFCs be such that for all k it holds that $g_{\min} \leq g_k \leq g_{\max}$. Let $\mathbf{A} \in \mathbb{C}^{L \times K_{\text{tot}}}$, be the pilot matrix with columns drawn uniformly i.i.d. from the sphere of radius \sqrt{L} and let $K_{\text{tot}} > L^2$. For any $\mathbf{b}^\circ \in \Theta_{K_a}$ the estimate \mathbf{b}^* , defined in (3.11), satisfies $\mathbf{b}^* = \mathbf{b}^\circ$ with probability exceeding $1 - 2\epsilon - \exp(-CL)$ (jointly, on a draw of \mathbf{A} and*

3. Block Fading Multi-User-MIMO Channel

a random channel realization), provided that

$$K_a \leq c \frac{L^2}{\log^2(eK_{tot}/L^2)}, \quad (3.12)$$

and

$$M \geq \frac{4}{1-\delta} \left(\frac{C' g_{max} \left(2 \log\left(\frac{eK_{tot}}{2K_a}\right) + \frac{\log(2/\epsilon)}{\max\{K_a, L\}} \right) \max\left\{1, \frac{K_a}{L}\right\} + \frac{N_0}{PL}}{g_{min}} \right)^2 \log\left(3eK_{tot}K_a \frac{1+\epsilon}{\epsilon}\right) \quad (3.13)$$

where $0 < \delta < 1$ and $0 < c, C, C'$ are universal constants that may depend on each other but not on the system parameters. The precise relation is given in the proof. \square

Proof. See Appendix F \square

Theorem 11 gives sufficient conditions under which the error probability of the estimator (3.11) vanishes and it shows that K_a can be larger than L , although then M has to grow at least as fast as $(K_a/L)^2$. Simple algebra (omitted for the sake of brevity) shows the following:

Corollary 3. Let \mathbf{A} be as above and let $M, K_a, L \rightarrow \infty$, then it is possible to choose

$$K_a = \mathcal{O}(L^2 / \log^2(K_{tot}/L^2)) \quad (3.14)$$

and

$$M = \mathcal{O}\left(K_a (g_{max}/g_{min})^2 \log^2(K_{tot}/K_a) \log(K_{tot}K_a)\right) \quad (3.15)$$

such that the estimation error of the ML estimator (3.11) vanishes. \square

Note, that the scaling condition (3.14) can be replaced with the stricter condition

$$K_a = \mathcal{O}(L^2 / \log^2(K_{tot}/K_a)). \quad (3.16)$$

This is because $K_a \leq L^2$ and therefore $L^2 / \log^2(K_{tot}/K_a) \leq L^2 / \log^2(K_{tot}/L^2)$, which implies

$$L^2 / \log^2(K_{tot}/K_a) = \mathcal{O}(L^2 / \log^2(K_{tot}/L^2)). \quad (3.17)$$

As said, the minimization in (3.9) or (3.11) is in general computationally unfeasible (beyond the problem of not knowing K_a). Next, we consider the relaxed ML estimator (3.10), where the domain of search is relaxed to the whole non-negative orthant. This estimator is formally equivalent to the ML estimator of the model parameters in the sparse

Bayesian learning framework, posed in [51]. In [51] a low-complexity iterative algorithm was given and it was shown in [52] that the iterative algorithm is guaranteed to converge to at least a local minimum of (3.7). We derive the iterative update equations here for completeness and show that they can be efficiently implemented by rank-1 updates. While this algorithm is not known to converge to the exact minimum of (3.7), empirical evidence suggests it converges very well. The algorithm proceeds as follows:

For each coordinate $k \in [K_{\text{tot}}]$, define the scalar function $f_k(d) = f(\gamma + d\mathbf{e}_k)$ where $f(\gamma)$ is the likelihood function (3.8) and \mathbf{e}_k denotes the k -th canonical basis vector with a single 1 at its k -th coordinate and zero elsewhere. Setting $\Sigma = \Sigma(\gamma) = \mathbf{A}\mathbf{\Gamma}\mathbf{A}^H + N_0\mathbf{I}_L$ where $\mathbf{\Gamma} = \text{diag}(\gamma)$ and applying the well-known Sherman-Morrison rank-1 update identity [106] we obtain that

$$(\Sigma + d\mathbf{a}_k\mathbf{a}_k^H)^{-1} = \Sigma^{-1} - \frac{d\Sigma^{-1}\mathbf{a}_k\mathbf{a}_k^H\Sigma^{-1}}{1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}. \quad (3.18)$$

Using (3.18) and applying the well-known determinant identity

$$|\Sigma + d\mathbf{a}_k\mathbf{a}_k^H| = (1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k)|\Sigma|, \quad (3.19)$$

we can simplify $f_k(d)$ as follows

$$f_k(d) = c + \log(1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k) - \frac{\mathbf{a}_k^H\Sigma^{-1}\widehat{\Sigma}_{\mathbf{y}}\Sigma^{-1}\mathbf{a}_k}{1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}d \quad (3.20)$$

where $c = \log|\Sigma| + \text{trace}(\Sigma^{-1}\widehat{\Sigma}_{\mathbf{y}})$ is a constant term independent of d . $f_k(d)$ is well-defined only when $d > d_0 := -\frac{1}{\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}$. Taking the derivative of $f_k(d)$ yields

$$f'_k(d) = \frac{\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}{1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k} - \frac{\mathbf{a}_k^H\Sigma^{-1}\widehat{\Sigma}_{\mathbf{y}}\Sigma^{-1}\mathbf{a}_k}{(1 + d\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k)^2}. \quad (3.21)$$

The only solution of $f'_k(d) = 0$ is given by

$$d^* = \frac{\mathbf{a}_k^H\Sigma^{-1}\widehat{\Sigma}_{\mathbf{y}}\Sigma^{-1}\mathbf{a}_k - \mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}{(\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k)^2}. \quad (3.22)$$

Note, that $d^* \geq d_0 = -\frac{1}{\mathbf{a}_k^H\Sigma^{-1}\mathbf{a}_k}$, thus, one can check from (3.20) that f_k is indeed well-defined at $d = d^*$. Moreover, we can check from (3.20) that $\lim_{\epsilon \rightarrow 0^+} f_k(d_0 + \epsilon) = \lim_{d \rightarrow \infty} f_k(d) = \infty$, thus, $d = d^*$ must be the global minimum of $f_k(d)$ in (d_0, ∞) . After the update we have $\gamma_k \leftarrow \gamma_k + d$. So, to preserve the positivity of γ_k , the optimal update step d is in fact given by $\max\{d^*, -\gamma_k\}$. The procedure is summarized in Algorithm 1.

The exact characterization of the performance of this algorithm remains, at the moment, an open problem. Specifically, it is not known under which conditions the iterative algorithm actually reaches the global minimum of (3.7). A heuristic intuition for why the local minima become rare in the large scale limit may be obtained as follows. Let us first note some property of the negative log-likelihood cost function (3.7). Define

$$\mathbf{\Sigma}(\boldsymbol{\gamma}) := \mathbf{A}\boldsymbol{\Gamma}\mathbf{A}^H + N_0\mathbf{I}_L \quad (3.23)$$

and let

$$\phi(\mathbf{\Sigma}) := -\log |\mathbf{\Sigma}| + \text{trace}(\mathbf{\Sigma}\hat{\mathbf{\Sigma}}_y). \quad (3.24)$$

Since $\mathbf{\Sigma}(\boldsymbol{\gamma})$ is positive definite for every non-negative vector $\boldsymbol{\gamma}$ and $N_0 > 0$, it is also invertible and the negative log-likelihood cost function can be expressed as $f(\boldsymbol{\gamma}) = \phi((\mathbf{\Sigma}(\boldsymbol{\gamma}))^{-1})$. Now $\phi : \mathbb{C}^{L \times L} \rightarrow \mathbb{R}$ is *strictly* convex. Hence, it has a unique minimal value over a convex set. Let $\mathbf{\Sigma}_*^{-1}$ denote the unique positive definite matrix with $0 \prec \mathbf{\Sigma}_*^{-1} \preceq 1/N_0$ that minimizes (3.24), and let $\mathbf{\Sigma}_*$ be its inverse. Now, if the set of pilot sequences $\{\mathbf{a}_k : k \in \mathcal{K}_{\text{tot}}\}$ is such that the set $\{\sum_{k=1}^{K_{\text{tot}}} \gamma_k \mathbf{a}_k \mathbf{a}_k^H : \gamma_k \geq 0\}$ spans the whole cone of positive semidefinite matrices, then $\mathbf{\Sigma}_*$ can be represented as $\mathbf{\Sigma}_* = \mathbf{\Sigma}(\boldsymbol{\gamma}^*)$ and therefore $\boldsymbol{\gamma}^*$ is a global minimizer of $f(\boldsymbol{\gamma})$ over $\{\boldsymbol{\gamma} : \gamma_i \geq 0\}$, i.e. $\boldsymbol{\gamma}_{\text{r-ML}}^* = \boldsymbol{\gamma}^*$. Since there are no local minimizers, the componentwise optimization algorithm will necessarily converge to a global minimizer. We cannot apply this argument though, because $\{\sum_{k=1}^{K_{\text{tot}}} \gamma_k \mathbf{a}_k \mathbf{a}_k^H : \gamma_k \geq 0\}$ will never span the whole cone of positive semidefinite matrices for any finite K_{tot} . Nonetheless, if K_{tot} is large enough we expect the approximation of the cone of positive semidefinite matrices to be good enough such that the log-likelihood function has few and small local minima. That explains, at least heuristically, the good convergence behavior of the componentwise optimization algorithm.

Another open problem are the conditions, under which it is guaranteed that the solutions of (3.9) and (3.10) coincide. It is only possible to confirm the validity of $\boldsymbol{\gamma}_{\text{r-ML}}^*$ a-posteriori, i.e. , if $\boldsymbol{\gamma}_{\text{r-ML}}^*$ happens to be K_a -sparse, then it follows that $\boldsymbol{\gamma}_{\text{c-ML}}^* = \boldsymbol{\gamma}_{\text{r-ML}}^*$. Hence, if $\boldsymbol{\gamma}_{\text{r-ML}}^*$ is K_a -sparse *and* the conditions on \mathbf{A}, K_a, M, L and K_{tot} of Theorem 11 are fulfilled, then $\boldsymbol{\gamma}_{\text{r-ML}}^*$ coincides with the correct solution $\boldsymbol{\gamma}^\circ$ with high probability. Intuitive explanations for the sparsity inducing nature of (3.10) have been provided in [51] for the SMV case and in [59] for the MMV case.

Non-Negative Least Squares

In this section we investigate a different approach to estimate $\boldsymbol{\gamma}$ which can be directly analyzed and for which we can provide a rigorous non-asymptotic bound on the ℓ_1 recovery error. Interestingly, analyzing this bound we find that the estimation error vanishes for

$M \rightarrow \infty$ under the same scaling condition (3.12) for K_a, L and K_{tot} as in Theorem 11. The strict convexity of $\phi(\cdot)$ defined in (3.24) suggests the following approach: first, find the matrix $\arg \min_{\Xi \succeq 0} \phi(\Xi)$, where $\{\Xi : \Xi \succeq 0\}$ denotes the set of positive semidefinite matrices. A simple calculation shows that the minimizer is simply the inverse of the empirical covariance matrix $\hat{\Sigma}_{\mathbf{y}}$. Then, find the estimate of γ as

$$\gamma^* = \arg \min_{\gamma \in \mathbb{R}_+^K} \|\Sigma(\gamma) - \hat{\Sigma}_{\mathbf{y}}\|_{\text{F}}^2. \quad (3.25)$$

Let us introduce the matrix $\mathbb{A} \in \mathbb{C}^{L^2 \times K_{\text{tot}}}$, whose k -th column is defined by:

$$\mathbb{A}_{:,k} := \text{vec}(\mathbf{a}_k \mathbf{a}_k^{\text{H}}). \quad (3.26)$$

and let $\mathbf{w} = \text{vec}(\hat{\Sigma}_{\mathbf{y}} - N_0 \mathbf{I}_L)$ denote the $L^2 \times 1$ vector obtained by stacking the columns of $\hat{\Sigma}_{\mathbf{y}} - N_0 \mathbf{I}_L$. Then, we can write (3.25) in the convenient form

$$\gamma^* = \arg \min_{\gamma \in \mathbb{R}_+^K} \|\mathbb{A}\gamma - \mathbf{w}\|_2^2, \quad (3.27)$$

as a linear *least squares* problem with non-negativity constraint, known as *non-negative least squares* (NNLS). Such an algorithm was proposed for the activity detection problem in [50].

A key property of the matrix \mathbb{A} is that a properly centered and rescaled version of it has the restricted isometry property (RIP). Let us define the centered version of \mathbb{A} , denoted by $\mathring{\mathbb{A}}$ as the $L(L-1) \times K_{\text{tot}}$ dimensional matrix, with the k -th column given by

$$\mathring{\mathbb{A}}_{:,k} := \text{vec}_{\text{non-diag}}(\mathbf{a}_k \mathbf{a}_k^{\text{H}} - \text{diag}(\mathbf{a}_k \mathbf{a}_k^{\text{H}})). \quad (3.28)$$

Where $\text{vec}_{\text{non-diag}}(\cdot)$ denotes the vectorization of only the non-diagonal elements, which in the case of $\mathbf{a}_k \mathbf{a}_k^{\text{H}} - \text{diag}(\mathbf{a}_k \mathbf{a}_k^{\text{H}})$, are zero anyway. Let $m = L(L-1)$, then the restricted isometry constant $\delta_{2s} = \delta_{2s}(\mathring{\mathbb{A}}/\sqrt{m})$ of $\mathring{\mathbb{A}}/\sqrt{m}$ of order $2s$ is defined as:

$$\delta_{2s} := \sup_{0 < \|\mathbf{v}\|_0 \leq 2s} \left| \frac{\|\mathring{\mathbb{A}}\mathbf{v}\|_2^2}{m\|\mathbf{v}\|_2^2} - 1 \right| \quad (3.29)$$

and if $\delta_{2s} \in [0, 1)$ the matrix $\mathring{\mathbb{A}}/\sqrt{m}$ is said to have RIP of order $2s$. The normalization is necessary to ensure that the expected norm of the columns of $\mathring{\mathbb{A}}$ is of order $\mathcal{O}(1)$ for all L , which is a necessary condition for the RIP to hold with high probability. It is well known that matrices with iid sub-Gaussian entries satisfy the RIP of order $2s$ with high probability for $s = \mathcal{O}(m/\log(eK_{\text{tot}}/s))$ [65]. The entries of $\mathring{\mathbb{A}}$ though are neither

3. Block Fading Multi-User-MIMO Channel

sub-Gaussian nor independent which makes the analysis more complicated. Nonetheless, recent results in [107, 108] show that matrices which have independent columns (with possibly correlated entries) satisfy the RIP with high probability if the columns have a bounded sub-exponential norm. Using the results of [107] we can establish the following Theorem, which is central for both Theorem 11 and Theorem 13.

Theorem 12. *Let $\mathbf{A} \in \mathbb{C}^{L \times K_{tot}}$, be the pilot matrix with columns drawn uniformly i.i.d. from the sphere of radius \sqrt{L} . Then, with probability exceeding $1 - \exp(-c_\delta \sqrt{m})$ on a draw of \mathbf{A} , it holds that $\mathring{\mathbf{A}}/\sqrt{m}$ has the RIP of order $2s$ with RIP-constant $\delta_{2s}(\mathring{\mathbf{A}}/\sqrt{m}) < \delta$ as long as*

$$2s \leq C_\delta \frac{m}{\log^2(eK_{tot}/m)} \quad (3.30)$$

for some constants $c, c_\delta, C_\delta > 0$ depending only on δ . □

Proof. See Appendix G. □

NNLS has a special property, as discussed for example in [109] and referred to as the \mathcal{M}^+ -criterion in [110], which makes it particularly suitable for recovering sparse vectors: If the row span of \mathbf{A} intersects the positive orthant, NNLS implicitly also performs ℓ_1 -regularization. Because of these features, NNLS has recently gained interest in many applications in signal processing [111], compressed sensing [110], and machine learning. In our case the \mathcal{M}^+ -criterion is fulfilled in an optimally-conditioned manner. Combined with the RIP of $\mathring{\mathbf{A}}$ it allows us to establish the following result:

Theorem 13. *Let $\mathbf{A} \in \mathbb{C}^{L \times K_{tot}}$, be the pilot matrix with columns drawn uniformly i.i.d. from the sphere of radius \sqrt{L} . There exist universal constants $c_i > 0$, $i = 1, \dots, 5$, depending only on some common parameter, but not on the system parameters, (see the proof in Appendix H for details) such that, if*

$$s \leq c_1 \frac{L^2}{\log^2(eK_{tot}/L^2)}, \quad (3.31)$$

then with probability exceeding $1 - \exp(-c_5 L)$ (on a draw of \mathbf{A}) the following holds: For all s -sparse activity pattern vectors $\boldsymbol{\gamma}^\circ$ and all realizations of $\hat{\boldsymbol{\Sigma}}_{\mathbf{y}}$, the solution $\boldsymbol{\gamma}^*$ of (3.27) fulfills for $1 \leq p \leq 2$ the bound:

$$\|\boldsymbol{\gamma}^\circ - \boldsymbol{\gamma}^*\|_p \leq \frac{c_2}{s^{1-\frac{1}{p}}} \sigma_s(\boldsymbol{\gamma}^\circ)_1 + \frac{c_3}{s^{\frac{1}{2}-\frac{1}{p}}} \left(\frac{\sqrt{L}}{\sqrt{s}} + c_4 \right) \frac{\|\mathbf{d}\|_2}{L}, \quad (3.32)$$

where $\sigma_s(\gamma^\circ)_1$ denotes the ℓ_1 -norm of γ° after removing its s largest components and where

$$\mathbf{d} = \text{vec} \left(\hat{\Sigma}_{\mathbf{y}} - \sum_{k=1}^{K_{\text{tot}}} \gamma_k^\circ \mathbf{a}_k \mathbf{a}_k^H - N_0 \mathbf{I}_L \right). \quad (3.33)$$

□

The proof is based on the NNLS results of [110] suitably adapted to our case. The common parameter on which the constants c_i depend is the RIP constant of a properly centered version of \mathbb{A} , defined in (3.26). We state this dependence explicitly to emphasize that Theorem 13 holds also for more general random models for \mathbf{A} , for which \mathbb{A} can be shown to have the RIP. The constants c_2, c_3, c_4 can be computed explicitly (see Appendix H) depending on the RIP constant of the other matrix model. The probability term $1 - \exp(-c_5 L)$ is precisely the probability that the centered version of the random matrix \mathbb{A} has the RIP. The result is uniform meaning that with high probability (on a draw of \mathbf{A}) it holds *for all* γ° and *for all* realizations of the random variable $\hat{\Sigma}_{\mathbf{y}}$. For $s = K_a = \|\gamma^\circ\|_0$ it implies (up to the $\|\mathbf{d}\|_2$ -term) exact recovery since in this case $\sigma_s(\gamma^\circ)_1 = 0$.

The analysis of the random variable $\|\mathbf{d}\|_2$ given in Appendix I shows that, for every realization of \mathbf{A} it holds that

$$\mathbb{E}_{\mathbf{Y}|\mathbf{A}}[\|\mathbf{d}\|_2] = \frac{L}{\sqrt{M}}(\|\gamma^\circ\|_1 + N_0) \quad (3.34)$$

with a deviation tail distribution satisfying

$$\mathbb{P}_{\mathbf{Y}|\mathbf{A}} \left(\|\mathbf{d}\|_2 > \sqrt{\alpha_\epsilon} \mathbb{E}_{\mathbf{Y}|\mathbf{A}}[\|\mathbf{d}\|_2] \right) \leq \epsilon \quad (3.35)$$

for

$$\alpha_\epsilon = c \log((eL)^2/\epsilon) \quad (3.36)$$

with some universal constant $c > 0$. The bounds (3.34) and (3.35) are independent of the realization of \mathbf{A} , so the conditional expectation/probability can be replaced by the total expectation/probability. Assuming that \mathbf{A} is chosen independent of the channel realization, it holds that with probability $(1 - \epsilon)(1 - \exp(-c_5 L)) \geq 1 - \epsilon - \exp(-c_5 L)$ the pilot matrix \mathbf{A} satisfies the condition in Theorem 13 *and* the channel realization \mathbf{d} satisfies $\|\mathbf{d}\|_2 \leq \sqrt{\alpha_\epsilon} \mathbb{E}_{\mathbf{Y}|\mathbf{A}}[\|\mathbf{d}\|_2]$. Setting $s = K_a$ in Theorem 13 (yielding $\sigma_s(\gamma^\circ) = 0$), for $p = 1$ we get the following:

Corollary 4. *With the assumptions as in Theorem 13, the following holds: For any K_a -sparse γ° with*

$$K_a \leq c_1 \frac{L^2}{\log^2(eK_{\text{tot}}/L^2)}, \quad (3.37)$$

3. Block Fading Multi-User-MIMO Channel

the NNLS estimate γ^* fulfills:

$$\frac{\|\gamma^\circ - \gamma^*\|_1}{\|\gamma^\circ\|_1} \leq c_3 \left(\sqrt{L} + c_4 \sqrt{K_a} \right) \frac{1 + \frac{N_0}{\|\gamma^\circ\|_1}}{\sqrt{M/\alpha_\epsilon}} \quad (3.38)$$

with probability at least $1 - \epsilon - \exp(-c_5 L)$, where c_1, c_3, c_4, c_5 are the same constants as in Theorem 13. \square

Using the well-known inequality $\|\gamma^\circ\|_1 \leq \sqrt{K_a} \|\gamma^\circ\|_2$, Theorem 13 for the case $p = 2$ gives:

Corollary 5. *Under the same conditions as in Corollary 4*

$$\frac{\|\gamma^\circ - \gamma^*\|_2}{\|\gamma^\circ\|_2} \leq c_3 \left(\sqrt{L} + c_4 \sqrt{K_a} \right) \frac{\left(1 + \frac{N_0}{\sqrt{K_a} \|\gamma^\circ\|_2} \right)}{\sqrt{M/\alpha_\epsilon}} \quad (3.39)$$

holds with probability at least $1 - \epsilon - \exp(-c_5 L)$ where c_3, c_4, c_5 are the same constants as in Theorem 13 provided that (3.37) holds. \square

In conclusion, the following scaling law is sufficient to achieve a vanishing estimation error.

Corollary 6. *Let $M, K_a, L \rightarrow \infty$ with K_a as in (3.14) and $M = K_a^\kappa$ for $\kappa > 1$ then for $p = 1, 2$ it holds with probability 1 that*

$$\lim_{M \rightarrow \infty} \frac{\|\gamma^\circ - \gamma^*\|_p}{\|\gamma^\circ\|_p} = 0. \quad (3.40)$$

\square

This shows that the NNLS estimator (3.25) can identify up to $O(L^2)$ active users by paying only a poly-logarithmic penalty $O(\log^2(\frac{K_{\text{tot}}}{K_a}))$ for increasing the number of potential users K_{tot} . This is a very appealing property in practical IoT setups where, as already mentioned in the introduction, K_{tot} may be very large. Note, that the scaling of the identifiable users is the same as that of the (uncomputable) restricted ML estimator, see Corollary 3, while the scaling of the minimum required M agrees up to poly-logarithmic factors.

Iterative Solution

Finding the ML estimate γ^* in (3.10) or the NNLS estimate (3.25) requires the optimization of a function over the positive orthant $\mathbb{R}_+^{K_{\text{tot}}}$. In Section 3.2.2 we have derived the

componentwise minimization condition (3.22) of the log-likelihood cost function. Starting from an initial point γ , at each step of the algorithm we minimize $f(\gamma)$ with respect to only one of its arguments γ_k according to (3.22). We refer to the resulting scheme as an *iterative componentwise minimization algorithm*. As discussed before, hopefully this will converge to the solution of (3.10). Variants of the algorithm may differ in the way the initial point is chosen and in the way the components are chosen for update. The noise variance N_0 can also be included as an additional optimization parameter and estimated along γ [51].

The same iterative componentwise minimization approach can be used to solve (iteratively) the NNLS problem (3.25). Of course, the component update step is different in the case of ML and in the case of NNLS. We omit the derivation of the NNLS component update since it consists of a straightforward differentiation operation. Since NNLS is convex, in this case the componentwise minimization algorithm is guaranteed to converge to the solution of the NNLS problem (3.25). Given the analogy of the two iterative componentwise minimization algorithms for ML and for NNLS, we summarize them in a unified manner in Algorithm 1.

ML and NNLS with Knowledge of the LSFCs

Since the ML and NNLS algorithms are non-Bayesian in nature, they work well without any a-priori information on the LSFCs. If \mathbf{g}° (true values of the LSFCs of all users, active and not) is known, the algorithms can be slightly improved by projecting each k -th coordinate update on the interval $[0, g_k^\circ]$ (see step 8) in Algorithm 1. In this case the thresholding step can be improved by choosing the thresholds relative to the channel strength $\hat{\mathcal{A}}_{\mathbf{g}^\circ} = \{i : \hat{\gamma}_i > \theta g_k^\circ\}$.

3.2.3. Empirical Comparison: MF, ML, NNLS and MMV-AMP

Simulation Setting and Performance Criteria

We assume that the output of each algorithm is an estimate γ^* of the active LSFC pattern of the users. We use the relative ℓ_1 norm of the difference $\|\gamma^* - \gamma^\circ\|_1 / \|\gamma^\circ\|_1$ as a measure of the estimation quality. The ℓ_1 norm is the natural choice here, since the coefficients γ_i represent the received signal power, i.e., they are related to the square of the signal amplitudes. Therefore, a more traditional “Square Error” (ℓ_2 norm), related to the 4th power of the signal amplitude, does not really have any relevant physical meaning for the underlying communication system. We define $\hat{\mathcal{A}}_c(\nu) := \{i : \gamma_i^* > \nu N_0\}$, with $\nu > 0$, as the estimate of the set of active users. We also define the misdetection and false-alarm

3. Block Fading Multi-User-MIMO Channel

Algorithm 1 Activity Detection via Coordinate-wise Optimization

- 1: **Input:** The sample covariance matrix $\hat{\Sigma}_{\mathbf{y}} = \frac{1}{M} \mathbf{Y} \mathbf{Y}^H$ of the $L \times M$ matrix of samples \mathbf{Y} .
 - 2: **Input:** The LSFCs of K_{tot} users $(g_1, \dots, g_{K_{\text{tot}}})$ if available.
 - 3: **Initialize:** $\Sigma = N_0 \mathbf{I}_L$, $\gamma = \mathbf{0}$.
 - 4: **for** $i = 1, 2, \dots$ **do**
 - 5: Select an index $k \in [K_{\text{tot}}]$ corresponding to the k -th component of $\gamma = (\gamma_1, \dots, \gamma_{K_{\text{tot}}})^\top$ randomly or according to a specific schedule.
 - 6: **If ML:** Set $d_0^* = \max \left\{ \frac{\mathbf{a}_k^H \Sigma^{-1} \hat{\Sigma}_{\mathbf{y}} \Sigma^{-1} \mathbf{a}_k - \mathbf{a}_k^H \Sigma^{-1} \mathbf{a}_k}{(\mathbf{a}_k^H \Sigma^{-1} \mathbf{a}_k)^2}, -\gamma_k \right\}$.
 - 7: **If NNLS:** Set $d_0^* = \max \left\{ \frac{\mathbf{a}_k^H (\hat{\Sigma}_{\mathbf{y}} - \Sigma) \mathbf{a}_k}{\|\mathbf{a}_k\|_2^4}, -\gamma_k \right\}$.
 - 8: Set $d^* = \min\{d_0^*, g_k - \gamma_k\}$ if LSFC g_k is available and $d^* = d_0^*$ otherwise.
 - 9: Update $\gamma_k \leftarrow \gamma_k + d^*$.
 - 10: Update $\Sigma^{-1} \leftarrow \Sigma^{-1} - \frac{d^* \Sigma^{-1} \mathbf{a}_k \mathbf{a}_k^H \Sigma^{-1}}{1 + d^* \mathbf{a}_k^H \Sigma^{-1} \mathbf{a}_k}$
 - 11: **end for**
 - 12: **Output:** The resulting estimate γ .
-

probabilities as

$$P_{\text{md}}(\nu) = 1 - \frac{\mathbb{E}[|\mathcal{K}_a \cap \hat{\mathcal{A}}_c|]}{K_a}, \quad P_{\text{fa}}(\nu) = \frac{\mathbb{E}[|\hat{\mathcal{A}}_c \setminus \mathcal{K}_a|]}{K_{\text{tot}} - K_a} \quad (3.41)$$

where K_a and K_{tot} denote the number of active and the number of potential users, respectively. By varying $\nu \in \mathbb{R}_+$, we get the *Receiver Operating Characteristic* (ROC) [112] of the algorithms. For simplicity of comparison, in the results presented here we have restricted to the point of the ROC where $P_{\text{md}}(\nu) = P_{\text{fa}}(\nu)$.

We consider several models for the distribution of the LSFCs g_k . The simplest case is when all LSFCs are constant, $g_k \equiv 1$, this corresponds to a scenario with perfect power control. We also consider the case of variable signal strengths such that $10 \log_{10}(g_k)$ is randomly distributed uniformly in some range $[10 \log_{10}(g_{\min}), 10 \log_{10}(g_{\max})]$ (uniform distribution in dB scale). This corresponds to the case of partial power control, where users partially compensate for their physical pathloss and reach some target SNR out of a set of possible values. In practice, these prefixed target SNR values corresponds to the various Modulation and Coding Schemes (MCS) of a given communication protocol, which in turn correspond to different data transmission rates (see for example the MCS modes of standards such as IEEE 802.11 [113] or 3GPP-LTE [114]). In passing, we notice here the importance of estimating not only the user activity pattern but their LSFCs, in order to perform rate allocation.

Matched Filter

A simple sub-optimal approach to recover the LSFCs of the active users is to use the pilot sequences as a bank of matched filters (MF) and apply them one by one to \mathbf{Y} . An estimate of the LSFC γ_k is then obtained as the average of the squares:

$$\hat{\gamma}_k^{\text{MF}} = \frac{\|\mathbf{a}_k^H \mathbf{Y}\|_2^2}{M} \quad (3.42)$$

This approach was recently analysed in [115] where a scaling law on the sample complexity was proven for sub-Gaussian \mathbf{A} . The results show as similar scaling behavior as in Theorem 11, specifically recovery is possible if $M = \mathcal{O}(K_a^2/L^2)$. It was shown in [116] that the estimator (3.42) is heavily biased, but the relative order of the LSFCs is preserved on average, i.e. $\mathbb{E}[\hat{\gamma}_1] \leq \dots \leq \mathbb{E}[\hat{\gamma}_{K_{\text{tot}}}]$ whenever $\gamma_1 \leq \dots \leq \gamma_{K_{\text{tot}}}$, where the expectation is over the M measurements. Therefore we can detect the active users by picking the K_a indices with the largest values $\hat{\gamma}_k$.

MMV-AMP

This version of AMP, as introduced in [117], is a Bayesian iterative recovery algorithm for the MMV problem, i.e., it aims to recover an unknown matrix with i.i.d. rows from linear Gaussian measurements. As said in the introduction, the use of MMV-AMP has been proposed in [4, 118] for the AD problem in a Bayesian setting, where the LSFCs are either known, or its distribution is known. Since unfortunately the formulation of MMV-AMP is often lacking details and certain terms (e.g., derivatives of matrix-valued functions with matrix arguments) are left indicated without explanations, for the sake of clarity and in order to provide a self-contained exposition we briefly review this algorithm here in the notation of this paper.

We can rewrite the received signal in (3.2) as

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{Z} \quad (3.43)$$

with $\mathbf{X} = \mathbf{G}\mathbf{B}\mathbf{H}$. We assume here that $P = 1$ for simplicity. Let $\mathbf{X}_{k,:}$ denote the k -th row of \mathbf{X} . Letting $\lambda = \frac{K_a}{K_{\text{tot}}}$ be the fraction of active users, in the Bayesian setting underlying the MMV-AMP algorithm it is assumed that the rows of \mathbf{X} are mutually statistically independent and identically distributed according to

$$p_X(\mathbf{x}) = (1 - \lambda)\delta_0 + \lambda \int_0^{+\infty} \frac{e^{-\frac{\|\mathbf{x}\|_2^2}{\zeta}}}{\pi\zeta} dp_G(\zeta), \quad (3.44)$$

where $p_G(\cdot)$ is the distribution of the LSFCs, i.e., for each k , it is assumed that $\mathbf{X}_{k,:}$ is either the identically zero vector (with probability λ) or a conditionally complex i.i.d. M -dimensional Gaussian vector with mean 0 and conditional variance g_k . Furthermore, the g_k 's are i.i.d. $\sim p_G(\cdot)$. The conditional distribution of $\mathbf{X}_{k,:}$ given g_k is obviously given by

$$p_{X|g}(\mathbf{x}|g_k) = (1 - \lambda)\delta_0 + \lambda \frac{e^{-\frac{\|\mathbf{x}\|_2^2}{g_k}}}{\pi g_k}. \quad (3.45)$$

The MMV-AMP iteration is defined as follows:

$$\mathbf{X}^{t+1} = \eta_t(\mathbf{A}^H \mathbf{Z}^t + \mathbf{X}^t) \quad (3.46)$$

$$\mathbf{Z}^{t+1} = \mathbf{Y} - \mathbf{A} \mathbf{X}^{t+1} + \frac{K_{\text{tot}}}{L} \mathbf{Z}^t \langle \eta'_t(\mathbf{A}^H \mathbf{Z}^t + \mathbf{X}^t) \rangle \quad (3.47)$$

with $\mathbf{X}^0 = \mathbf{0}$ and $\mathbf{Z}^0 = \mathbf{Y}$. The function $\eta_t : \mathbb{C}^{K_{\text{tot}} \times M} \rightarrow \mathbb{C}^{K_{\text{tot}} \times M}$ is defined row-wise as

$$\eta_t(\mathbf{R}) = \begin{bmatrix} \eta_{t,1}(\mathbf{R}_{1,:}) \\ \vdots \\ \eta_{t,K_{\text{tot}}}(\mathbf{R}_{K_{\text{tot}},:}) \end{bmatrix}, \quad (3.48)$$

where each row function $\eta_{t,k} : \mathbb{C}^M \rightarrow \mathbb{C}^M$ is chosen as the posterior mean estimate of the random vector \mathbf{x} , with a priori distribution as the rows of \mathbf{X} as given above, in the *decoupled* Gaussian observation model

$$\mathbf{r} = \mathbf{x} + \mathbf{z}, \quad (3.49)$$

where \mathbf{z} is an i.i.d. complex Gaussian vector with components $\sim \mathcal{CN}(0, \Sigma_t)$. When \mathbf{g} is known, such posterior mean estimate is conditional on the knowledge of g_k , i.e., we define

$$\eta_{t,k}(\mathbf{r}) = \tilde{\eta}_t(\mathbf{r}, g_k) := \mathbb{E}[\mathbf{x}|\mathbf{r}, g_k]. \quad (3.50)$$

If \mathbf{g} is not known, the posterior mean estimate is unconditional, i.e., we define (with some abuse of notation)

$$\eta_{t,k}(\mathbf{r}) = \tilde{\eta}_t(\mathbf{r}) := \mathbb{E}[\mathbf{x}|\mathbf{r}]. \quad (3.51)$$

Notice that in the latter case $\eta_{t,k}(\cdot)$ does not depend on k , i.e., the same mapping $\tilde{\eta}_t(\cdot)$ is applied to all the rows in (3.48). The noise variance in the decoupled observation model, Σ_t is provided at each iteration t by the following recursive equation termed *State Evolution* (SE),

$$\Sigma_{t+1} = N_0 \mathbf{I}_M + \frac{K_{\text{tot}}}{L} \mathbb{E}[\mathbf{e}_t \mathbf{e}_t^H] \quad (3.52)$$

where

$$\mathbf{e}_t = \begin{cases} (\tilde{\eta}_t(\mathbf{x} + \mathbf{z}, g_k) - \mathbf{x})^\top & \text{if } \mathbf{g} \text{ is known} \\ (\tilde{\eta}_t(\mathbf{x} + \mathbf{z}) - \mathbf{x})^\top & \text{if } \mathbf{g} \text{ is not known} \end{cases} \quad (3.53)$$

The initial value of the SE is given by $\mathbf{\Sigma}_0 = N_0 \mathbf{I}_M + \frac{K_{\text{tot}}}{L} \mathbb{E}[\mathbf{x}\mathbf{x}^H]$. The sequence $(\mathbf{\Sigma}_t)_{t=0,1,2,\dots}$ does not depend on a specific input \mathbf{X} and can be precomputed. The SE equation has the important property that it predicts the estimation error of the AMP output $\{\mathbf{X}^t\}_{t=0,1,\dots}$ asymptotically in the sense that in the limit of $K_{\text{tot}}, L \rightarrow \infty$ with $L/K_{\text{tot}} = \text{const.}$ it holds that [119]

$$\lim_{K_{\text{tot}} \rightarrow \infty} \frac{\|\mathbf{X}^{t+1} - \mathbf{X}\|_F^2}{K_{\text{tot}}} = \text{tr}(\mathbb{E}[\mathbf{e}_t \mathbf{e}_t^H]) = \text{tr}(\mathbf{\Sigma}_t - N_0 \mathbf{I}_M) \frac{L}{K_{\text{tot}}}. \quad (3.54)$$

Formally this was proven for the case when the entries of \mathbf{A} are Gaussian iid. In practice this property holds also when the columns of \mathbf{A} are sampled uniformly from the sphere, as in our case. Note, that $\text{tr}(\mathbb{E}[\mathbf{e}_t \mathbf{e}_t^H])$ is the MSE of the estimator $\tilde{\eta}$ in the Gaussian vector channel (3.49) and therefore the choice (3.50) (or (3.51) resp.) is asymptotically optimal as it minimizes the MSE in each iteration.

Since there is no spatial correlation between the receive antennas, $\mathbf{\Sigma}_0$ is diagonal and it can be shown (see [118]) that $\mathbf{\Sigma}_t$ is diagonal for all t . In the case of \mathbf{g} is known to the AD estimator, a simple calculation yields the function $\tilde{\eta}_{t,k}(\mathbf{r})$ defined in (3.50) in the form

$$\tilde{\eta}_{t,k}(\mathbf{r}) = \phi_{t,k}(\mathbf{r}) g_k (g_k \mathbf{I}_M + \mathbf{\Sigma}_t)^{-1} \mathbf{r}, \quad (3.55)$$

where the coefficient $\phi_{t,k}(\mathbf{r}) \in [0, 1]$ is the posterior mean estimate of the k -th component b_k of the activity pattern \mathbf{b} , when rewriting the decoupled observation model (3.49) as $\mathbf{r} = \sqrt{g_k} b_k \mathbf{h} + \mathbf{z}$. In particular, we have (details are omitted and can be found in [118])

$$\begin{aligned} \phi_{t,k}(\mathbf{r}) &= \mathbb{E}[b_k | \mathbf{r}, g_k] \\ &= p(b_k = 1 | \mathbf{r}, g_k) \\ &= \left\{ 1 + \frac{1 - \lambda}{\lambda} \prod_{i=1}^M \left[\frac{g_k + \tau_{t,i}^2}{\tau_{t,i}^2} \exp \left(-\frac{g_k |r_i|^2}{\tau_{t,i}^2 (g_k + \tau_{t,i}^2)} \right) \right] \right\}^{-1} \end{aligned} \quad (3.56)$$

The term $\langle \eta'(\cdot) \rangle$ in (3.47) is defined as

$$\langle \eta'_t(\mathbf{R}) \rangle = \frac{1}{K_{\text{tot}}} \sum_{k=1}^{K_{\text{tot}}} \eta'_{t,k}(\mathbf{R}_{k,:}), \quad (3.57)$$

where $\eta'_{t,k}(\cdot) \in \mathbb{C}^{M \times M}$ is the Jacobi matrix of the function $\eta_{t,k}(\cdot)$ evaluated at the k -th row $\mathbf{R}_{k,:}$ of the matrix argument \mathbf{R} . For known LSFCs and uncorrelated antennas (yielding

diagonal $\mathbf{\Sigma}_t = \text{diag}(\tau_{t,1}^2, \dots, \tau_{t,M}^2)$ for all t), the derivative is explicitly given by

$$\eta'_{t,k}(\mathbf{r}) = \phi_{t,k}(\mathbf{r})\text{diag}(\mathbf{\Xi}_{t,k}) + (\mathbf{\Xi}_{t,k}\mathbf{r})(\tilde{\mathbf{\Xi}}_{t,k}\mathbf{r})^H(\phi_{t,k}(\mathbf{r}) - \phi_{t,k}(\mathbf{r})^2) \quad (3.58)$$

where we define $\mathbf{\Xi}_{t,k} = \text{diag}\left(\frac{g_k}{g_k + \tau_{t,i}^2} : i \in [M]\right)$ and $\tilde{\mathbf{\Xi}}_{t,k} = \text{diag}\left(\frac{g_k}{\tau_{t,i}^2(g_k + \tau_{t,i}^2)} : i \in [M]\right)$. Analogous expressions for the case where the LSFCs \mathbf{g} are unknown to the receiver can be found, but their expression cannot be generally given in a compact form and in general depends on the LSFC distribution $p_G(\cdot)$ (see [118] for more details).

MMV-AMP Scaling

For the single measurement vector (SMV) case ($M = 1$) it was shown in [34] that in the asymptotic limit $L, K_{\text{tot}}, K_a \rightarrow \infty$ with fixed ratios L/K_{tot} and K_a/K_{tot} the estimate $\mathbf{A}^H \mathbf{z}^t + \mathbf{x}^t$ in the AMP algorithm in the t -th iteration is indeed distributed like the true target signal in Gaussian noise with noise variance Σ^t given by the SE. A generalized version of this statement that includes the MMV case was proven in [119]. It was shown in [4] that, based on the state evolution equation (3.52), the error of activity detection vanishes in the limit $M \rightarrow \infty$ for any number of active users. It is important to notice that, in this type of SE-based analysis, first the limit $K_a, L \rightarrow \infty$ is taken at fixed M and then the limit $M \rightarrow \infty$ is taken. This makes it impossible to derive a scaling relation between M and K_a . Furthermore, this order of taking limits assumes that K_a is much larger than M . Hence, this type of analysis does not generally describe the case when M scales proportional to K_a or even a bit faster. Finally, it is implicit in this type of analysis that L , K_a and K_{tot} are asymptotically in linear relation, i.e., $\frac{K_a}{L} \rightarrow \alpha$ and $\frac{K_{\text{tot}}}{L} \rightarrow \beta$ for some $\alpha, \beta \in (0, \infty)$. Hence again, it is impossible to capture the scaling studied in our work, where K_a is essentially quadratic in L , K_{tot} can be much larger than K_a , and M scales to infinity slightly faster than K_a .

The above observation is a possible explanation for the behavior described in Section 3.2.3, which is in fact quite different from what is predicted by the SE and in fact reveals an annoying non-convergent behavior of MMV-AMP when M is large with respect to L and the dimensions are of “practical interest”, i.e., not extremely large.

MMV-AMP Approximations

Instead of pre-computing the sequence $(\mathbf{\Sigma}_t)_{t=0,1,\dots}$, in the SMV case, where $\mathbf{\Sigma}_t$ reduces to a single parameter τ_t^2 , it is common to use the norm of the residual $\|\mathbf{Z}_t\|_2^2/K_{\text{tot}}$ as an empirical estimate of $\mathbf{\Sigma}_t$ [31,33], since it leads to faster convergence [120] while disposing the need of pre-computing the state evolution recursion. We find empirically that, analogous

to the SMV case, estimating the i -th diagonal entry of $\mathbf{\Sigma}_t = \text{diag}(\tau_{t,1}^2, \dots, \tau_{t,M}^2)$ as $\tau_{t,i}^2 = \|\mathbf{Z}_{:,i}^t\|_2^2 / K_{\text{tot}}$ (i.e., the empirical variance of the i -th column of the matrix \mathbf{Z}^t in (3.47)) leads to a good performance.

Another possible approximation arises from the observation that in the derivative (3.58), the diagonal terms are typically much larger than the off-diagonal terms, which is to be expected, since in expectation the off-diagonal entries of the term $(\Xi_{t,k}\mathbf{r})(\tilde{\Xi}_{t,k}\mathbf{r})^H$ vanish. So we find empirically that reducing the calculation of the derivative to just the diagonal entries, barely alters the performance in a large parameter regime, while significantly reducing the complexity of the MMV-AMP iterations from $\mathcal{O}(M^2)$ to $\mathcal{O}(M)$.

Activity detection with MMV-AMP

For known LSFCs an estimate of the activity pattern can be obtained directly by thresholding the posterior mean estimate of b_k (3.56). For statistically known LSFCs we have to calculate the integral of (3.56) over the distribution of the LSFCs. For large M this integral may become numerically unstable, in that case we can also use the following method: Let \mathbf{X}^{t_0} and \mathbf{Z}^{t_0} denote the output of the MMV-AMP algorithm at the final iteration. Let $\mathbf{R}^{t_0} := \mathbf{A}^H \mathbf{Z}^{t_0} + \mathbf{X}^{t_0}$. Under the assumption that the asymptotic decoupling phenomenon described in Section 3.2.3 holds, i.e. that the decoupled observation model represents faithfully the statistics of the rows of \mathbf{R}^{t_0} , each row $\mathbf{R}_{k,:}^{t_0}$ is distributed as $\sqrt{\gamma_k} \mathbf{h}_k + \mathbf{z}_k$ with $\mathbf{z}_k \sim \mathcal{CN}(0, \mathbf{\Sigma}^{t_0})$ and \mathbf{h}_k has the statistics of the Gaussian MIMO i.i.d. channel vector of user k . Furthermore we assume that $\mathbf{\Sigma}^{t_0}$ is diagonal, with entries $\tau_{t_0,i}^2 : i = 1, \dots, M$, which are estimated as described in the previous section. Then the ML estimate of γ_k from \mathbf{R}^{t_0} is given by

$$\hat{\gamma}_k = \max \left(0, \frac{\|\mathbf{R}_{k,:}^{t_0}\|_2^2}{M} - \frac{\sum_{i=1}^M \tau_{t_0,i}^2}{M} \right). \quad (3.59)$$

Then, the activity pattern as well as the active LSFC pattern can be obtained by thresholding the $\hat{\gamma}_k$.

Instability of MMV-AMP

In simulations, we have observed that the MMV-AMP algorithm as described in Section 3.2.3, for certain parameter settings, exhibits an annoying non-convergent behavior that occurs at random with some non-negligible probability (according to the realization of the random pilot matrix \mathbf{A} , the random channel matrix \mathbf{H} , and the random observation noise). We find that this behavior occurs most frequently for either small $K_a \ll L$ and M similar to or larger than K_a , or for $M > K_a > L$. Also the dynamic range of the LSFCs plays an important role. While this behavior occurs less frequently or completely vanishes for

a small dynamic range or constant LSFCs, it occurs more frequently for large dynamic ranges. For example if we let g_k be distributed uniformly in dB scale between 0 and 20dB, known at the receiver, for $K_a = 20$ the algorithm is stable for $M = 4$, in the sense that the effective noise variance τ_t^2 decreases consistently, but unstable for $M = 10$, i.e. for many instances the actual measured values of $\|\mathbf{X}^t - \mathbf{X}\|_F^2 / (MK_{\text{tot}})$ diverge a lot from their SE prediction (3.52). This behavior is illustrated in Figure 3.1, where $\|\mathbf{X}^t - \mathbf{X}\|_F^2 / (MK_{\text{tot}})$ is plotted for $t = 1, 2, \dots$ for several samples along with τ_t^2/M , where $\mathbf{\Sigma}_t = \tau_t^2 \mathbf{I}_M$ is calculated according to the SE (3.52). For $K_a < L$ one may argue that this is an artificial behavior, which can be circumvented by simply discarding the information from some of the antennas, but this is certainly not possible for $K_a > L$, where $M > K_a$ measurements are necessary. We find that specifically in this regime $M > K_a > L$ the MMV-AMP performance differs significantly from its state evolution prediction, which is consistent with what was argued before. These outliers occur even if none of the approximations mentioned in Section 3.2.3 are applied. Although we find that approximating the derivative $\eta'(\cdot)$ as described in Section 3.2.3 helps to reduce the number of samples that do not converge to the state evolution prediction. Another observation is that the use of normalized pilots ($\|\mathbf{a}_k\|_2^2 = L$) improves the convergence to the SE prediction compared to Gaussian iid pilots.

Complexity Comparison

The complexity of the discussed covariance-based AD algorithms (ML and NNLS) scales with the size of the covariance matrix and the total number of users, i.e. $\mathcal{O}(K_{\text{tot}}L^2)$, plus the complexity of once calculating the empirical covariance matrix which is linear in ML .

The complexity of the MF approach is $\mathcal{O}(MLK_{\text{tot}})$ or, with a sub-sampled FFT matrix as pilot matrix, $\mathcal{O}(MK_{\text{tot}} \log K_{\text{tot}})$.

The complexity of MMV-AMP in each iteration scales like $\mathcal{O}(M^2LK_{\text{tot}})$ or, with a sub-sampled FFT matrix as pilot matrix, like $\mathcal{O}(M^2K_{\text{tot}} \log K_{\text{tot}})$. Using the simplified derivative as described in paragraph 3.2.3 the complexity is reduced to $\min(\mathcal{O}(MK_{\text{tot}} \log K_{\text{tot}}), \mathcal{O}(MK_{\text{tot}}L))$. In any case the covariance-based algorithms scale better with M , while MF and MMV-AMP scales better with L .

Scaling

The performance of AD is visualized in Figure 3.2 ('CS regime', i.e. $K_a \leq L$) and Figure 3.3 ($K_a > L$). Here we assumed all the LSFCs to be identically equal to 1, MMV-AMP was run with the full knowledge of the LSFCs and the ML and NNLS algorithms were run with the box-constraints described in Section 3.2.2. In Figure 3.2 the NNLS algorithm is comparably worse than MMV-AMP and ML. This is to be expected, since M is small com-

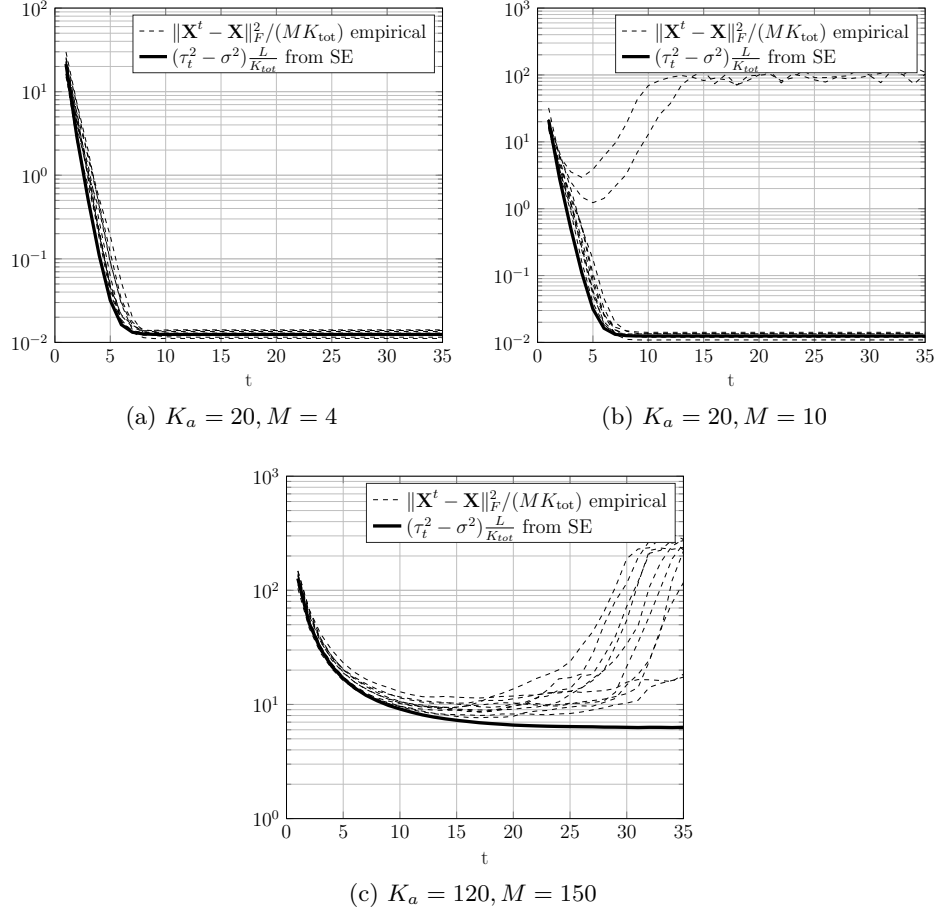


Figure 3.1.: Evolution of the normalized MSE in the AMP iterations (3.46)-(3.47) for 10 sample runs and its state evolution prediction from (3.52). $L = 100$, $K_{\text{tot}} = 2000$ and the LSFCs are chosen such that snr_k (see (3.4)) are uniformly distributed between 0 and 20dB and are assumed to be known at the receiver.

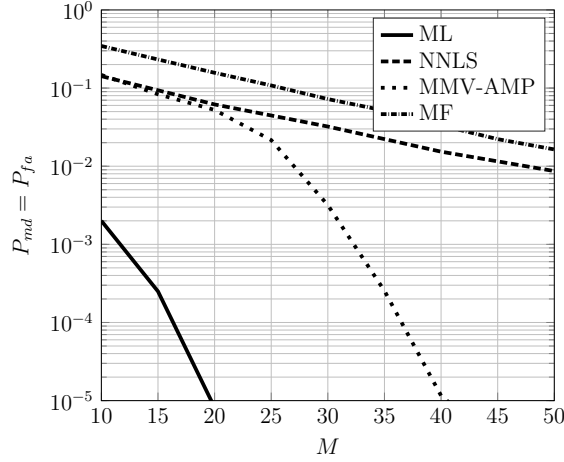


Figure 3.2.: Scaling of the support detection error vs. M at the border of the CS regime for $K_a = L = 100$, $K_{\text{tot}} = 2000$ with constant LSFCs at $\text{SNR}_k = 0$ dB.

pared to L , which leads to a significant gap between the true and the empirical covariance matrix $\|\hat{\Sigma}_{\mathbf{y}} - \Sigma_{\mathbf{y}}\|_F$. Interestingly, although the ML algorithm is also covariance based, it still outperforms MMV-AMP. In Figure 3.3 we see that beyond the CS regime, the performance of MMV-AMP and MF significantly deteriorates, while the activity detection error probability of ML and NNLS still decays exponentially with M . In Figure 3.4 we compare the LSFC estimation performance of the ML and NNLS algorithms. The simulations confirm Corollary 4 and show that the relative ℓ_1 recovery error of NNLS indeed decays like $1/\sqrt{M}$. We see that the same decay behavior holds for the ML algorithm only with significantly better constants. Note, that the number of required antennas for the ML algorithm scales fundamentally different depending on whether $K_a \leq L$ or $K_a > L$. In the first case the probability of error decays a lot faster with increasing M , matching qualitatively the scaling derived in Theorem 11, which states that (up to constant or logarithmic factors) $M = \mathcal{O}((K_a/L)^2)$.

Corollary 4 predicts that, in the limit $M \rightarrow \infty$, the recovery error of NNLS vanishes, as long as the number of active users fulfils condition (3.14). We confirm this behavior empirically in Figure 3.5a, where we solve the NNLS problem (3.25) using the true covariance matrix $\Sigma^\circ = \mathbf{A} \text{diag}(\gamma^\circ) \mathbf{A}^H + N_0 \mathbf{I}_L$ instead of the empirical covariance matrix $\hat{\Sigma}_{\mathbf{y}}$. In this case, $\|\mathbf{d}\|_2 = 0$ in (3.32) and the recovery error should be identically zero when the true vector γ° is K_a -sparse and the system parameters are such that Theorem 13 holds. This is confirmed by Figure 3.5a, showing a quadratic curve, below which the recovery error vanishes. We also observe a very similar behavior for the ML algorithm, (see Figure 3.5b). This suggests that the condition (3.25) is indeed necessary independent of the algorithm.

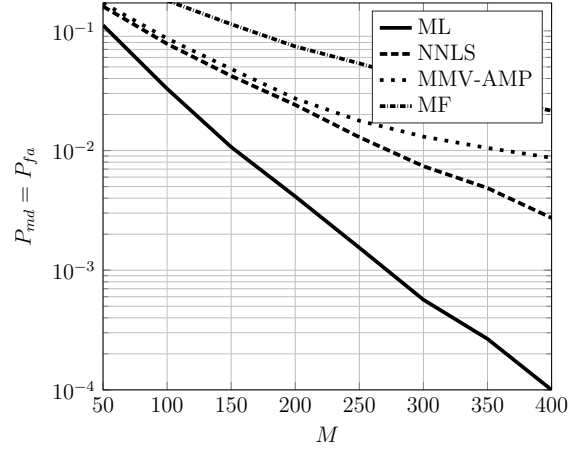


Figure 3.3.: Scaling of the support detection error vs. M beyond the CS regime (i.e. $K_a > L$). Here $K_a = 300, L = 100, K_{\text{tot}} = 2000$ with constant LSFCs at $\text{snr}_k = 0$ dB.

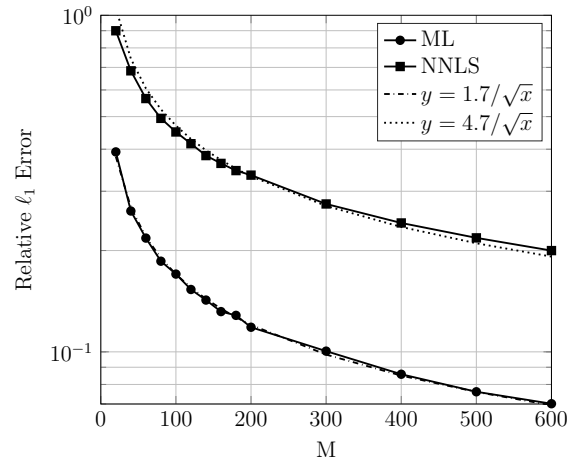


Figure 3.4.: Relative ℓ_1 error of the estimation of the LSFCs of the active users for $D_c = 100, K_a = 200, K_{\text{tot}} = 2000$. The LSFCs are chosen such that snr_k are uniform in the range 0-20dB. The dotted lines show that the curves are well represented by a c/\sqrt{M} behavior, for some constant c , as predicted by Corollary 4.

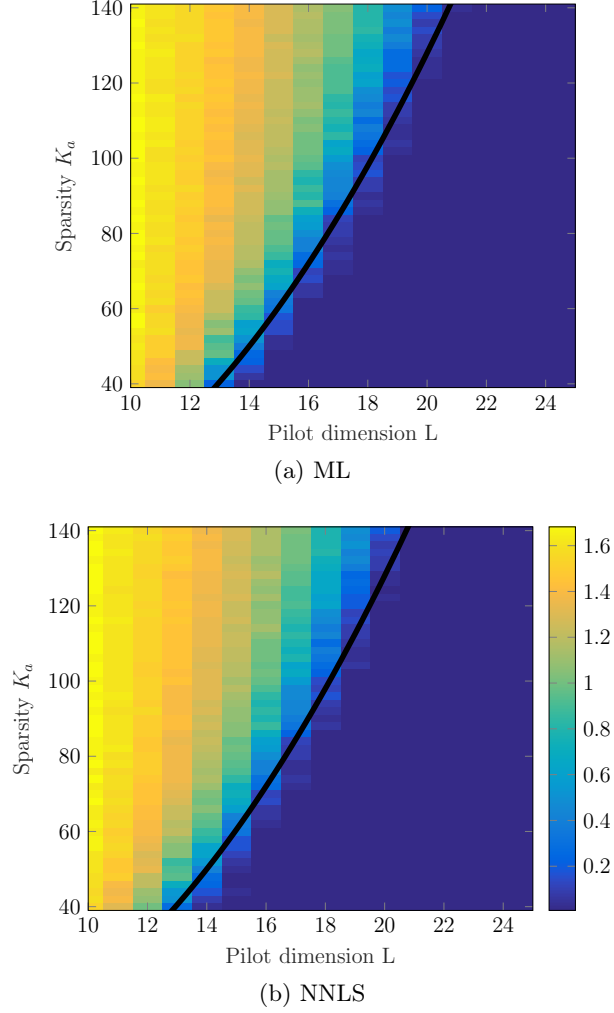


Figure 3.5.: Phase transition of the recovery error for NNLS and ML in the limit $M \rightarrow \infty$ for $K_{\text{tot}} = 1000$. The function $x \rightarrow (x-4)^2/2$ is overlaid in black to emphasize the super-linear scaling. The color indicates the normalized ℓ_1 -error as it is subject of Corollary 4 in the NNLS case. The LSFC are constant and the activity pattern is chosen uniformly at random from all K_a -sparse vectors. The results are obtained by averaging over random pilot matrices and activity patterns.

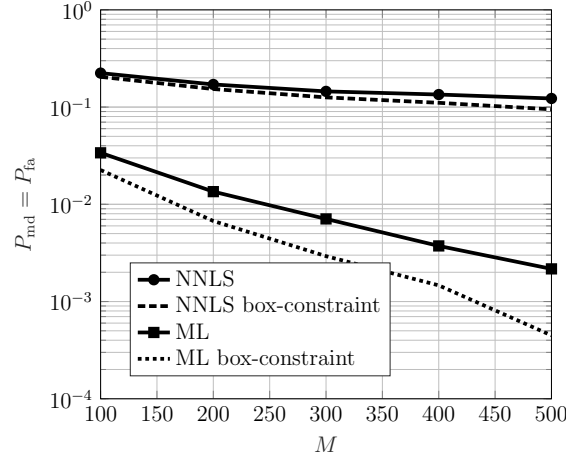


Figure 3.6.: Effect of using the box-constraint (see step 8 in Algorithm 1) when the LSFCs g_k are known at the receiver. Here $K_a = 150, L = 100, K_{\text{tot}} = 2000$ and the LSFCs are distributed such that snr_k are uniform in the range $0 - 20\text{dB}$.

Figure 3.6 shows the gain in performance when the LSFCs are known at the receiver and the box-constraint (step 8 in Algorithm 1) is employed.

3.3. Massive MIMO Unsourced Random Access

In this section the U-RA concept is extended to the case of a block-fading channel with a massive MIMO BS receiver. The algorithms for AD and their analysis from the previous Section are used to show that the ML scheme (see Algorithm 1) provides an efficient inner decoder for the CCS scheme. The presented scaling properties in Corollary 6 allow to estimate the required per-user-power, in terms of E_b/N_0 , and the required number of receive antennas M for reliable transmission.

The channel model is the same as described in Section 3.2.1, i.e., a Rayleigh block-fading channel with deterministic LSFCs and multiple receive antennas. For this Section assume that L spans the whole coherence block-length. Let $n = SL$, for some integer S , such that the transmission of a codeword spans S fading blocks. In contrast to the AD setup, now again a grant-free U-RA setting is considered, i.e. a fixed number K_a of users transmit their messages, picked from a common codebook $\mathcal{C} = \{\mathbf{c}(m) : m \in [2^{nR}]\}$, formed by 2^{nR} , now complex valued, codewords $\mathbf{c}(m) \in \mathbb{C}^n$. A fixed but unknown number K_a of users transmit their messages over multiple coherence blocks. The BS must then produce a list \mathcal{L} of the transmitted messages $\{m_k : k \in K_a\}$ (i.e., the messages of the active users).

For this section the error criterion is slightly relaxed to allow for more flexibility in the decoder design. The system performance is expressed in terms of the *Per-User Probability*

of *Misdetction*, defined as the average fraction of transmitted messages not contained in the list, i.e.,

$$p_{\text{md}}^{\text{msg}} = \frac{1}{K_a} \sum_{k \in \mathcal{K}_a} \mathbb{P}(m_k \notin \mathcal{L}), \quad (3.60)$$

and the *Probability of False-Alarm*, defined as the average fraction of decoded messages that were indeed not sent, i.e.,

$$p_{\text{fa}}^{\text{msg}} = \mathbb{E} \left[\frac{|\mathcal{L} \setminus \{m_k : k \in \mathcal{K}_a\}|}{|\mathcal{L}|} \right]. \quad (3.61)$$

The size of the list is also an outcome of the decoding algorithm, and therefore it is a random variable. As customary, the average error probabilities of false-alarm/misdetction are defined as the expected values of $p_{\text{fa}}^{\text{msg}}/p_{\text{md}}^{\text{msg}}$ resp. over all involved random variables. That is in this case the Rayleigh fading coefficients, the AWGN noise and the choice of messages, where the messages are assumed to be chosen uniformly and independent of each other. Notice, that in this problem formulation the total number of users K_{tot} is completely irrelevant, as long as it is much larger than the number of active user K_a (e.g., we may consider $K_{\text{tot}} = \infty$).

3.3.1. Unsourced Random Access as AD Problem

First, we analyse the case $S = 1$, i.e. each user transmits his codeword in a single block of length L . This model is also known as the quasi-static fading AWGN [28], in contrast to the block-fading model, where a message is encoded over multiple independent fading blocks. Further fix $J = LR$ and let $\mathbf{A} \in \mathbb{C}^{L \times 2^J} = [\mathbf{a}_1, \dots, \mathbf{a}_{2^J}]$, be a matrix with columns normalized such that $\|\mathbf{a}_i\|_2^2 = L$. Each column of \mathbf{A} represents one codeword. Let i_k denote the J -bit messages produced by the active users $k \in \mathcal{K}_a$, represented as integers in $[1 : 2^J]$, user k simply sends the column \mathbf{a}_{i_k} of the coding matrix \mathbf{A} . The received signal at the M -antennas BS takes on the form

$$\begin{aligned} \mathbf{Y} &= \sum_{k \in \mathcal{K}_a} \sqrt{P g_k} \mathbf{a}_{i_k} \mathbf{h}_k^\top + \mathbf{Z} \\ &= \sqrt{P} \mathbf{A} \mathbf{\Phi} \mathbf{G}^{1/2} \mathbf{H} + \mathbf{Z} \end{aligned} \quad (3.62)$$

where, as for the AD model in (3.3), $\mathbf{G} = \text{diag}(g_1, \dots, g_{K_{\text{tot}}})$ is the diagonal matrix of LSFCs, $\mathbf{H} \in \mathbb{C}^{K_{\text{tot}} \times M}$ is the matrix containing, by rows, the user channel vectors \mathbf{h}_k formed by the small-scale fading antenna coefficients (Gaussian i.i.d. entries $\sim \mathcal{CN}(0, 1)$), $\mathbf{Z} \in \mathbb{C}^{L \times M}$ is the matrix of AWGN samples (i.i.d. entries $\sim \mathcal{CN}(0, N_0)$), and $\mathbf{\Phi} \in \{0, 1\}^{2^J \times K_{\text{tot}}}$ is a binary selection matrix where for each $k \in \mathcal{K}_a$ the corresponding column $\mathbf{\Phi}_{:,k}$ is all-zero

but a single one in position i_k , and for all $k \in [K_{\text{tot}}] \setminus \mathcal{K}_a$ the corresponding column $\Phi_{:,k}$ contains all zeros.

$\mathbf{X} = \sqrt{P}\Phi\mathbf{G}^{1/2}\mathbf{H}$ is a matrix of dimension $2^J \times M$. The r -th row of such a matrix is given by

$$\mathbf{X}_{r,:} = \sum_{k \in \mathcal{K}_a} \sqrt{Pg_k \phi_{r,k}} \mathbf{h}_k^\top, \quad (3.63)$$

where $\phi_{r,k}$ is the (r,k) -th element of Φ , equal to one if $r = i_k$ and zero otherwise. It follows that $\mathbf{X}_{r,:}$ is Gaussian with i.i.d. entries $\sim \mathcal{CN}(0, \sum_{k \in \mathcal{K}_a} g_k \phi_{r,k})$. Since the messages are uniformly distributed over $[1 : 2^J]$ and statistically independent across the users, the probability that $\mathbf{X}_{r,:}$ is identically zero is given by $(1 - 2^{-J})^{K_a}$. Hence, for 2^J significantly larger than K_a , the matrix \mathbf{X} is row-sparse.

In order to map the decoding into a problem completely analogous to the AD problem already discussed before, with some abuse of notation the modified LSFC-activity coefficients are defined as $\gamma_r := \sum_{k \in \mathcal{K}_a} Pg_k \phi_{r,k}$ and $\mathbf{\Gamma} = \text{diag}(\gamma_1, \dots, \gamma_{2^J})$. Then, (3.62) can be written as

$$\mathbf{Y} = \mathbf{A}\mathbf{\Gamma}^{1/2}\tilde{\mathbf{H}} + \mathbf{Z}, \quad (3.64)$$

where $\tilde{\mathbf{H}} \in \mathbb{C}^{2^J \times M}$ with i.i.d. elements $\sim \mathcal{CN}(0, 1)$. Notice that in (3.64) the number of total users K_{tot} plays no role. In fact, none of the matrices involved in (3.64) depends on K_{tot} .

The task of the inner decoder at the BS is to identify the non-zero elements of the modified active LSFC pattern $\boldsymbol{\gamma}$, the vector of diagonal coefficients of $\mathbf{\Gamma}$. The active (non-zero) elements correspond to the indices of the transmitted messages. Notice that even if two or more users choose the same sub-message, the corresponding modified LSFC γ_r is positive since it corresponds to the sum of the signal powers. In other words, since the detection scheme is completely non-coherent (it never explicitly estimates the complex channel matrix) and active signals add in power, there is no risk of signal cancellation or destructive interference.

At this point, it is clear that the problem of identifying the set of transmitted messages from observation (3.64) is completely analogous to the AD problem from the observation in (3.3), where the role of the total number of users K_{tot} in the AD problem is replaced by the number of messages 2^J in the inner decoding problem. Building on this analogy, we shall use the discussed ML algorithm to decode the inner code.

It is interesting to notice that the modified LSFCs in $\boldsymbol{\gamma}$ are random sums of the individual user channel gains $\{g_k\}$. Hence, even if the g_k 's were exactly individually known, or their statistics was known, these random sums would have unknown values and unknown statistics (unless averaging over all possible active subsets, which would involve an

exponential complexity in K_{tot} which is clearly infeasible in our context). Hence, Bayesian approaches such as MMV-AMP (see Section 3.2.3) as advocated in [4, 49, 118, 121] do not find a straightforward application here. In contrast, the proposed non-Bayesian approaches (in particular, the ML algorithm in Algorithm 1), that treats γ as a deterministically unknown vector.

Notice also that in a practical unsourced random access scenario such as a large-scale IoT application, the slot dimension L may be of the order of 100 to 200 symbols, while for a city-wide IoT data collector it is not unreasonable to have M of the order of 500 to 1000 antennas (especially when considering narrowband signals such as in LoRA-type applications [122, 123]). This is precisely the regime where we have observed a critical behavior of MMV-AMP, while our algorithm uniformly improves as M increases, for any slot dimension L .

3.3.2. Discussion and Analysis for One-Slot Transmission

In this section we discuss the performance of the ML decoder in a single slot ($S = 1$). For the sake of simplicity, in the discussion of this section we assume $g_k = 1$ for all k . In this case, the SNR P/N_0 is also the SNR at the receiver, for each individual (active) user.

Corollary 4 shows that, if the coding matrix \mathbf{A} is chosen randomly, the probability of an error in the estimation of the support of γ vanishes in the limit $M \rightarrow \infty$ for any SNR $\frac{P}{N_0} > 0$ as long as $K_a = \mathcal{O}(L^2 / \log^2(e2^J/L^2))$. Then Corollary 4 gives the following bound for the reconstruction error of

$$\frac{\|\gamma - \gamma^*\|_1}{\|\gamma\|_1} \leq \kappa \left(1 + \left(K_a \frac{P}{N_0} \right)^{-1} \right) \sqrt{\frac{K_a}{M}} \quad (3.65)$$

where κ is some universal constant and γ^* denotes the estimate of γ by the NNLS algorithm (see Section 3.2.2). Our numerical results (Section 3.2.3) suggest that the reconstruction error of the ML algorithm is at least as good as that of NNLS (in practice it is typically *much better*). This bound is indeed very conservative. Nevertheless, this is enough to give achievable scaling laws for the probability of error of the inner decoder. It follows from (3.65) that $\frac{\|\gamma - \gamma^*\|_1}{\|\gamma\|_1} \rightarrow 0$ for $(M, K_a, \frac{P}{N_0}) \rightarrow (\infty, \infty, 0)$ as long as

$$\frac{K_a(1 + (K_a P/N_0)^{-1})^2}{M} = o(1), \quad (3.66)$$

which is satisfied if M grows as

$$M = \max(K_a, (P/N_0)^{-1})^\kappa \quad (3.67)$$

for some $\kappa > 1$. Assuming that J scales such that $2^J = \delta L^2$ for some fixed $\delta \geq 1$, i.e. $J = \mathcal{O}(\log L)$, then the condition in Corollary 4 becomes $K_a = \mathcal{O}(L^2)$ and we can conclude that the recovery error vanishes for sum spectral efficiencies up to

$$\frac{K_a J}{L} = \mathcal{O}(L \log L). \quad (3.68)$$

This shows that we can achieve a total spectral efficiency that grows without bound, by encoding over larger and larger blocks of dimension L , as long as the number of messages per user and the number of active users both grow proportionally to L^2 and the number of BS antennas scales as in (3.67). The achievable sum spectral efficiency grows as $L \log(L)$ and the error probability can be made as small as desired, for any given $E_b/N_0 > 0$. Of course, in this regime the rate per active user vanishes as $\log(L)/L$.

We wish to stress again that this system is completely non-coherent, i.e., there is no attempt to either explicitly (via pilot symbols) or implicitly to estimate the channel matrix (small-scale fading coefficients).

3.3.3. Coding over Multiple Coherence Blocks

In practice it is not feasible to transmit even small messages (e.g. $J \sim 100$) within one coherence block ($S = 1$), because the number of columns of the coding matrix \mathbf{A} grows exponentially in J . Aside from the computational complexity L may also be limited physically by the coherence time of the channel. In both cases it is necessary to transmit the message over multiple blocks. Let each user transmit his message over a *frame* of S fading blocks and within each block use the code described in Section 3.3.1 as *inner* code with the ML decoder as inner decoder.

We can again use the outer tree code described in Section 2.6.1. Let $B = nR$ denote the number of bits per user message. Given J and the slot length L , the inner code is used to transmit in sequence the S (outer-encoded) blocks forming a frame. Let $\mathbf{A} \in \mathbb{C}^{L \times 2^J}$ be the coding matrix as defined in Section 3.3.1. Each column of \mathbf{A} now represents one inner codeword. Letting $i_k(1), \dots, i_k(S)$ denote the sequence of S (outer-)encoded J -bit messages produced by the outer encoder of active user $k \in \mathcal{K}_a$. The user k now simply sends in sequence, over consecutive slots of length L , the columns $\mathbf{a}_{i_k(1)}, \mathbf{a}_{i_k(2)}, \dots, \mathbf{a}_{i_k(S)}$ of the coding matrix \mathbf{A} . As described in Section 3.3.1, the inner decoding problem is equivalent to the AD problem (3.64). Note, that in contrast to the sparse regression structure used in Chapter 2, here the sub-message in each slot is independently encoded by the same inner coding matrix \mathbf{A} .

For each subslot s , let $\hat{\gamma}[s] = (\hat{\gamma}_1[s], \dots, \hat{\gamma}_{2^J}[s])^\top$ denote the ML estimate of γ in subslot s obtained by the inner decoder. Then, the list of active messages at subslot s is defined

as

$$\mathcal{S}_s = \{r \in [2^J] : \hat{\gamma}_r[s] \geq \nu_s\}, \quad (3.69)$$

where ν_1, \dots, ν_S are suitable pre-defined thresholds. Let $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_S$ the sequence of lists of active subblock messages. Since the subblocks contain parity bits with parity profile $\{0, p_2, \dots, p_S\}$, not all message sequences in $\mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_S$ are possible and the outer tree decoder described in Section 2.6.1 is used to create a list of messages that fulfill all parity checks.

3.3.4. Asymptotic Analysis of the Concatenated Code

It was shown in Section 2.6 that outer code rates up to $1 - \alpha^{-1}$ are achievable and optimal with the tree code in the scaling regime where $J = \alpha \log_2 K_a$. The resulting achievable sum spectral efficiency can be calculated as in Section 3.3.2 with a subtle but important difference, because the results on the outer code are valid only in the logarithmic regime $J = \alpha \log_2 K_a$, i.e. $2^J = K_a^\alpha$ for $\alpha > 1$. According to Corollary 4 the error probability of the inner code vanishes if the number of active users scale no faster than $K_a = \mathcal{O}(L^2 / \log^2(e2^J/L^2))$. Using the scaling condition $J = \alpha \log_2 K_a$ and that $K_a \leq L^2$, this implies that in the logarithmic regime the error probability of the inner code vanishes if the number of active users scales as $K_a = \mathcal{O}(L^2 / \log^2(L))$. This gives a sum spectral efficiency of

$$\frac{K_a R_{\text{out}} J}{L} = \mathcal{O}\left(\frac{K_a \log K_a}{L}\right) = \mathcal{O}\left(\frac{L}{\log L}\right). \quad (3.70)$$

The order of this sum spectral efficiency is, by a factor $\log^2 L$, smaller than the one we calculated in Section 3.3.2. This is because the order of supported active users is smaller by exactly the same $\log^2 L$ factor. In Section 3.3.2 we assumed that J scales as $2^J = \delta L^2 = \mathcal{O}(K_a)$ for some $\delta > 1$, so that the ratio $K_a/2^J$ remains constant. It is not clear from the analysis in [87], whether the probability of error of the outer tree code would vanish in the regime. We can get a converse by evaluating the entropy bound (2.57) from Section 2.6. Let $2^J = \delta K_a$ with $\delta > 1$, then $(1 - 2^{-J})^{K_a} = (1 - \delta/K_a)^{K_a} \xrightarrow{K_a \rightarrow \infty} \exp(-\delta)$. Therefore the binary entropy $\mathcal{H}_2((1 - 2^{-J})^{K_a})$ remains a constant in the limit $J, K_a \rightarrow \infty$ and we get that

$$K_a R_{\text{out}} J \leq \delta K_a \mathcal{H}_2(\exp(-\delta)). \quad (3.71)$$

This shows that $R_{\text{out}} \rightarrow 0$ in the limit $J, K_a \rightarrow \infty$ is the best achievable asymptotic per-user outer rate, but the outer sum rate $K_a R_{\text{out}} J$ is proportional to K_a . The resulting sum

spectral efficiencies scale as

$$\frac{K_a R_{\text{out}} J}{L} = \mathcal{O}\left(\frac{K_a}{L}\right) = \mathcal{O}(L). \quad (3.72)$$

This means it could be possible to increase the achievable sum spectral efficiencies by a factor of $\log L$ by using an outer code that is able to achieve the entropy bound (2.57) in the regime $2^J = \delta K_a$. It is not clear though whether the code of [87] or some other code can achieve this.

3.3.5. Simulations

The outer decoder requires a hard decision on the support of the estimated $\hat{\gamma}[s]$. The decision problem is slightly different than the one discussed in Section 2.5 because the users can potentially have different channel gains and an estimate of the LSFCs is available at the receiver. When K_a is known, one approach consists of selecting the $K_a + \Delta$ largest entries in each section, where $\Delta \geq 0$ can be adjusted to balance between false alarm and misdetection in the outer channel. However, the knowledge of K_a is a very restrictive assumption in such type of systems. An alternative approach, which does not require this knowledge, consists of fixing a sequence of thresholds $\{\nu_s : s \in [S]\}$ and let $\rho[s]$ to be the binary vector of dimension 2^J with elements equal to 1 for all components of $\hat{\gamma}[s]$ above the threshold ν_s . By choosing the thresholds, one can balance between missed detections and false alarms. Nonetheless, for simplicity we stick to the first approach and choose $\Delta = 50$.

For the simulations in Figure 3.7 we choose $B = 96$ bits as payload size for each user, a frame of $S = 32$ slots of $L = 100$ dimensions per slot, yielding an overall block length $n = 3200$. Choosing the binary subblock length $J = 12$, the inner coding matrix \mathbf{A} has dimension 100×4096 and therefore is still quite manageable. We choose the columns of \mathbf{A} uniformly i.i.d. from the sphere of radius \sqrt{L} . For the outer code, we choose the following parity profile $p = [0, 9, 9, \dots, 9, 12, 12, 12]$, yielding an outer coding rate of $R_{\text{out}} = 0.25$ information bits per binary symbol. Notice also that if one wishes to send the same payload message using the piggybacking scheme of [49, 121], each user should make use of 2^{96} columns, which is totally impractical. All large scale fading coefficients are fixed to $g_k \equiv 1$. In Figure 3.7 we fix $N_0 = 1$ and choose the transmit power (energy per symbol), such that $E_b/N_0 = 0$ dB and plot the sum of the two types of message error probabilities $P_e = p_{\text{md}}^{\text{msg}} + p_{\text{fa}}^{\text{msg}}$, (see (3.60) and (3.61)) as a function of the number of active users for different numbers of receive antennas M . Figure 3.8 shows how P_e falls as a function of E_b/N_0 for different values of K_a and M . Figure 3.9 shows the required values of E_b/N_0 as a function of K_a to achieve a total error probability $P_e < 0.05$ for the code parameters in Table 3.1. We use three different settings here, depending on the

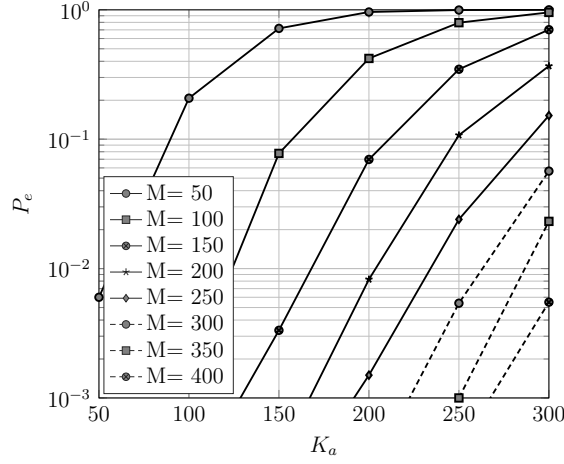


Figure 3.7.: Error probability ($P_e = p_{\text{md}}^{\text{msg}} + p_{\text{fa}}^{\text{msg}}$) as a function of the number of active users for different numbers of receive antennas. $E_b/N_0 = 0$ dB, $L = 100$, $n = 3200$, $B = 96$ bits, $S = 32$, $J = 12$.

values of the coherence block-length L . In all three settings the total block-length is fixed to $n = 3200$ and $B \approx 100$, which gives a per-user spectral efficiency of $R \approx 0.031$ bits per channel use. With $K_a = 300$ the total spectral efficiency is $\mu \approx 9$ bits per channel use, which is significantly larger than today's LTE cellular systems (in terms of bit/s/Hz per sector) and definitely much larger than IoT-driven schemes such as LoRA [122, 123]. According to the Shannon-limit for the scalar Gaussian multiple access channel (only one receive antenna, no fading) $E_b/N_0 > (2^{K_a R} - 1)/(K_a R)$, and therefore each user needs at least ≈ 17.5 dB to achieve a total spectral efficiency of 9 bits per channel use. Here we find that gains of 20 dB or more are possible even with non-coherent detection by the use of multiple receive antennas. The simulations confirm qualitatively the behavior predicted in Sections 3.3.2 and 3.3.4. The performance of MF under the different settings in Table 3.1 is depicted Figure 3.10 and the performances of ML, MMV-AMP and MF are compared in Figure 3.11. The ML algorithm is consistently superior to the other two. The achievable total spectral efficiencies seem to be mainly limited by the coherence block-length L , and for a given total spectral efficiency the required energy-per-bit can be made arbitrary small by increasing M . This shows also quantitatively that the non-coherent massive MIMO channel is very attractive for unsourced random access, since it preserves the same desirable characteristics of unsourced random access as in the non-fading Gaussian model of [9] (users transmit without any pre-negotiation, and no use of pilot symbols is needed), while the total spectral efficiency can be made as large as desired simply by increasing the number of receiver antennas.

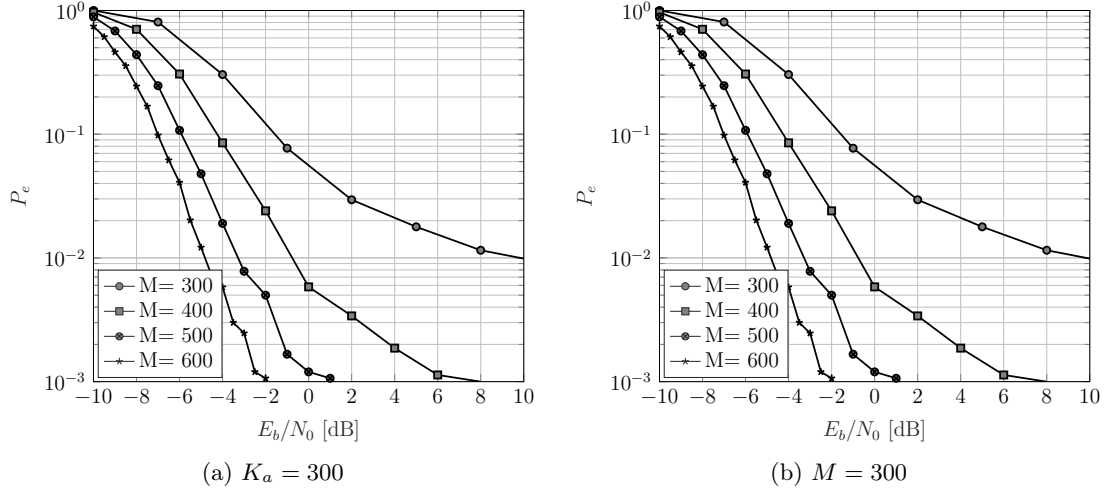


Figure 3.8.: Error probability ($P_e = p_{\text{md}}^{\text{msg}} + p_{\text{fa}}^{\text{msg}}$) as a function of E_b/N_0 with ML as inner decoder. $L = 100$, $n = 3200$, $b = 96$ bits, $S = 32$, $J = 12$.

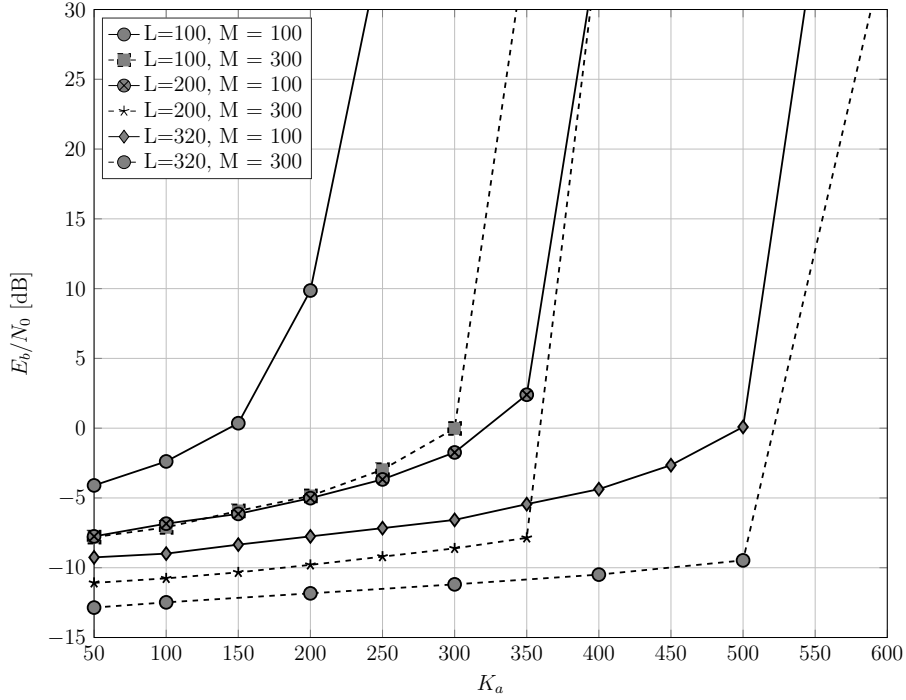


Figure 3.9.: Required energy-per-bit to achieve $P_e < 0.05$ with ML as inner decoder. L and S are varied, while $n = 3200$ and $B \approx 100$ are fixed. The precise parameters are given in Table 3.1.

| | S | J | R_{out} | Parity profile | B |
|-----------|----|----|------------------|-----------------------------|-----|
| $L = 100$ | 32 | 12 | 0.25 | $[0,9,\dots,9,12,12,12]$ | 96 |
| $L = 200$ | 16 | 15 | 0.42 | $[0,7,8,8,9,\dots,9,13,14]$ | 100 |
| $L = 320$ | 10 | 19 | 0.52 | $[0,9,\dots,9,19]$ | 99 |

Table 3.1.: Parameters for Figure 3.9

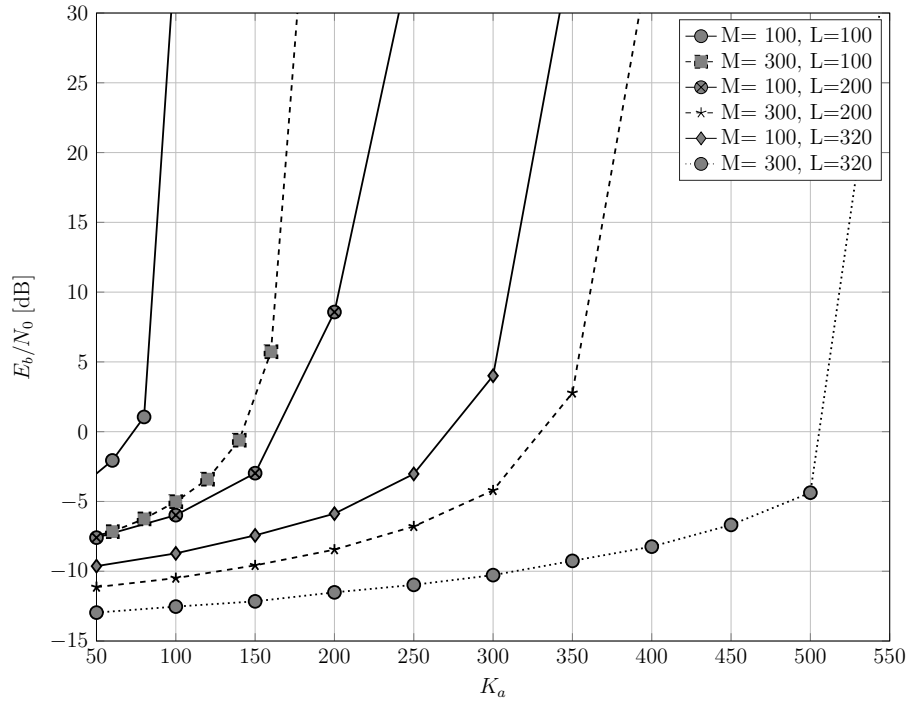


Figure 3.10.: Required energy-per-bit to achieve $P_e < 0.05$ with a simple matched filter approach as inner decoder.

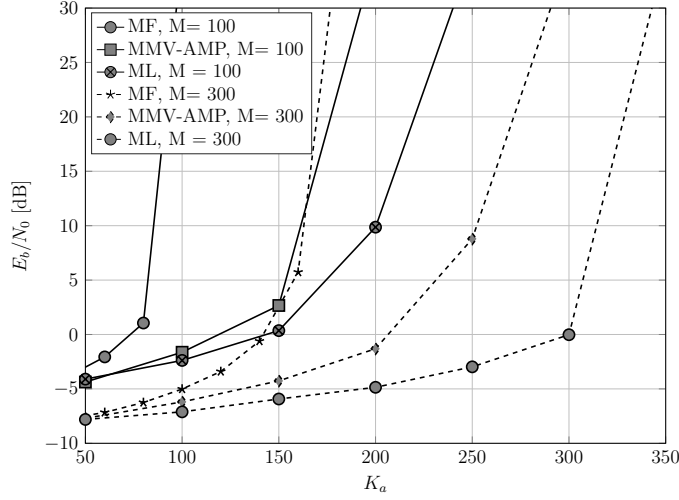


Figure 3.11.: Required energy-per-bit to achieve $P_e < 0.05$. Comparison of MF, MMV-AMP and ML as inner decoder. $L = 100$. The precise parameters are given in Table 3.1.

3.3.6. Impact of Large-Scale Fading

In a wireless random-access scenario the received LSFCs of the active users may vary over several orders of magnitude. Nonetheless, the assumptions of constant LSFCs, known at the receiver, can be justified by assuming an initial power control phase where users adjust their power based on feedback from the BS until an acceptable power distribution is reached. In many of the proposed pilot based MU-MIMO grant-free random-access schemes, e.g. [3–7, 49], power control is crucial to avoid the near-far effect and allow for reliable communication. Note, that such a power control phase goes against the basic idea behind grant-free random access, that is to reduce the access latency when messages are short and access is very sporadic. Especially in a high-mobility scenario, as it is typical when the coherence times are small, the LSFCs of the users will change between the sporadic activations, which would require to do power control before each activation slot. This would significantly increase the transmission delay and may even destroy the gains of the grant-free approach. Furthermore, a high dynamic range at the transmitter comes with increased battery consumption and a higher hardware complexity, which goes against the underlying idea of U-RA to enable low-cost low-energy transmitters.

To investigate the effect of a more realistic LSFC distribution on the CCS U-RA scheme we assume the following widely used model for the LSFCs [124]:

$$g_k[\text{dB}] = -\alpha - 10\beta \log_{10}(d_k) + \sigma_{\text{shadow}}z \quad (3.73)$$

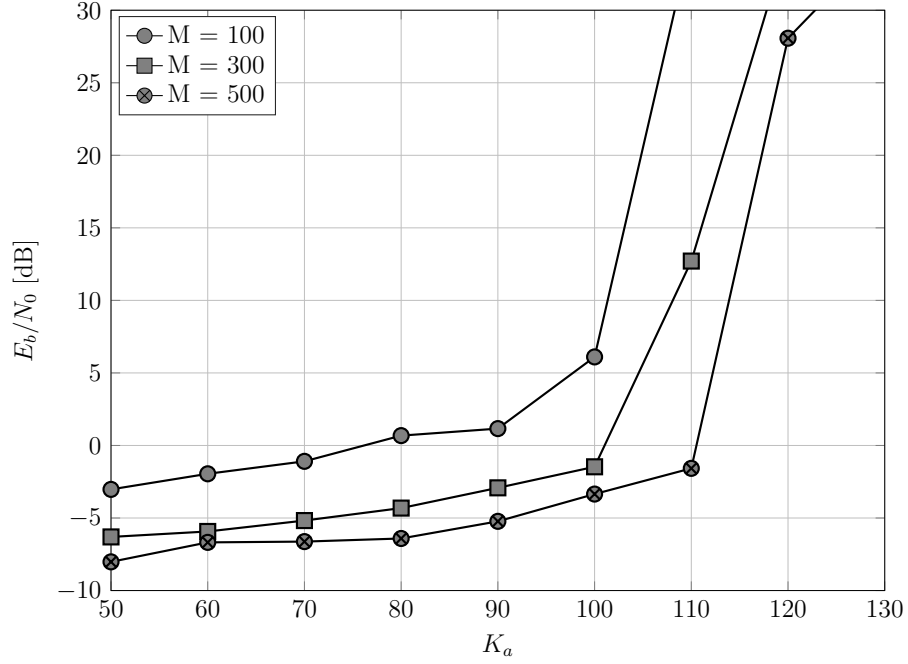


Figure 3.12.: Impact of LSF: Required energy-per-bit to achieve $P_e < 0.05$ with $L = 100$ and $n = 3200$ with the ML algorithm. The LSFCs are chosen according to the model (3.73) with users distributed uniformly on a circle of radius 1km.

where d_k is the distance of user k to the BS in km and $z \sim \mathcal{N}(0, 1)$. For the simulations in Figure 3.12 we set $\alpha = 100$, $\beta = 3.76$ and $\sigma_{\text{shadow}}^2 = 8$. For the noise density we assume a spectral density of -170 dBm/Hz and a bandwidth of 10 MHz. The positions of the active users are assumed to be distributed uniformly around the BS in a circle of radius 1 km. We fix $L = 100$ and the remaining parameters are chosen as in Table 3.1.

The simulations show that the AD performance of the ML algorithm decreases significantly in the presence of severe power imbalances, but only in the overloaded regime $K_a > L$. A possible explanation is that interference cancellation, which is implicitly done by the ML algorithm and is generally known to improve performance under power imbalances, cannot be done perfectly when channel estimation is not possible, as it happens in the overloaded regime $K_a > L$. In Section 3.5 I will discuss an approach to overcome this limitation.

The MF and the MMV-AMP algorithm are not considered in the faded setting because their performance is very bad. For the MF algorithm this is expected, because it suffers heavily under a near far effect, which is why the weakest 20% of the users are almost never correctly recovered. The instability of the MMV-AMP algorithm under varying receive powers has already been discussed in Section 3.2.3.

3.4. Massive MIMO Unsourced Random Access - Slow Fading

Although we have shown in the previous section that the sum spectral efficiency can grow unbounded with the coherence block length L , in practice, L is limited by the coherence time of the channel. In typical wireless systems the coherence block-length n may range from a couple of hundred to a couple of thousand, depending mainly on the speed of the transmitters. At a carrier frequency of 2 GHz the coherence times, according to the model $T_c \approx 1/(4D_s)$ [45], range from 45 ms at 3 km/h to 1 ms at 120 km/h. The coherence bandwidth depends on the maximal delay spread and, in an outdoor environment, typically ranges from 100 to 500 kHz depending on the propagation conditions. Therefore, the number of OFDM symbols in a coherence block may range from 100 to 20000, depending mainly on the assumed speed and the geometry of the environment. Unfortunately, the ML algorithm of Section 3.2.2 has a run-time complexity that scales with L^2 , which makes it unfeasible to use at $L \gtrsim 300$.

In this section we present a conceptually simple algorithm that can be used when $L > K_a$. It is based on pilot transmission, AD, channel estimation, MRC and single-user decoding, very similar to the state-of-the art approach for massive MIMO grant-free random access [3–5]. In contrast to the scheme with fixed pilots allocated to all users, we use a pool of non-orthogonal pilots from which active users pick one pseudo-randomly based on the first bits of their message.

We show that a collision of users, i.e. two users picking the same pilot sequence, can be resolved by using a polar single-user code with a successive-cancellation-list (SCL) decoder [73–75]. Finite-length simulations show that the performance of the coding scheme can be well predicted by analytical calculations. Despite its simplicity the suggested scheme has an energy efficiency that is comparable to existing approaches.

Note, that the problem treated here is formally almost equivalent to grant-free random-access with fixed pilots allocated to each user. Differences arise only in the possibility of collisions and the associated use of an list decodable single-user code. The error probability of AD and MRC in the asymptotic limit $K_a, K_{\text{tot}}, L \rightarrow \infty$ with fixed ratios K_a/K_{tot} and K_{tot}/L has been analysed in [5]. In this work we focus on the finite-blocklength regime and the combination of MRC with a single-user polar code.

3.4.1. Pilot-based Massive MIMO U-RA

In this section it is assumed, again, that $S = 1$, i.e. the complete message is transmitted in one coherence block of length L . Let the coherence block of length n be divided into two periods of lengths L and n_d . In the first period each users chooses one of $N = 2^J$ (non-orthogonal) pilot sequences based on the first J bits of its message. The received

signal in the identification phase is the same as in (3.3). The BS uses an AD algorithm from Section 3.2 to estimate the indices $\hat{\mathcal{I}}$ of the set \mathcal{I} of used pilots and the corresponding LSFCs. Then a linear MMSE estimate of the channel matrix is computed as

$$\hat{\mathbf{H}} = \hat{\mathbf{\Gamma}}_{\hat{\mathcal{I}}}^{1/2} \mathbf{A}_{\hat{\mathcal{I}}}^H \left(\mathbf{A}_{\hat{\mathcal{I}}} \hat{\mathbf{\Gamma}}_{\hat{\mathcal{I}}} \mathbf{A}_{\hat{\mathcal{I}}}^H + N_0 \mathbf{I}_L \right)^{-1} \mathbf{Y}_p \quad (3.74)$$

where $\mathbf{A}_{\hat{\mathcal{I}}}$ denotes a sub-matrix of the pilot matrix \mathbf{A} which contains only the columns which have been estimated as active and $\hat{\mathbf{\Gamma}}_{\hat{\mathcal{I}}}$ contains the LSFCs of the active users on the diagonal. In the second period each users encodes its remaining $B - J$ -bit message with a binary $(B - J, 2n_d)$ block code and modulates the $2n_d$ coded bits via QPSK on a sequence of n_d complex symbols \mathbf{s}_k . These are transmitted over the n_d channel uses in the second phase. The matrix of received signals in the second phase is

$$\mathbf{Y}_d = \sum_{k=1}^{K_{\text{tot}}} \sqrt{P_{\text{data}} \gamma_k} \mathbf{s}_k \mathbf{h}_k + \mathbf{Z}_d. \quad (3.75)$$

The BS uses the channel estimate $\hat{\mathbf{H}}$ from the first phase to perform multiuser detection via maximum-ratio-combining (MRC) [46], i.e. it computes

$$\hat{\mathbf{S}} = \hat{\mathbf{\Gamma}}_{\hat{\mathcal{I}}}^{-1/2} \hat{\mathbf{H}} \mathbf{Y}_d^H \quad (3.76)$$

The rows of $\hat{\mathbf{S}}$ correspond to noisy estimates of the transmitted sequences \mathbf{s}_k . Note, that it is also possible to use zero-forcing [46] instead of MRC but this would require that $M > K_a$. The rows of $\hat{\mathbf{S}}$ are individually demodulated, the bit-wise log-likelihood ratios are computed and fed into a soft-input single-user decoder. If the decoder finds a valid codeword, the index of the corresponding pilot is converted back to bits and the concatenation of the two parts is added to the output list. The use of a polar code with CRC-bits and a successive-cancellation-list decoder has the additional benefit that we can include all the valid codewords in the output list of the SCL decoder in the U-RA output list. This allows to recover the messages of colliding users which have chosen the same pilot in the first phase. The ability of polar codes to resolve sums of codewords has been observed and used for U-RA on the AWGN in combination with spreading sequences [26] and a slotted Aloha approach [25, 125].

3.4.2. Analysis

In this section we calculate a finite-blocklength lower bound on the error probability and on the energy efficiency of the MRC approach. We assume that the AD and LSFC estimation can be done without errors. This gives a lower bound on the error probability and in the

regime where $L > K_a$ we expect it to be tight, as in this regime the AD error rates and the error of the LSFC estimation are very low [4, 82]. For simplicity we consider $P = P_{\text{data}}$ here. The covariance of the channel estimation error of the LMMSE estimation in (3.74) is given by

$$\mathbf{C}_e = \mathbf{I}_{K_a} - \mathbf{\Gamma}_{\mathcal{I}}^{1/2} \mathbf{A}_{\mathcal{I}}^H \left(\mathbf{A}_{\mathcal{I}} \mathbf{\Gamma}_{\mathcal{I}} \mathbf{A}_{\mathcal{I}}^H + N_0 \mathbf{I}_L \right)^{-1} \mathbf{A}_{\mathcal{I}} \mathbf{\Gamma}_{\mathcal{I}}^{1/2} \quad (3.77)$$

and the MSE for each active user $k \in \mathcal{K}_a$ is given by $\text{mse}_k = \mathbb{E}\{|h_{k,m} - \hat{h}_{k,m}|^2\} = (\mathbf{C}_e)_{k,k}$. If the pilots would be orthogonal this error reduces to

$$\text{mse}_k = \frac{N_0}{N_0 + L\gamma_k} \quad (3.78)$$

which lower bounds the error and will show, in the simulations, to be a sufficiently tight approximation for practical random non-orthogonal pilots. A lower bound on the effective SINR of each user after MRC is given by [46]

$$\text{SINR}_k \geq \frac{M\gamma_k(1 - \text{mse}_k)}{N_0 + \sum_{k=1}^{K_a} \gamma_k} \quad (3.79)$$

An approximation of the achievable rates of a block-code with block-length $2n_d$ and error probability p_e on a real AWGN channel with power SINR is given by the normal approximation [21]

$$R \approx 0.5 \log(1 + \text{SINR}) - \sqrt{\frac{V}{2n_d}} Q^{-1}(p_e) \quad (3.80)$$

where

$$V = \frac{\text{SINR}}{2} \frac{\text{SINR} + 2}{(\text{SINR} + 1)^2} \log^2 e \quad (3.81)$$

and $Q(\cdot)$ is the Q-function. Using the normal approximation we can find the required SINR to achieve a certain error probability at a given block-length and then we can find the required input power to achieve the target SINR.

3.4.3. Complexity

For the large coherence block-lengths considered here we use the MMV-AMP algorithm described in Section 3.2.3. The complexity of AD with the modified MMV-AMP algorithm is in the order of $\mathcal{O}(MN \log N)$ when the pilots are chosen as the columns of a randomly sub-sampled DFT matrix, which allows to replace the matrix multiplications in the AMP iterations by Fast-Fourier-Transforms. Also we employ the approximate calculation of the derivatives in the MMV-AMP as described in Section 3.2.3 which reduces the complexity of MMV-AMP from $\mathcal{O}(M^2 N \log N)$ to $\mathcal{O}(MN \log N)$.

3.4.4. Simulations

For the simulation in Figure 3.13 we choose $n = 3200$, $P_e = 0.05$, $B = 100$, $L = 1152$, $n_d = 2048$ and $B_p = 16$. Also all LSFCs are considered to be constant $g_k = 1$. We use a polar code [73–75] with a state-of-the-art SCL decoder with 16 CRC bits and a list size of 32. The simulations show an excellent overlap with the theoretical results, although we have neglected the non-orthogonality of the pilots in the analysis. For comparison we add the reported values of the tensor-based-modulation (TBM) approach [126], although the latter have been obtained with the higher value $P_e = 0.1$. For $M = 50$ the results pilot based scheme and TBM show a similar shape, although the TBM approach achieves better results for $K_a \geq 400$. The results show that with only $M = 100$ receive antennas over thousand users can be served concurrently, which leads to sum-spectral efficiencies beyond 30 bits per channel use. This is not surprising, since the scheme essentially resembles pilot based massive MU-MIMO with MRC, which is known to achieve very high spectral efficiencies. Nonetheless, the results show that MMV-AMP provides an algorithm that can scale to thousands of concurrent users even with a large number of non-orthogonal pilots. Furthermore, the combination with a single-user polar code with list decoding can efficiently reduce the effect of pilot collisions, as we will further investigate in the following.

3.4.5. Collisions

The average number of collisions of k users on one pilot is given by

$$\mathbb{E}[C_k] = \frac{\binom{K_a}{k}}{N^{k-1}}. \quad (3.82)$$

We can safely ignore the collisions of more than two users since their number is much smaller than 1 for the considered parameters. For $K_a = 1000$ and $N = 2^{16}$ (3.82) gives an average number of 7 – 8 collisions of order two. If all of the colliding messages would result in an error, this would lead to an per-user error probability of 0.016 on average. This could be incorporated into the above analysis by subtracting this values from the target error probability p_e in the normal approximation (3.80). Nonetheless, if a list decoder is used as a single-user decoder, it is possible to recover both of the colliding messages as demonstrated in Figure 3.14.

In Figure 3.14 we simulate a collisions between two users on a non-fading AWGN channel and visualize the probability distribution of the number of correctly decoded messages. We can see that at a per-user SNR of about -11 dB the probability that both messages are lost drops below 0.01, where in roughly 70% of the cases both codewords are correctly

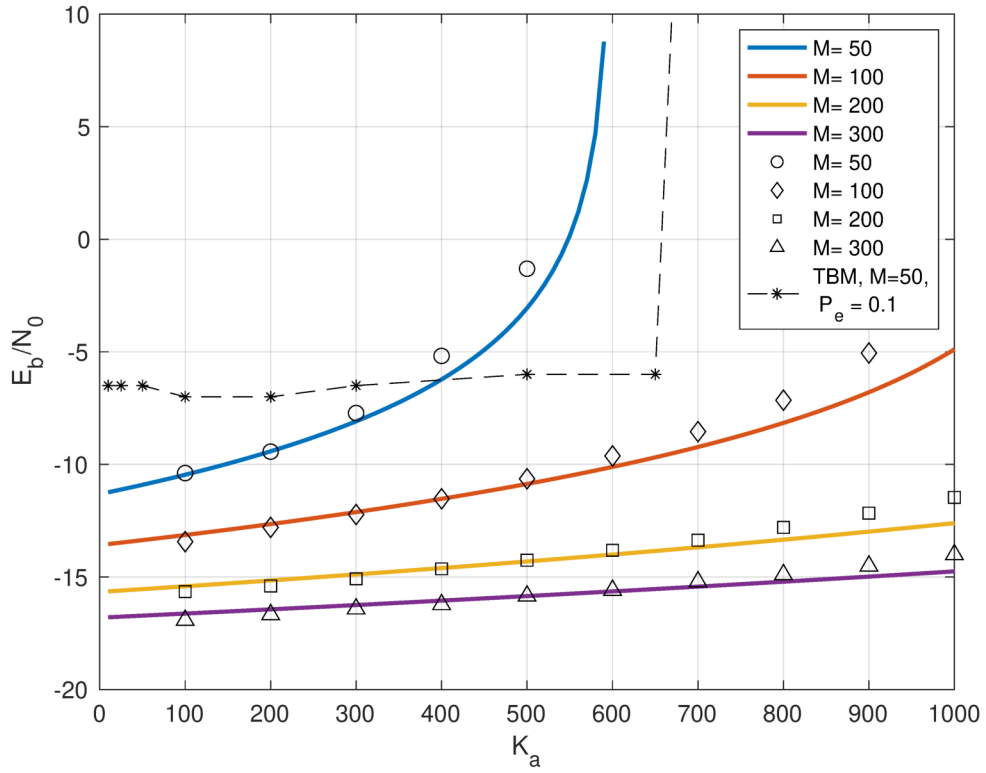


Figure 3.13.: Required energy-per-bit with the MRC approach to achieve $P_e = p_{\text{md}} + p_{\text{fa}} < 0.05$. $n = 3200$ and $B = 100$ are fixed. Solid lines represent the theoretical estimates from Section 3.4.2. The tensor-based-modulation (TBM) scheme from [126] is included for comparison.

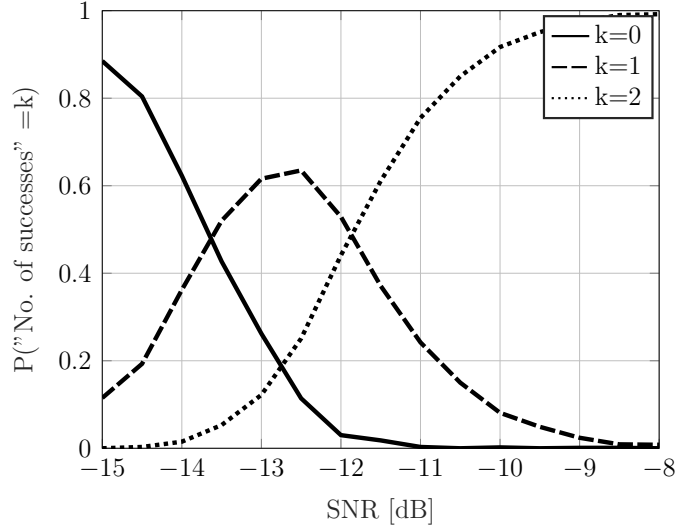


Figure 3.14.: Probability distribution of recovered codewords in the case of a two-user collision in the non-fading case. Polar code with $B - B_p = 84$ message bits, $n_d = 2048$ complex QPSK coded symbols (4096 real BPSK symbols) and SCL decoding with 16 CRC-bits and list-size 32.

recovered. The simulation shows that at high SNR values the SCL decoder can reliably recover both messages. When there is no fading, half of the coded bits are erased on average when two codewords are added. Since the rate of the polar code $R = (B - B_p)/(2n_d)$ is much smaller than $1/2$, these erasures can be recovered, see also [26]. This situation changes when fading is involved. In Figure 3.15 we take the uncertainty of the channel estimation and the MRC into account via the following simplified two user collision model. Let \mathbf{s}_1 and \mathbf{s}_2 denote the QPSK modulated sequences of two users and $\mathbf{h}_1, \mathbf{h}_2 \in \mathbb{C}^M$ their iid Rayleigh channel vectors. We model the channel estimates as $\hat{\mathbf{h}} = \mathbf{h}_1 + \mathbf{h}_2 + \mathbf{e}$ where $\mathbf{e} \sim \mathcal{CN}(0, \sigma_{\text{est}}^2)$ is the channel estimation error with variance σ_{est}^2 . The model for the estimated single-user sequence for both users is then

$$\hat{\mathbf{s}}^\top = \hat{\mathbf{h}}^H (\sqrt{P}\mathbf{h}_1\mathbf{s}_1^\top + \sqrt{P}\mathbf{h}_2\mathbf{s}_2^\top + \mathbf{z}) \quad (3.83)$$

with $\mathbf{z} \sim \mathcal{CN}(0, 1)$. We fix the per-user SNR to -10 dB and vary the variance of the channel estimation error. The simulation in Figure 3.15 shows that the probability of recovering both codewords saturates to a non-zero value, in contrast to the non-fading case. This effect persists, even when the base per-user SNR is increased. Nonetheless, the probability that at least one codeword is recovered converges to 1.

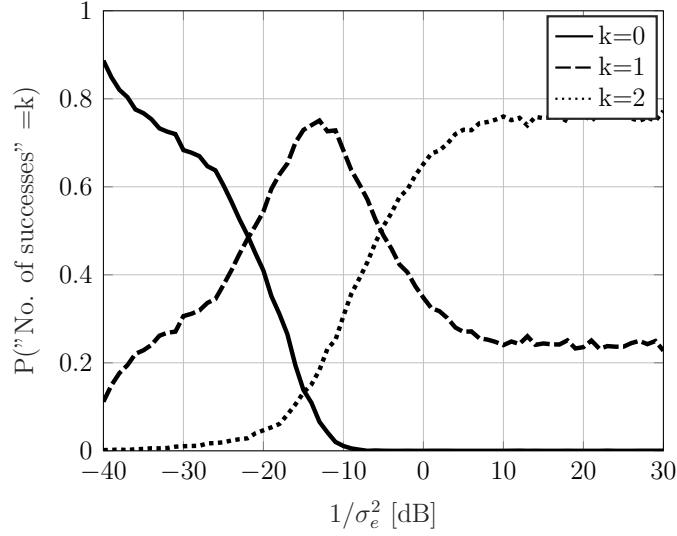


Figure 3.15.: Probability distribution of recovered codewords in the case of a two-user collision with MRC according to the model (3.83). Polar code with $B - B_p = 84$ message bits, $n_d = 2048$ complex QPSK coded symbols and SCL decoding with 16 CRC-bits and list-size 32.

3.5. Going Cell-Free

The simulations in Section 3.3.6 have shown that power imbalances in the received LSFCs significantly worsen the performance of the ML AD algorithm when $K_a > L$ compared to constant LSFCs. With the presented U-RA scheme it is possible to circumvent the near-far effect by combining the measurements of several BSs in a very simple way. Each BS uses the ML algorithm to estimate the active columns in each slot. A central processor then forms the union of the active sets of geographically close BSs in each slot and runs the outer tree decoder. Alternatively each BS can run the outer decoder individually and the central processor forms the union of the message lists. While the first approach leads to lower error probabilities, as the simulations will show, the second approach is better suited to a completely decentralized approach. This can be done without any a-priori cell management. Indeed, the single BSs will naturally identify the strongest signals in their vicinity without the need for BS assignment or power control. This allows for the design of a truly scalable random access system, since users can be added to the system without modifying the existing system and the number of supported active users per slot can be scaled arbitrarily by increasing the number of antennas per BS and the density of BSs. There problem of pilot contamination does not arise.

In the last part of this section I present a new concept which allows for joint AD and positioning. The idea is to share the distant dependent LSFCs between BSs, use them to

estimate the locations of the active users and use the location estimates to refine the ML estimation of the LSFCs. This process can be iterated until convergence. Although, the simulation shows a better AD performance compared to completely independent processing the improved AD algorithm cannot readily be used in the CCS U-RA scheme, mainly because the complexity is too high to allow for sufficient simulation numbers. Nonetheless, the concept seems promising and future research may improve the complexity of the algorithm.

3.5.1. Model

We assume that K_{tot} users are distributed at random in some two dimensional area of which K_a are active at each time slot. Q base stations (BS) with M receive antennas per BS are placed at positions $\mathbf{z}^q = (z_x^q, z_y^q)$, $q = 1, \dots, Q$. Let $\mathbf{p}^k = (p_x^k, p_y^k)$ denote the position of user k . Let $d_{kq} = \|\mathbf{p}^k - \mathbf{z}^q\|_2$ be the distance between user k and BS q . The large-scale fading coefficients are assumed to follow the distribution introduced in (3.73). We define the function

$$h(d_{kq}) = -\alpha - 10\beta \log_{10}(d_{kq}), \quad (3.84)$$

such that the distribution of the LSFCs can be expressed as

$$p(g_{kq} = g \text{ [dB]} | \mathbf{p}^k) = \mathcal{N}(h(d_{kq}), \sigma_{\text{shadow}}^2). \quad (3.85)$$

Furthermore, g_{kq} are assumed to be conditionally independent given \mathbf{p}^k

$$p(g_{k1}, \dots, g_{kQ} | \mathbf{p}^k) = (2\pi\sigma_{\text{shadow}}^2)^{-\frac{Q}{2}} \exp\left(-\frac{\sum_{q=1}^Q \|h(d_{kq}) - 10\log_{10}(g_{kq})\|_2^2}{2\sigma_{\text{shadow}}^2}\right) \quad (3.86)$$

Then the measurements at the BSs are created according to

$$\begin{aligned} \mathbf{Y}_1 &= \mathbf{A}\mathbf{\Gamma}_1^{1/2}\mathbf{H}_1 + \mathbf{Z}_1 \\ &\vdots \\ \mathbf{Y}_Q &= \mathbf{A}\mathbf{\Gamma}_Q^{1/2}\mathbf{H}_Q + \mathbf{Z}_Q \end{aligned} \quad (3.87)$$

with $\mathbf{\Gamma}_q = \text{diag}(\gamma_q)$ and $\gamma_q = (\gamma_{1q}, \dots, \gamma_{K_{\text{tot}}q})$ with $\gamma_{kq} = Pb_k g_{kq}$. For the AD scenario $\mathbf{A} \in \mathbb{C}^{L \times K_{\text{tot}}}$ is the matrix of pilot sequences. In the U-RA scenario \mathbf{A} is the inner coding matrix with $K_{\text{tot}} = 2^J$, as described in Section 3.3.1. $\mathbf{H}_q \in \mathbb{C}^{K_{\text{tot}} \times M}$ are the small-scale fading coefficients and, as before, are assumed to have iid $\mathcal{CN}(0, 1)$ entries. The entries of the noise matrices \mathbf{Z}_q are iid $\mathcal{CN}(0, 1)$. The propagation model can be summarized by the

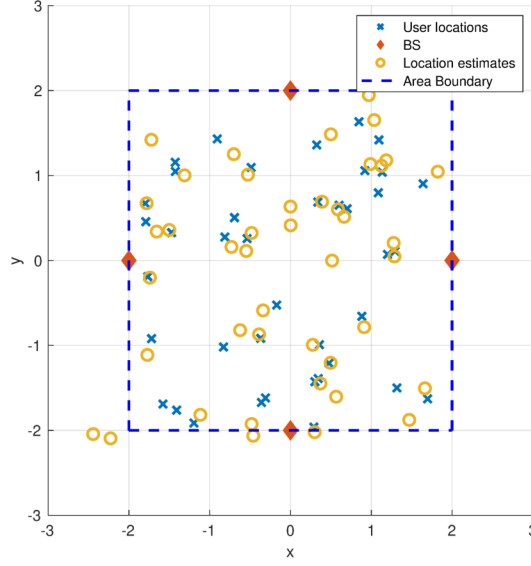


Figure 3.16.: Exemplary snapshot of a square cell geometry with $Q = 4$ BSs and the estimates positioning algorithm (Section 3.5.3) after 4 iterations. $R_{\text{cell}} = 2\text{km}$, $Q = 4$, $K_a = 40$, $M = 50$, $L = 100$, $\sigma_{\text{shadow}}^2 = 1$

Markov-Chain

$$(\mathbf{p}^1, \dots, \mathbf{p}^{K_a}) \rightarrow (\gamma_1, \dots, \gamma_Q) \rightarrow (\mathbf{Y}_1, \dots, \mathbf{Y}_Q). \quad (3.88)$$

3.5.2. U-RA Simulations

For simplicity this work will consider a fixed square area with side-length R_{cell} and $Q = 4$ BSs, one placed at the center of each edge. In Figure 3.17 we compare two different decoding schedules. In both of them each BS first creates, independent of the others, a list of active columns in each slot $\mathcal{S}_1^q, \dots, \mathcal{S}_S^q$ according to (3.69). In the first schedule, referred to as *individual outer decoding*, each BS runs the outer tree decoder on their tuple $(\mathcal{S}_1^q, \dots, \mathcal{S}_S^q)$ to create a list of messages \mathcal{L}^q . A central processor collects these lists and forms the final output list as the union

$$\mathcal{L} = \bigcup_{q=1}^Q \mathcal{L}^q \quad (3.89)$$

of all lists. In the second schedule, called *joint outer decoding*, the lists of active columns are first concatenated by the central processor $\mathcal{S}_s = \bigcup_{q=1}^Q \mathcal{S}_s^q$ and then the outer tree code is run on $\mathcal{S}_1, \dots, \mathcal{S}_S$.

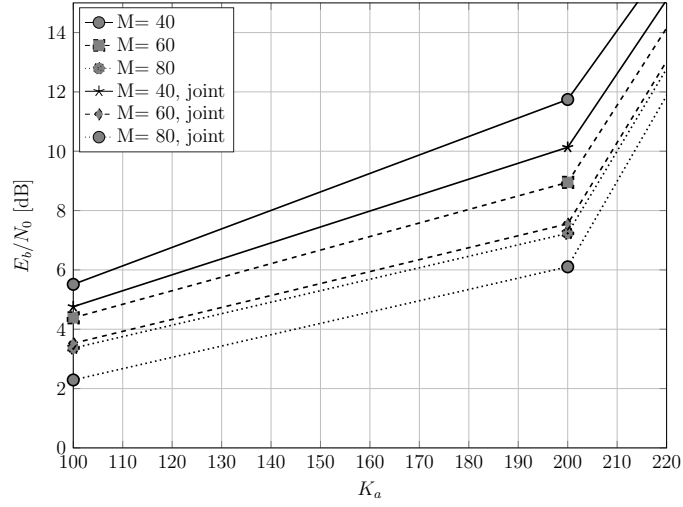


Figure 3.17.: Comparison between independent outer decoding and joint outer decoding in the decentralized setting with $Q = 4$ and a square cell with side length 2 km. $L = 100, n = 3200$, remaining parameters as in Section 3.3.6.

The parameters of the LSFCs are chosen as in Section 3.3.6. We fix $L = 100, n = 3200$ and $S = 32$. The remaining parameters are given in Table 3.1. The results in Figure 3.17 show that the joint outer decoding is generally better than the individual decoding, even though the difference is not large. A more detailed evaluation of the energy efficiency of decentralized CCS with joint outer decoding at different values of M is given in Figure 3.18. We can observe that it is possible to increase the number of active users beyond L without any cell management even though each single BS on its own would not be able to decode all messages in the presence of severe power imbalances as the results in Section 3.3.6 have shown.

3.5.3. Joint AD and Position Estimation

We formulate the algorithm for the AD problem of detecting K_a active out of K_{tot} total users, as in Section 3.2.

MAP LSFC estimation

Following the propagation model (3.73), each BS can find an estimate of $\gamma_q = (\gamma_{q1}, \dots, \gamma_{qK_{\text{tot}}})$ from measurements \mathbf{Y}_q by maximizing the posterior probability $p(\gamma|\mathbf{Y}_q)$ instead of the likelihood $p(\mathbf{Y}_q|\gamma)$. The new cost function is given by

$$f_q^{\text{MAP}}(\gamma) = -\log(p(\mathbf{Y}_q|\gamma)) - \log(p(\gamma|\mathbf{p}_1, \dots, \mathbf{p}_{K_{\text{tot}}})) \quad (3.90)$$

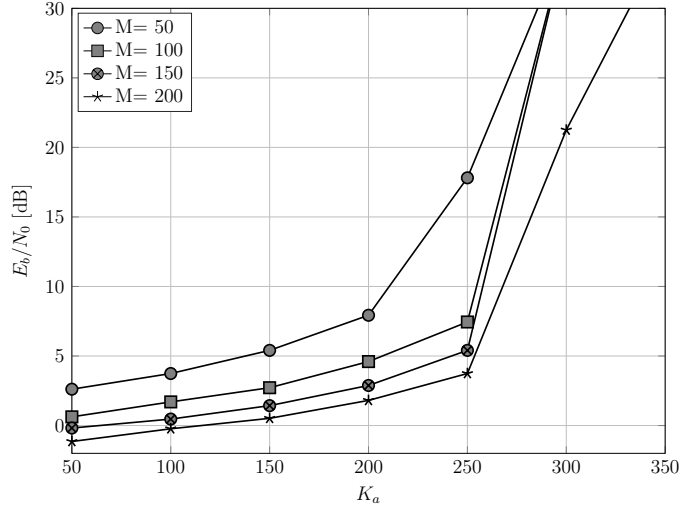


Figure 3.18.: Decentralized U-RA with CCS, ML as inner decoder and joint outer decoding in a square cell with $Q = 4$ and side length 2 km. $L = 100, n = 3200$, remaining parameters as in Section 3.3.6.

As in the ML algorithm, this estimate can be found efficiently by coordinate-wise optimization. When the assumed prior on γ is iid. the coordinate-wise optimization reduces to

$$\frac{d}{dd}[f(\gamma + d\mathbf{e}_k) - \log(p(\gamma_k + d|\mathbf{p}_k)))] = 0. \quad (3.91)$$

For Gaussian marginals the condition (3.91) leads to closed form update equations [127]. In the general case the solution to (3.91) can be found numerically.

Position estimation

Under the log-shadowing propagation model (3.73) the ML estimate of the positions $\mathbf{p}^1, \dots, \mathbf{p}^{K_{\text{tot}}}$ coincides with the least-squares solution:

$$\mathbf{p}^k = \arg \min_{\mathbf{p} \in \mathbb{R}^2} \sum_{q=1}^Q \|h(\|\mathbf{p} - \mathbf{z}^q\|) - 10 \log_{10} \left(\frac{\gamma_{kq}}{P} \right)\|_2^2 \quad (3.92)$$

Note that the model (3.73) holds more generally for $b_k g_{kq}$, instead of just g_{kq} , when the positions of inactive users are defined by $\mathbf{p}^k := (\infty, \infty)$ for $k \notin \mathcal{K}_a$. (3.92) is a non-convex optimization problem, so it cannot be guaranteed to have no local minima. Nonetheless, due to the small dimension of the problem the solution can be found by either grid search or a generic non-linear solver like MATLAB's LSQNONLIN function.

In general the accuracy of position estimation based on received signal strengths is limited due to the strong effect of shadowing, as can be seen from calculating the Cramer-Rao bound [77, 78]. Nonetheless, when a more accurate pathloss map, i.e. a function of the form $h_q(\mathbf{p})$ that gives the pathloss between BS q and a user at position \mathbf{p} , is available, it can simply be substituted in (3.92). Such maps can be generated for a given area by physical measurements, ray-tracing simulations or novel deep-neural-network based methods [79, 80].

Iterative algorithm

The MAP estimation of γ^q and the position estimation can be alternated in an iterative fashion if the BSs are able to share their estimates of γ^q with each other. Let μ_{kq}^t and $(\sigma^t)_k^2$ denote the mean and variance of the estimated log-normal distribution of the received LSFC at BS q at iteration t .

We initialize $(\sigma^0)_k^2 = \infty$ for all $k = 1, \dots, K_{\text{tot}}$, which is equivalent to an uninformative prior. μ_k^0 can be chosen arbitrarily. The iterations proceed as follows. At time step t

1. Each BS calculates an estimate γ_q^t of the LSFCs by coordinate-wise optimization of the log-posterior-probability (3.90) assuming a log-normal prior on γ with means and variances μ_k^t and $(\sigma^t)_k^2$.
2. The estimates $\gamma_1^t, \dots, \gamma_Q^t$ are shared and used to create position estimates \mathbf{p}_k^t of the users by (3.92).
3. The position estimates are used to update the means and variances as

$$\mu_{kq}^{t+1} = \max(h(\|\mathbf{p}_k^t - \mathbf{z}_q\|), \mu^*) \quad (3.93)$$

$$(\sigma^{t+1})_k^2 = \min_{\mathbf{p} \in \mathbb{R}^2} \sum_{q=1}^Q \|h(\|\mathbf{p}_k^t - \mathbf{z}_q\|) - 10 \log_{10} \left(\frac{\gamma_{kq}^t}{P} \right)\|_2^2 \quad (3.94)$$

4. Users with $\max_q(\mu_{kq}^{t+1}) \leq \mu^*$ are marked as inactive, i.e. μ_{kq}^{t+1} is set to zero and $\mathbf{p}_k := (\infty, \infty)$.

The last step is necessary to prune weak users from the system. The maximum in Step 3 is only taken to avoid too small numbers which would lead to numerical problems. μ^* is a parameter that controls the minimum signal strength that is detected as active. The empirical estimation of σ^t in (3.94) as the estimation error of the positions makes the algorithm stable: If the variance in the position estimation is too high the prior is not taken into account in the MAP estimation step.

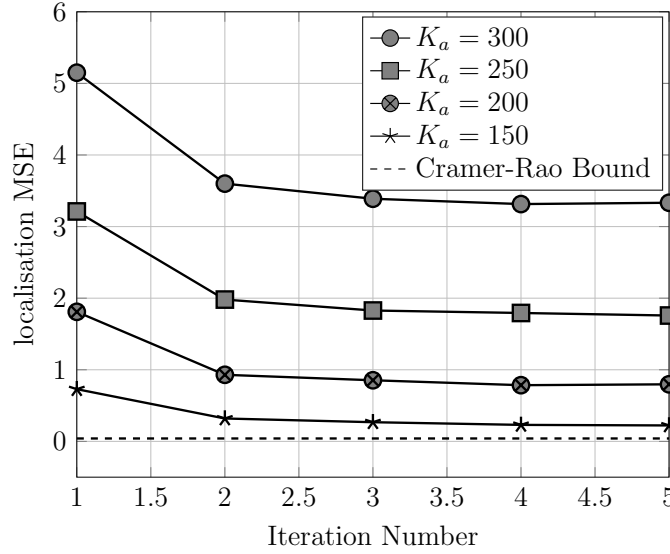


Figure 3.19.: Localisation MSE for square cells with side length 2 km and $Q = 4$ BSs at the center of each side. $L = 100, M = 50, K_{\text{tot}} = 2000, \sigma_{\text{shadow}}^2 = 8$.

For the AD each BS creates an estimated active set \mathcal{A}_q by picking the indices of the $L + \Delta$ largest components of γ_q . The final estimated set of active users is taken as the union of all \mathcal{A}_q . Then the error is defined as in Section 3.2.3.

For the simulations in Figure 3.19, Figure 3.20 and Figure 3.21 we chose the square geometry as in Figure 3.16 with side length 2 km. The parameters of the pathloss model are chosen as in Section 3.3.6. We let $L = 100, M = 50, K_{\text{tot}} = 2000, \Delta = 10, \mu^* = -30$ dB and fix $P = 15$ dBm. Figure 3.19 shows that the localization MSE indeed decreases with each iteration, although the biggest improvement occurs after the first iteration. The Cramer-Rao bound was calculated as described in [78] and gives the best MSE that can be achieved by an unbiased estimator that has access to the LSFCs at the BSs. However, it does not take into account the uncertainty of the estimation of the LSFCs by the ML algorithm. We can see that for active user numbers below $K_a = 150$ the MSE of the position estimation is very close to the Cramer-Rao bound. The AD error is visualized in Figure 3.20 and Figure 3.21. We can see that false alarm are efficiently pruned by the cooperation of the BSs. The missed detection rate does not improve though or gets worse. This is expected because a user who is missed by all the BSs in the first step will very likely be missed in subsequent iterations as well. On the other hand, a user with a weak connection to all BSs can be detected initially by at least one BS but dropped later by the thresholding step. This happens only rarely and to user with a very weak signal.

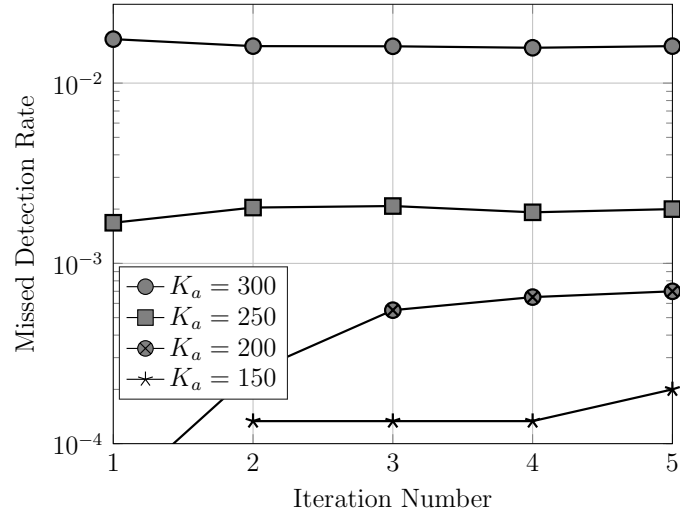


Figure 3.20.: Missed detection rate for square cells with side length 2 km and $Q = 4$ BSs at the center of each side. $L = 100, M = 50, K_{\text{tot}} = 2000, \sigma_{\text{shadow}}^2 = 8$.

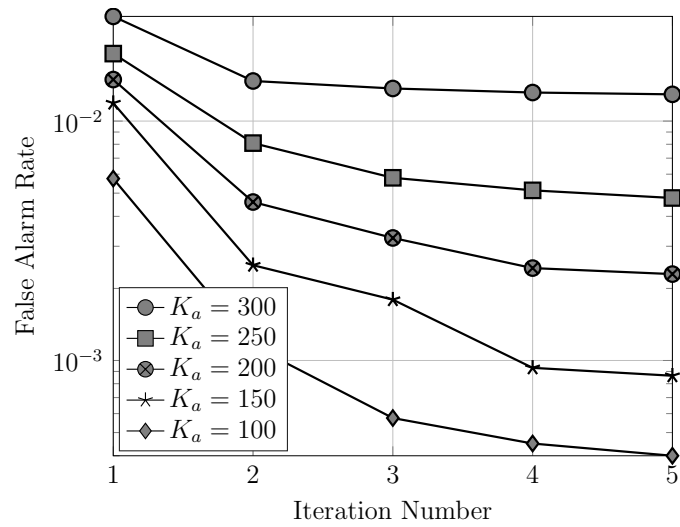


Figure 3.21.: False alarm rate for square cells with side length 2 km and $Q = 4$ BSs at the center of each side. $L = 100, M = 50, K_{\text{tot}} = 2000, \sigma_{\text{shadow}}^2 = 8$.

3.6. Summary

In the beginning of this chapter the problem of user activity detection in a massive MIMO setup was introduced. It was shown that with a covariance-based algorithm and L signal dimensions one can reliably estimate the activity of $K_a = O(L^2/\log^2(K_{\text{tot}}/K_a))$ active users in a set of K_{tot} users, which is a much larger than the previous bound $K_a = O(L)$ obtained via traditional compressed sensing techniques. In particular, one needs to pay only a poly-logarithmic penalty $O(\log^2(K_{\text{tot}}/K_a))$ with respect to the number of potential users K_{tot} , which makes the scheme ideally suited for activity detection in IoT setups where the number of potential users can be very large. We discuss low-complexity algorithms for activity detection and provided numerical simulations to illustrate our results. In particular, as a byproduct of numerical investigation, we also showed a curious unstable behavior of MMV-AMP in the regime where the number of receiver antennas is large, which is precisely the case of interest with a massive MIMO receiver. We proposed a scheme for unsourced random access where we use the introduces activity detection scheme(s) directly. We showed that an arbitrarily fixed probability of error can be achieved at any E_b/N_0 for sufficiently large number of antennas, and a total spectral efficiency that grows as $O(L \log L)$, where L is the code block length, can be achieved. Such a one-shot scheme is conceptually nice but not suited for typical practical applications with message payload of the order of $B \approx 100$ bits, since it would require a codebook matrix with 2^B columns. Hence, we have also considered the application of the concatenated approach pioneered in [87], where the message is broken into a sequence of smaller blocks and the activity detection scheme is applied as an inner encoding/decoding stage at each block, while an outer tree code takes care of “stitching together” the sequence of decoded submessages over the blocks. Numerical simulations show the effectiveness of the proposed method. It should be noticed that these schemes are completely non-coherent, i.e., the receiver never tries to estimate the massive MIMO channel matrix of complex fading coefficients. Therefore, the scheme pays no hidden penalty in terms of pilot symbol overhead, often connected with the assumption of ideal coherent reception, i.e., channel state information known to the receiver.

The complexity of the covariance-based inner decoder grows proportional to $\mathcal{O}(L^2)$ which can become troublesome at very large L . For this case we introduce and analyse an alternative scheme based on transmitting non-orthogonal signature sequences randomly chosen from a given set and subsequent maximum-ratio-combining. The use of a single-user polar code with a list decoder allows to resolve possible pilot collisions. Simulations show that such a setup, despite its conceptual simplicity and low complexity gives results similar to other proposed approaches for this setting. Furthermore, the analytical results

show an excellent agreement with the simulations, which will allow to optimize the scheme in the future. Finally, it is shown that the coupled-compressed-sensing based scheme can be readily extended to a setting with distributed base stations, where only minimal information exchange is required between the receivers. A novel algorithm is presented that allows to improve the activity detection and estimate the positions of the active users by exchanging the estimated large-scale fading coefficients. Simulations show the efficiency of such a setup.

4. Conclusion

The first chapter of this work analysed in-depth the properties of the coupled-compressed-sensing algorithm for unsourced random access on the real AWGN channel without fading, while the second chapter introduced concepts and algorithms for unsourced random access on a block-fading channel with multiple receive antennas. The technical content of these chapter was already recollectd in their respective summaries. Here, I will focus on relations between the two chapters, points that remained open, possible future research directions and a final conclusion.

Although the scaling law for the sum-spectral efficiency, derived in Chapter 3, has important implications for the design of massive random access systems, it is unsatisfying in the sense that it does not provide concrete numbers on achievable rates, i.e. the constants in the scaling laws are unspecified. An extension of the replica/state evolution analysis in Chapter 2 to the fading case with multiple receive antennas seems challenging because the existing analysis methods do not account for the relative scaling of L and M , as discussed in Section 3.2.3. A precise analysis of the error probabilities may allow for the design of a power allocation that improves the performance similar to the $M = 1$ case. The results in Section 2.8 have shown that a specifically matched power allocation is necessary to improve the performance compared to the uniform power allocation, which is very hard to find by empiric search.

An important bottleneck in the usage of the covariance-based algorithm for activity detection with large L is the $\mathcal{O}(L^2)$ scaling of the complexity, caused by the size of the covariance matrix. This scaling does not seem to be fundamental and a smart choice of signature sequences may lead to fast operations which speed up the algorithm, similar to the replacement of matrix multiplications with fast-Fourier transforms in AMP.

Another issue that was ignored in this and comparable works is the one on synchronisation. Usually timing synchronisation has to be acquired in the initial access phase. So for a random access scheme to be truly grant-free would require to avoid the initial timing acquisition, since it is not reasonable in the massive IoT setting to keep all potential users synchronised. Because the CCS scheme can be used together with OFDM modulation, in principle, the synchronisation problem can be dealt with in the same way that timing offsets are detected in the initial access phase of 4G-LTE, i.e. by adding a sufficiently long

4. Conclusion

guard interval and a cyclic prefix to the pilot sequences. A detailed study on the impact of different timing delays on the achievable rates remains open.

Overall, from an information theoretic point of view, the "unsourced" paradigm seems to be the correct model for a grant-free random access setting, because the assumption of individual codebooks assumes implicitly that some sort of initial access protocol has already taken place, which is not compliant with a truly grant-free setting and unfeasible for a very large amount of users with sporadic activity. This work has shown that unsourced random access, and the coupled compressed sensing approach in particular, has a lot of favorable properties and enables an alternative design of grant-free random access systems that can get around several of the problems which pilot based protocols suffer from.

Appendices

A. Optimal product distribution

Lemma 1. Let $\Omega \subset \mathbb{R}$ be some discrete set. Let $p(\mathbf{s})$ with $p : \Omega^{2^J} \rightarrow \mathbb{R}_+$ be a probability mass function on Ω^{2^J} . Let $p_{s_1}(s_1), \dots, p_{s_{2^J}}(s_{2^J})$ denote the marginals of $p(\mathbf{s})$. Further let,

$$\mathcal{P}_{\text{prod}} := \left\{ q(\mathbf{s}) = \prod_{i=1}^{2^J} q_i(s_i), q_i : \Omega \rightarrow \mathbb{R}_+ \left| \sum_{\mathbf{s} \in \Omega} q(\mathbf{s}) = 1 \right. \right\} \quad (\text{A.1})$$

denote the space of product distributions on Ω^{2^J} . Then

$$\arg \min_{q \in \mathcal{P}_{\text{prod}}} D(p \parallel q) = \prod_{i=1}^{2^J} p_{s_i}(s_i) \quad (\text{A.2})$$

Proof. For a product distribution $q \in \mathcal{P}_{\text{prod}}$, $D(p \parallel q)$ can be expressed as:

$$D(p \parallel q) = \sum_{\mathbf{s}} p(\mathbf{s}) \log \frac{p(\mathbf{s})}{q(\mathbf{s})} \quad (\text{A.3})$$

$$= \sum_{\mathbf{s}} p(\mathbf{s}) \log \frac{p(\mathbf{s})}{\prod_i p_{s_i}(s_i)} \frac{\prod_i p_{s_i}(s_i)}{q(\mathbf{s})} \quad (\text{A.4})$$

$$= D\left(p \parallel \prod_i p_{s_i}\right) + \sum_{\mathbf{s}} \sum_{i=1}^{2^J} p(\mathbf{s}) \log \frac{p_{s_i}(s_i)}{q_i(s_i)} \quad (\text{A.5})$$

The first term is independent of q and the second term can be rewritten as

$$\sum_{\mathbf{s}} \sum_{i=1}^{2^J} p(\mathbf{s}) \log \frac{p_{s_i}(s_i)}{q_i(s_i)} \quad (\text{A.6})$$

$$= \sum_{i=1}^{2^J} \sum_{s_i} \left(\sum_{\mathbf{s} \setminus s_i} p(\mathbf{s}) \right) \log \frac{p_{s_i}(s_i)}{q_i(s_i)} \quad (\text{A.7})$$

$$= \sum_{i=1}^{2^J} \sum_{s_i} p_{s_i}(s_i) \log \frac{p_{s_i}(s_i)}{q_i(s_i)} \quad (\text{A.8})$$

$$= \sum_{i=1}^{2^J} D(p_{s_i} \parallel q_i) \quad (\text{A.9})$$

A. Optimal product distribution

which is non-negative and minimized by $q_i \equiv p_{s_i}$.

□

B. $D_{\text{KL}}(\text{Mult} \parallel \text{Bin})$

A vector (z_1, \dots, z_{2^J}) is called multinomial distributed with parameter n and probabilities p_1, \dots, p_{2^J} if

$$p(\mathbf{z}) = \mathbb{P}(Z_1 = z_1, \dots, Z_{2^J} = z_{2^J}) \quad (\text{B.1})$$

$$= \begin{cases} \frac{n!}{z_1! \dots z_{2^J}!} p_1^{z_1} \dots p_{2^J}^{z_{2^J}} & \sum_{i=1}^{2^J} z_i = n \\ 0 & \text{else} \end{cases} \quad (\text{B.2})$$

It follows from the multinomial theorem, that the distribution is normalized $\sum_{\mathbf{z}} p(\mathbf{z}) = 1$. An important property of the multinomial distribution is that the marginals follow a binomial distribution:

$$p(Z_i = z_i) = \sum_{\mathbf{z} \setminus z_i} \mathbb{P}(z_1, \dots, z_{2^J}) = \binom{n}{z_i} p_i^{z_i} (1 - p_i)^{n - z_i} \quad (\text{B.3})$$

with covariance given by $\text{cov}(Z_i, Z_j) = -np_i p_j$.

Let $n = K_a$ and $p_i = 2^{-J}$ for all $i = 1, \dots, K_a$ and let $q(Z)$ denote the binomial distribution with parameters $n = K_a$ and $p_i = 2^{-J}$. Then the marginals of p are all identical and equal to $q(Z)$ and it holds

$$D_{\text{KL}} \left(p \parallel \prod_{i=1}^{K_a} q_i \right) = \sum_{\mathbf{z}} p(\mathbf{z}) \log \frac{p(\mathbf{z})}{q_i(z_i)} \quad (\text{B.4})$$

$$= -H(p) + 2^J H(q) \quad (\text{B.5})$$

where $H(p)$ denotes the entropy of the multinomial distribution p and $H(q)$ the entropy of the binomial distribution q . Both entropies are well known and given by

$$H(p) = JK_a - \log_2 K_a! + 2^J \sum_{t=0}^{K_a} \binom{K_a}{t} (2^{-J})^t (1 - 2^{-J})^{K_a - t} \log_2 t! \quad (\text{B.6})$$

and

$$H(q) = -\log_2 K_a! + \mathbb{E}_q[\log Z!] + \mathbb{E}_q[\log(K_a - Z)!] + J\mathbb{E}_q[Z] - \mathbb{E}_q[K_a - Z] \log_2(1 - 2^{-J}) \quad (\text{B.7})$$

We have $\mathbb{E}_q[Z] = K_a/2^J$ and $\mathbb{E}_q[K_a - Z] = K_a - K_a/2^J$. In the limit for large J , we can expand $\mathbb{E}_q[\log_2(K_a - Z)!]$ in terms of 2^{-J} and get:

$$\mathbb{E}_q[\log_2(K_a - Z)!] = \log_2 K_a! - \frac{K_a}{2^J} \log_2 K_a + \mathcal{O}\left(\frac{1}{2^{2J}}\right). \quad (\text{B.8})$$

Inserting this, (B.6) and (B.7) into (B.5), many terms cancel and we get:

$$D_{\text{KL}}\left(p \parallel \prod_{i=1}^{K_a} q_i\right) = \log_2 K_a! - K_a(2^J - 1) \log_2(1 - 2^{-J}) - K_a \log_2 K_a \quad (\text{B.9})$$

Using $\log_2(1 - 2^{-J}) = -\log_2 e/2^J + \mathcal{O}(1/2^{2J})$ for large J and the Stirling approximation $\log_2 K_a! = K_a \log_2 K_a - K_a \log_2 e + \mathcal{O}(\log K_a)$ we get

$$D_{\text{KL}}\left(p(\mathbf{x}) \parallel \prod_{i=1}^{K_a} q_i(x_i)\right) = \mathcal{O}(\log K_a) - \frac{K_a \log_2 e}{2^J}. \quad (\text{B.10})$$

which implies (2.25).

C. Exponential L^1 convergence implies exponential pointwise convergence a.e.

Theorem 14. *Let $(f_J)_{J=1,2,\dots}$ be a sequence of integrable functions s.t.*

$$\|f_J - f\|_{L^1} \leq \frac{c}{2^J} \quad (\text{C.1})$$

for some constant $c > 0$ and all large enough J . Then for any $\delta > 0$ there is a J_δ such that for all $J \geq J_\delta$

$$|f_J(t) - f(t)| \leq \frac{1}{2^{(1-\delta)J}} \quad (\text{C.2})$$

holds for all t except for a set of size $\mathcal{O}(2^{-\delta J})$. \square

Proof. Let $\epsilon > 0$ then

$$\sum_{j=J}^{\infty} \|f_j - f\|_{L^1} \leq \frac{c}{2^{J-1}} < \epsilon \frac{1}{2^{(1-\delta)J}} \quad (\text{C.3})$$

holds for all

$$J > J_\delta = \log_2(2c/\epsilon)/\delta \quad (\text{C.4})$$

where we have used condition (C.1) and the formula

$$\sum_{j=J}^{\infty} 2^{-j} = 2^{-(J-1)}. \quad (\text{C.5})$$

Now define the sets

$$A_j = \left\{ t : |f_j(t) - f(t)| > \frac{1}{2^{(1-\delta)J}} \right\} \quad (\text{C.6})$$

and let $\mu(A_j)$ denote the Lebesgue measure of A_j . Then it follows from elementary properties of the integral that

$$\sum_{j=J}^{\infty} \mu(A_j) \frac{1}{2^{(1-\delta)J}} < \sum_{j=J}^{\infty} \|f_j - f\|_{L^1} < \epsilon \frac{1}{2^{(1-\delta)J}} \quad (\text{C.7})$$

and so

$$\mu \left(\bigcup_{j=J}^{\infty} A_j \right) \leq \sum_{j=J}^{\infty} \mu(A_j) < \epsilon \quad (\text{C.8})$$

Let $\mathcal{A}_J := \bigcup_{j=J}^{\infty} A_j$, then

$$\mathcal{A}_J^c = \left\{ t : |f_j(t) - f(t)| \leq \frac{1}{2^{(1-\delta)J}}, \forall j \geq J \right\} \quad (\text{C.9})$$

and so (C.8) states that $\mu(\mathcal{A}_J)$, the measure of the set of points on which the pointwise convergence does not hold, can be made arbitrary small. More precisely, it follows from (C.4) that

$$\mu(\mathcal{A}_J) = \mathcal{O} \left(2^{-\delta J} \right) \quad (\text{C.10})$$

□

D. Proof of Theorem 5

Let

$$q(k) = p_k \tag{D.1}$$

for $k = 0, \dots, K_a$, where p_k are the binomial probabilities defined in (2.7) and

$$q_{\text{OR}}(0) = p_0 \tag{D.2}$$

$$q_{\text{OR}}(1) = 1 - p_0. \tag{D.3}$$

Let r, s, z be jointly distributed according to the Gaussian model

$$r = \sqrt{t}s + z, \tag{D.4}$$

with $z \sim \mathcal{N}(0, 1)$ independent of s for some fixed $t \geq 0$ and s distributed according to q . Let $\text{mmse}(t)$ be the MMSE of estimating s from the Gaussian observation r and let $f(r)$ be the PME of s , given by

$$f(r) = \frac{1}{Z(r)} \sum_{k=0}^{K_a} p_k k e^{-(r-k\sqrt{t})^2/2} \tag{D.5}$$

with

$$Z(r) = \sum_{k=0}^{K_a} p_k e^{-(r-k\sqrt{t})^2/2} \tag{D.6}$$

Let $f^{\text{OR}}(r)$ be the mismatched PME, which estimates s from r assuming that s is distributed according to q_{OR} . It is given by

$$f^{\text{OR}}(r) = \frac{(1 - p_0)e^{-(r-\sqrt{t})^2/2}}{Z^{\text{OR}}(r)} \tag{D.7}$$

with

$$Z^{\text{OR}}(r) = p_0 e^{-r^2/2} + (1 - p_0) e^{-(r-\sqrt{t})^2/2} \tag{D.8}$$

Let $\text{mmse}_{\text{OR}}(t)$ be the mean square error of $f^{\text{OR}}(r)$. Since $f^{\text{OR}}(r)$ is mismatched we have $\text{mmse}(t) \leq \text{mmse}_{\text{OR}}(t)$ and it holds:

$$0 \leq \text{mmse}_{\text{OR}}(t) - \text{mmse}(t) \quad (\text{D.9})$$

$$= \sum_{k=0}^{K_a} p_k \mathbb{E} \left\{ \left[k - f^{\text{OR}}(\sqrt{t}k + z) \right]^2 - \left[k - f(\sqrt{t}k + z) \right]^2 \right\} \quad (\text{D.10})$$

$$= p_0 \mathbb{E} \left\{ \left[f^{\text{OR}}(z) \right]^2 - \left[f(z) \right]^2 \right\} \quad (\text{D.11})$$

$$+ p_1 \mathbb{E} \left\{ \left[1 - f^{\text{OR}}(\sqrt{t} + z) \right]^2 - \left[1 - f(\sqrt{t} + z) \right]^2 \right\} \quad (\text{D.12})$$

$$+ \sum_{k=2}^{K_a} p_k \mathbb{E} \left\{ \left[k - f^{\text{OR}}(\sqrt{t}k + z) \right]^2 - \left[k - f(\sqrt{t}k + z) \right]^2 \right\} \quad (\text{D.13})$$

We can bound the terms (D.11) - (D.13) individually. Since $0 \leq f^{\text{OR}}(r) \leq 1$, (D.13) is bound by

$$(\text{D.13}) \leq \sum_{k=2}^{K_a} p_k k^2 \quad (\text{D.14})$$

$$= \text{Var}(s) + [\mathbb{E}(s)]^2 - p_1 \quad (\text{D.15})$$

$$= \frac{K_a}{2^J} \left(1 - \frac{K_a}{2^J} \right) + \frac{K_a^2}{2^{2J}} - \frac{K_a}{2^J} (1 - 2^{-J})^{K_a-1} \quad (\text{D.16})$$

$$= \mathcal{O} \left(\frac{K_a^2}{2^{2J}} \right) \quad (\text{D.17})$$

For the remaining terms we split the expected values over z in (D.11) and (D.12) in two parts, depending on whether $Z(r) = Z(\sqrt{t} + z)$ is smaller or larger than $Z^{\text{OR}}(r)$.

D.1. $Z(r) \leq Z^{\text{OR}}(r)$

We have:

$$f^{\text{OR}}(r) = \frac{1}{Z^{\text{OR}}(r)} (1 - p_0) e^{-(r-\sqrt{t})^2/2} \quad (\text{D.18})$$

$$\leq \frac{1}{Z(r)} (1 - p_0) e^{-(r-\sqrt{t})^2/2} \quad (\text{D.19})$$

and

$$f(r) = \frac{1}{Z(r)} \sum_{k=0}^{K_a} p_k k e^{-(r-\sqrt{t}k)^2/2} \quad (\text{D.20})$$

$$\geq \frac{1}{Z(r)} p_1 e^{-(r-\sqrt{t})^2/2} \quad (\text{D.21})$$

because all the summands are non-negative. It follows for all r :

$$[f^{\text{OR}}(r)]^2 - [f(r)]^2 \quad (\text{D.22})$$

$$\leq \left[\frac{p_1 e^{-(r-\sqrt{t})^2/2}}{Z(r)} \right]^2 \left[\left(\frac{1-p_0}{p_1} \right)^2 - 1 \right] \quad (\text{D.23})$$

$$\leq \left[\left(\frac{1-p_0}{p_1} \right)^2 - 1 \right] \quad (\text{D.24})$$

$$= \left[\left(\frac{1 - (1-2^{-J})^{K_a}}{\frac{K_a}{2^J} (1-2^{-J})^{K_a-1}} \right)^2 - 1 \right] \quad (\text{D.25})$$

$$\leq \mathcal{O} \left(\frac{K_a^2}{2^{2J}} \right) \quad (\text{D.26})$$

which bounds the integral in (D.11) on the set $\{Z(z) \leq Z^{\text{OR}}(z)\}$. For (D.12) notice that

$$[1 - f^{\text{OR}}(r)]^2 - [1 - f(r)]^2 \quad (\text{D.27})$$

$$= f^{\text{OR}}(r)^2 - f(r)^2 + 2[f(r) - f^{\text{OR}}(r)] \quad (\text{D.28})$$

The first term was already bound in (D.26), we bound the second term by

$$f(r) - f^{\text{OR}}(r) \quad (\text{D.29})$$

$$= [1 - f^{\text{OR}}(r)] - [1 - f(r)] \quad (\text{D.30})$$

$$= \frac{p_0 e^{-r^2/2}}{Z^{\text{OR}}(r)} - \frac{p_0 e^{-r^2/2} + \sum_{k=1}^{K_a} p_k (1-k) e^{-(r-\sqrt{tk})^2/2}}{Z(r)} \quad (\text{D.31})$$

$$\leq \frac{1}{Z(r)} \sum_{k=2}^{K_a} p_k (k-1) e^{-(r-\sqrt{tk})^2/2} \quad (\text{D.32})$$

$$\leq \frac{1}{Z(r)} \max_{k \geq 2} \left\{ e^{-(r-\sqrt{tk})^2/2} \right\} \sum_{k=2}^{K_a} k p_k \quad (\text{D.33})$$

$$= \frac{1}{Z(r)} \max_{k \geq 2} \left\{ e^{-(r-\sqrt{tk})^2/2} \right\} \left(\frac{K_a}{2^J} - p_1 \right) \quad (\text{D.34})$$

$$= \frac{e^{-(r-\sqrt{tk^*})^2/2}}{p_1 e^{-(r-\sqrt{t})^2/2}} \mathcal{O} \left(\frac{K_a^2}{2^{2J}} \right) \quad (\text{D.35})$$

where $k^* = \arg \max_{k \geq 2} \{\exp(-(r - \sqrt{t}k)^2/2)\}$. For the last line, notice that $Z(r) \geq p_k \exp(-(r - \sqrt{t}k)^2/2)$ for all k and especially for $k = 1$. In (D.35) we have $r = \sqrt{t} + z$. For the expected value in (D.12) we have to integrate over z , restricted to the values of z for which $Z(r) \leq Z^{\text{OR}}(r)$. Since

$$\int_{-\infty}^{\infty} e^{-z^2/2} \frac{e^{-(z - \sqrt{t}(k^* - 1))^2/2}}{p_1 e^{-z^2/2}} = \frac{1}{p_1} \quad (\text{D.36})$$

and the integrand is non-negative, the same integral, restricted to $\{z : Z(\sqrt{t} + z) \leq Z^{\text{OR}}(\sqrt{t} + z)\}$, is also bounded by $1/p_1$.

D.2. $Z(r) > Z^{\text{OR}}(r)$

Let r be such that, $Z(r) > Z^{\text{OR}}(r)$. It holds that

$$f^{\text{OR}}(r) = 1 - \frac{p_0 e^{-r^2/2}}{Z^{\text{OR}}(r)} \quad (\text{D.37})$$

$$\leq 1 - \frac{p_0 e^{-r^2/2}}{Z(r)} \quad (\text{D.38})$$

$$= \frac{1}{Z(r)} \sum_{k=1}^{K_a} p_k e^{-(r - k\sqrt{t})^2/2} \quad (\text{D.39})$$

$$\leq \frac{1}{Z(r)} \sum_{k=1}^{K_a} k p_k e^{-(r - k\sqrt{t})^2/2} \quad (\text{D.40})$$

$$= f(r) \quad (\text{D.41})$$

Since both terms are non-negative we get

$$f^{\text{OR}}(r)^2 - f(r)^2 \leq 0 \quad (\text{D.42})$$

which, together with (D.26), shows that (D.11) is bounded by a term of order $\mathcal{O}\left(\frac{K_a^2}{2^{2J}}\right)$.

For (D.12), the same argumentation as in (D.28) holds and it remains to bound $f(r) - f^{\text{OR}}(r)$. We have that

$$f(r) - f^{\text{OR}}(r) \quad (\text{D.43})$$

$$= \frac{1}{Z(r)} \sum_{k=1}^{K_a} k p_k e^{-(r - k\sqrt{t})^2/2} - \frac{1}{Z^{\text{OR}}(r)} (1 - p_0) e^{-(r - \sqrt{t})^2/2} \quad (\text{D.44})$$

$$\leq \frac{1}{Z(r)} \sum_{k=1}^{K_a} k p_k e^{-(r-k\sqrt{t})^2/2} - \frac{1}{Z(r)} (1-p_0) e^{-(r-\sqrt{t})^2/2} \quad (\text{D.45})$$

$$\leq \frac{1}{Z(r)} \sum_{k=2}^{K_a} k p_k e^{-(r-k\sqrt{t})^2/2} \quad (\text{D.46})$$

$$\leq \frac{1}{Z(r)} \max_{k \geq 2} \left\{ e^{-(r-\sqrt{t}k)^2/2} \right\} \left(\frac{K_a}{2^J} - p_1 \right) \quad (\text{D.47})$$

$$= \frac{e^{-(r-\sqrt{t}k^*)^2/2}}{p_1 e^{-(r-\sqrt{t})^2/2}} \mathcal{O} \left(\frac{K_a^2}{2^{2J}} \right) \quad (\text{D.48})$$

This is the same term as in (D.48), for which we have shown that its expected value over z is bounded by $p_1^{-1} \mathcal{O}(K_a^2/2^{2J})$. (D.48), together with (D.35), shows that $f(r) - f^{\text{OR}}(r)$ is bounded by $p_1^{-1} \mathcal{O}(K_a^2/2^{2J})$ for all r . (D.35) and (D.42) show that $f^{\text{OR}}(r)^2 - f(r)^2 = \mathcal{O}(K_a^2/2^{2J})$ for all r , and therefore, by (D.28), also (D.12) is bound by a term of order $\mathcal{O} \left(\frac{K_a^2}{2^{2J}} \right)$. This concludes the proof of Theorem 5.

E. Proof of Theorem 6

The RS-potential (2.32), rescaled by $\beta/2^J$ takes the form

$$i^{\text{RS}}(\eta) = \frac{R_{\text{in}} 2^J}{J} I(\eta \hat{P}) + \frac{\log_2 e}{2} [(\eta - 1) - \ln \eta] \quad (\text{E.1})$$

with the mutual information

$$I(\eta \hat{P}) := I(X; Y) = H(Y) - H(Y|X) \quad (\text{E.2})$$

for $P(X = 0) = p_0$, $P(X = 1) = 1 - p_0$ and $Y = (\eta \hat{P})^{\frac{1}{2}} X + Z$, for $Z \sim \mathcal{N}(0, 1)$ independent of X . The mutual information $I(\eta \hat{P})$ can be evaluated as follows. First, note that in an additive channel $H(Y|X) = H(Z)$, so $H(Y|X)$ is independent of η and therefore we can ignore it. The distribution of Y is given by

$$\begin{aligned} p(y) &= p_0 p(y|x=0) + (1 - p_0) p(y|x=1) \\ &= \frac{p_0}{\sqrt{2\pi}} \exp\left(-\frac{y^2}{2}\right) + \frac{1 - p_0}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(y - (\eta \hat{P})^{\frac{1}{2}}\right)^2\right), \end{aligned} \quad (\text{E.3})$$

so the differential output entropy $H(Y) = -\int p(y) \log_2 p(y) dy$ can be split into the sum of two parts. Define H_0 and H_1 respectively by

$$H_0 := -\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{y^2}{2}\right) \log_2(p(y)) dy \quad (\text{E.4})$$

and

$$\begin{aligned} H_1 &:= -\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2} \left(y - (\eta \hat{P})^{\frac{1}{2}}\right)^2\right) \log_2(p(y)) dy \\ &= -\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{y^2}{2}\right) \log_2\left(p\left(y + (\eta \hat{P})^{\frac{1}{2}}\right)\right) dy \end{aligned} \quad (\text{E.5})$$

such that the following relation holds:

$$I(\eta \hat{P}) = p_0 H_0 + (1 - p_0) H_1. \quad (\text{E.6})$$

Taking into account the scaling factor in (E.1) and using that $\lim_{J \rightarrow \infty} 2^J(1 - p_0) = K_a$ and $\lim_{J \rightarrow \infty} p_0 = 1$ we get that

$$\lim_{J \rightarrow \infty} \frac{R_{\text{in}} 2^J}{J} I(\eta \hat{P}) = \lim_{J \rightarrow \infty} \left(\frac{R_{\text{in}} 2^J}{J} H_0 + \frac{S}{J} H_1 \right) \quad (\text{E.7})$$

Now let us take a closer look at $\log_2 p(y) = \log_2(e) \ln p(y)$ which appears in both H_0 and H_1 . Let $x_1, x_2 > 0$ with $x_2 > x_1$. Then for the logarithm of the sum of exponentials it holds that

$$-\ln(e^{-x_1} + e^{-x_2}) = x_1 + \ln(1 + e^{-(x_2 - x_1)}). \quad (\text{E.8})$$

The error term $\ln(1 + e^{-(x_2 - x_1)})$ decays exponentially as the difference $x_2 - x_1$ grows. Since $p(y)$ is the sum of two exponentials we can approximate $\ln p(y)$ by:

$$-\ln p(y) = \min \left\{ \frac{y^2}{2} - \ln(p_0), \frac{1}{2} \left(y - (\eta \hat{P})^{\frac{1}{2}} \right)^2 - \ln(1 - p_0) \right\} \quad (\text{E.9})$$

This approximation is justified, since the difference of the two exponents in $p(y)$ is proportional to \sqrt{J} , and so it grows large with J .¹ First, note, that since $\min\{a, b\} \leq a$ and $\min\{a, b\} \leq b$ holds for all $a, b \in \mathbb{R}$, $-\ln p(y) \leq y^2/2 - \ln(1 - p_0)$ as well as $-\ln p(y + (\eta \hat{P})^{\frac{1}{2}}) \leq y^2/2 + \ln(2^J/K_a)$. This means that each of the integrands in H_0 and H_1/J resp. is bounded uniformly, for all J , by an integrable function. This allows us to evaluate the integrals by using Lebesgue's theorem on dominated convergence. For this purpose we need to calculate the pointwise limits of $\ln p(y)$ and $\ln p(y + (\eta \hat{P})^{\frac{1}{2}})/J$. The theorem on dominated convergence then states, that the limit of the integrals is given by the integral of the pointwise limits.

The minimum in (E.9) can be expressed as

$$-\ln p(y) = \begin{cases} \frac{y^2}{2} & y < \gamma \\ \frac{1}{2} \left(y - (\eta \hat{P})^{\frac{1}{2}} \right)^2 + \ln \left(\frac{2^J}{K_a} \right) & y \geq \gamma \end{cases} \quad (\text{E.10})$$

where we neglected $\ln(p_0) = \ln(1 - K_a/2^J) \sim K_a/2^J$ and γ is given by

$$\gamma = \frac{1}{2} \left(\eta \hat{P} \right)^{\frac{1}{2}} + \ln \left(\frac{2^J}{K_a} \right) \left(\eta \hat{P} \right)^{-\frac{1}{2}}. \quad (\text{E.11})$$

¹Technically, this approximation does not hold at the point where the two exponents in $p(y)$ are equal. However, since the integral of a function does not depend on the value of the function at points of measure zero, we can redefine $\ln p(y)$ arbitrary at that point.

Given the considered scaling constraints and $\hat{P} = JSNR/R_{\text{in}} = 2J\mathcal{E}_{\text{in}}$, γ can be rewritten as

$$\gamma = \sqrt{\frac{J}{2}} \left(\sqrt{\eta\mathcal{E}_{\text{in}}} + \frac{1 - \frac{1}{\alpha}}{\log e \sqrt{\eta\mathcal{E}_{\text{in}}}} \right) \quad (\text{E.12})$$

The term in parenthesis is strictly positive for all η so $\lim_{J \rightarrow \infty} \gamma = \infty$ and therefore the pointwise limit of $\ln p(y)$ is give by $\lim_{J \rightarrow \infty} \ln p(y) = -y^2/2$. It follows from Lebesgue's theorem on dominated convergence that

$$\lim_{J \rightarrow \infty} H_0 = \log_2 e \quad (\text{E.13})$$

which is independent of η , so we can ignore it when evaluating $i^{\text{RS}}(\eta)$. For the calculation of H_1 note that:

$$-\ln p\left(y + (\eta\hat{P})^{\frac{1}{2}}\right) = \begin{cases} \frac{1}{2} \left(y + (\eta\hat{P})^{\frac{1}{2}}\right)^2 & y < \gamma' \\ \frac{y^2}{2} + \ln\left(\frac{2^J}{K_a}\right) & y \geq \gamma' \end{cases} \quad (\text{E.14})$$

where we defined $\gamma' := \gamma - (\eta\hat{P})^{\frac{1}{2}}$. γ' is not non-negative anymore and therefore the asymptotic behavior of γ' depends on η in the following way:

$$\lim_{J \rightarrow \infty} \gamma' = \begin{cases} \infty & \text{if } \eta < \bar{\eta} \\ 0 & \text{if } \eta = \bar{\eta} \\ -\infty & \text{if } \eta > \bar{\eta} \end{cases} \quad (\text{E.15})$$

where $\bar{\eta}$ was defined in (2.37). This gives the following asymptotic behavior:

$$-\lim_{J \rightarrow \infty} \frac{\ln p(y + (\eta\hat{P})^{\frac{1}{2}})}{J} = \begin{cases} \eta\mathcal{E}_{\text{in}} + \frac{1}{2J} & \eta < \bar{\eta} \\ (1 - \alpha^{-1})/\log_2 e & \eta \geq \bar{\eta} \end{cases} \quad (\text{E.16})$$

Finally, using (E.13), (E.5), (E.7), (E.16) and the θ function defined in (2.36) we get:

$$\lim_{J \rightarrow \infty} \left(\frac{i^{\text{RS}}(\eta)}{\log_2 e} - \frac{R_{\text{in}} 2^J}{J} \right) = \eta S \mathcal{E}_{\text{in}} [1 - \theta(\eta - \hat{\eta})] + \frac{S}{\log_2 e} \left(1 - \frac{1}{\alpha} \right) \theta(\eta - \bar{\eta}) + \frac{1}{2} [(\eta - 1) - \ln \eta] \quad (\text{E.17})$$

This proves the statement of the theorem.

F. Proof of Theorem 11

The main line of arguments in this section is based on [66]. In turns, the proof in [66] is based on a RIP result which was claimed and successively retracted [67]. The result was applied to a non-centered matrix and therefore could not have the claimed property. We fix this here, using our own new RIP result (Theorem 12) and, for the sake of clarity and self-contained presentation, give a complete streamlined proof for the case of known LSFCs. At several points our proof technique differs from [66], which results in the slightly better bound on M . Let us first introduce some notation.

Definition 1. For $t > 1$ define the Renyi divergence of order t between two probability densities p and q as

$$\mathcal{D}_t(p, q) := \frac{1}{t-1} \ln \int p(x)^t q(x)^{1-t} dx \quad (\text{F.1})$$

◇

Definition 2. A differentiable function f is called strongly convex with parameter $m > 0$ if the following inequality holds for all points x, y in its domain:

$$f(y) \geq f(x) + \nabla f(x)^\top (y - x) + \frac{m}{2} \|x - y\|_2^2 \quad (\text{F.2})$$

◇

Without loss of generality we assume that $P = 1$, which can always be achieved by renormalizing $N_0 := N_0/P$. Let \mathbf{b}° denote the true activity pattern with known sparsity K_a , and \mathbf{b}^* be the output of the estimator (3.11). Using the union bound, we can write

$$\begin{aligned} \mathbb{P}(\mathbf{b}^* \neq \mathbf{b}^\circ) &= \mathbb{P} \left(\max_{\mathbf{b} \in \Theta_{K_a} \setminus \{\mathbf{b}^\circ\}} p(\mathbf{Y}|\mathbf{b}) \geq p(\mathbf{Y}|\mathbf{b}^\circ) \right) \\ &= \mathbb{P} \left(\bigcup_{\mathbf{b} \in \Theta_{K_a} \setminus \{\mathbf{b}^\circ\}} \{p(\mathbf{Y}|\mathbf{b}) \geq p(\mathbf{Y}|\mathbf{b}^\circ)\} \right) \\ &\leq \sum_{\mathbf{b} \in \Theta_{K_a} \setminus \{\mathbf{b}^\circ\}} \mathbb{P}(\mathbf{Y} : p(\mathbf{Y}|\mathbf{b}) - p(\mathbf{Y}|\mathbf{b}^\circ) \geq 0) \\ &\leq \sum_{\mathbf{b} \in \Theta_{K_a} \setminus \{\mathbf{b}^\circ\}} \mathbb{P}(\mathbf{Y} : p(\mathbf{Y}|\mathbf{b}) - p(\mathbf{Y}|\mathbf{b}^\circ) > -\alpha) \end{aligned} \quad (\text{F.3})$$

for any $\alpha > 0$. With slight abuse of notation we define $\mathbf{\Sigma}(\mathbf{b}) := \mathbf{A}\mathbf{B}\mathbf{G}^\circ\mathbf{A}^H + N_0\mathbf{I}_L$, the covariance matrix for a given binary pattern \mathbf{b} for a fixed vector of LSFCs \mathbf{g}° , with $\mathbf{B} = \text{diag}(\mathbf{b})$ and $\mathbf{G}^\circ = \text{diag}(\mathbf{g}^\circ)$. Let $p_{\mathbf{b}} := \mathcal{CN}(0, \mathbf{\Sigma}(\mathbf{b}))$ denote the Gaussian distribution with covariance matrix $\mathbf{\Sigma}(\mathbf{b})$, then $\log p(\mathbf{Y}|\mathbf{b}) = \sum_j \log p_{\mathbf{b}}(\mathbf{Y}_{:,j})$. The following large deviation property of $\log p(\mathbf{Y}|\mathbf{b})$ is established in [66, Corollary 1]:

Theorem 15.

$$\begin{aligned} \mathbb{P} \left(\log p(\mathbf{Y}|\mathbf{b}) - \log p(\mathbf{Y}|\mathbf{b}^\circ) > -\frac{M}{2} \mathcal{D}_{1/2}(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) \right) \\ \leq \exp \left(-\frac{M}{4} \mathcal{D}_{1/2}(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) \right) \end{aligned} \quad (\text{F.4})$$

where $\mathcal{D}_{1/2}(p_{\mathbf{b}}, p_{\mathbf{b}^\circ})$ is the Renyi divergence of order $1/2$ between $p_{\mathbf{b}}$ and $p_{\mathbf{b}^\circ}$ defined in Definition 1. \square

The result of Theorem 15 holds only if $\mathcal{D}_{1/2}(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) > 0$, so in the following we will establish conditions under which this is true. First, note that since $p_{\mathbf{b}}$ and $p_{\mathbf{b}^\circ}$ are zero-mean Gaussian distributions with covariance matrices $\mathbf{\Sigma}(\mathbf{b})$ and $\mathbf{\Sigma}(\mathbf{b}^\circ)$ resp., their Renyi divergence of order t can be expressed in closed form as:

$$\mathcal{D}_t(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) = \frac{1}{2(1-t)} \log \frac{|(1-t)\mathbf{\Sigma}(\mathbf{b}) + t\mathbf{\Sigma}(\mathbf{b}^\circ)|}{|\mathbf{\Sigma}(\mathbf{b})|^{1-t} |\mathbf{\Sigma}(\mathbf{b}^\circ)|^t} \quad (\text{F.5})$$

Let $\psi(\mathbf{b}) := -\log |\mathbf{\Sigma}(\mathbf{b})|$, then we can see that $\mathcal{D}_t(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) \geq t \frac{m^*}{4} \|\mathbf{b} - \mathbf{b}^\circ\|_2^2$, with m^* being the strong convexity constant of $\psi(\cdot)$, is equivalent to

$$\psi((1-t)\mathbf{b} + t\mathbf{b}^\circ) \leq (1-t)\psi(\mathbf{b}) + t\psi(\mathbf{b}^\circ) - \frac{1}{2}m^*t(1-t)\|\mathbf{b} - \mathbf{b}^\circ\|_2^2. \quad (\text{F.6})$$

Here we used the fact that

$$\begin{aligned} -\psi((1-t)\mathbf{b} + t\mathbf{b}^\circ) &= \log |\mathbf{\Sigma}((1-t)\mathbf{b} + t\mathbf{b}^\circ)| \\ &= \log |\mathbf{A}((1-t)\mathbf{B} + t\mathbf{B}^\circ)\mathbf{G}^\circ\mathbf{A}^H + N_0\mathbf{I}_L| \\ &= \log |(1-t)\mathbf{\Sigma}(\mathbf{b}) + t\mathbf{\Sigma}(\mathbf{b}^\circ)| \end{aligned} \quad (\text{F.7})$$

Inequality (F.6) is precisely the condition that $\psi(\cdot)$ is strongly convex along the line connecting \mathbf{b} and \mathbf{b}° . So if $\psi(\cdot)$ is strongly convex on the set of $2K_a$ -sparse vectors, then

$$\mathcal{D}_t(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) \geq t \frac{m^*}{4} \|\mathbf{b} - \mathbf{b}^\circ\|_2^2 \quad (\text{F.8})$$

holds for any K_a -sparse vectors \mathbf{b} and \mathbf{b}° . Let $\mathbf{b}_1, \mathbf{b}_2 \in \Theta_{K_a}$ be two arbitrary K_a -sparse vectors. Since $\log|\cdot|$ is differentiable on \mathbb{R}^+ , a Taylor expansion of $\psi(\mathbf{b}_1)$ around \mathbf{b}_2 gives:

$$\begin{aligned}\psi(\mathbf{b}_1) &= \psi(\mathbf{b}_2) + \langle \nabla \psi(\mathbf{b}_1), \mathbf{b}_2 - \mathbf{b}_1 \rangle \\ &\quad + \frac{1}{2}(\mathbf{b}_2 - \mathbf{b}_1)^\top \nabla^2 \psi(\mathbf{b}_r)(\mathbf{b}_2 - \mathbf{b}_1)\end{aligned}\tag{F.9}$$

for $\mathbf{b}_r = (1-r)\mathbf{b}_1 + r\mathbf{b}_2$ with some $r \in [0, 1]$. Let $\Delta \mathbf{b} := \mathbf{b}_2 - \mathbf{b}_1$, then the strong convexity of $\psi(\cdot)$ is equivalent to

$$\sum_{i,j} \frac{\partial^2 \psi}{\partial b_i \partial b_j} \Big|_{\mathbf{b}=\mathbf{b}_r} \Delta b_i \Delta b_j \geq m^* \|\mathbf{b}_2 - \mathbf{b}_1\|_2^2.\tag{F.10}$$

The derivatives of ψ are given by:

$$\frac{\partial \psi}{\partial b_i} \Big|_{\mathbf{b}=\mathbf{b}_r} = -\text{trace}(\Sigma(\mathbf{b}_r)^{-1} g_i^\circ \mathbf{a}_i \mathbf{a}_i^H)\tag{F.11}$$

$$\frac{\partial^2 \psi}{\partial b_i \partial b_j} \Big|_{\mathbf{b}=\mathbf{b}_r} = \text{trace}(\Sigma(\mathbf{b}_r)^{-1} g_i^\circ \mathbf{a}_i \mathbf{a}_i^H \Sigma(\mathbf{b}_r)^{-1} g_j^\circ \mathbf{a}_j \mathbf{a}_j^H)\tag{F.12}$$

Next we will calculate m^* . It holds that

$$\begin{aligned}\sum_{i,j} \frac{\partial^2 \psi}{\partial b_i \partial b_j} \Big|_{\mathbf{b}=\mathbf{b}_r} \Delta b_i \Delta b_j &= \sum_{i,j} \text{trace} \left(\Sigma(\mathbf{b}_r)^{-1} \Delta b_i g_i^\circ \mathbf{a}_i \mathbf{a}_i^H \Sigma(\mathbf{b}_r)^{-1} \Delta b_j g_j^\circ \mathbf{a}_j \mathbf{a}_j^H \right) \\ &= \text{trace} \left(\Sigma(\mathbf{b}_r)^{-1} \left(\sum_i \Delta b_i g_i^\circ \mathbf{a}_i \mathbf{a}_i^H \right) \Sigma(\mathbf{b}_r)^{-1} \left(\sum_j \Delta b_j g_j^\circ \mathbf{a}_j \mathbf{a}_j^H \right) \right) \\ &= \text{trace} \left(\Sigma(\mathbf{b}_r)^{-1} (\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)) \Sigma(\mathbf{b}_r)^{-1} (\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)) \right) \\ &\geq \sigma_{\min}(\Sigma(\mathbf{b}_r)^{-1}) \text{trace} \left((\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)) \Sigma(\mathbf{b}_r)^{-1} (\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)) \right) \\ &\geq \sigma_{\min}^2(\Sigma(\mathbf{b}_r)^{-1}) \|\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)\|_F^2 \\ &= \frac{\|\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)\|_F^2}{\sigma_{\max}^2(\Sigma(\mathbf{b}_r))}.\end{aligned}\tag{F.13}$$

Here $\sigma_{\min}(\mathbf{A})$ (resp., $\sigma_{\max}(\mathbf{A})$) denotes the minimum (resp., maximum) singular value of \mathbf{A} . In the first and the second inequality in (F.13) we used the fact that $\text{trace}(\mathbf{A}\mathbf{B}) \geq \sigma_{\min}(\mathbf{A})\text{trace}(\mathbf{B})$ for positive semi-definite matrices \mathbf{A}, \mathbf{B} , and in the second inequality in (F.13) we used the fact that the covariance matrix is symmetric and $\text{trace}(\mathbf{A}^\top \mathbf{A}) = \|\mathbf{A}\|_F^2$.

We can rewrite $\|\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)\|_F^2 = \|\mathbb{A}(\mathbf{g}^\circ \odot (\mathbf{b}_2 - \mathbf{b}_1))\|_2^2$, where $\mathbb{A} \in \mathbb{C}^{L^2 \times K_{\text{tot}}}$ is the matrix defined in (3.26), obtained by stacking the L^2 -dimensional vectors $\text{vec}(\mathbf{a}_k \mathbf{a}_k^H)$ by

columns. We show in (H.16) that $\|\mathbb{A}\mathbf{x}\|_2 \geq \|\mathring{\mathbb{A}}\mathbf{x}\|_2$ holds $\forall \mathbf{x} \in \mathbb{R}^{K_{\text{tot}}}$, with $\mathring{\mathbb{A}}$ being the centered version of \mathbb{A} , which is defined in (3.28).

We show in Theorem 12, that, with probability at least $1 - \exp(-c_\delta L)$, $\mathring{\mathbb{A}}/\sqrt{L(L-1)}$, the centered and rescaled version of \mathbb{A} has RIP of order $2K_a$ with constant $\delta_{2K_a} < \delta$ if condition (3.12) is fulfilled. In particular $\|\mathring{\mathbb{A}}\mathbf{x}\|_2^2 \geq (1 - \delta_{2K_a})L(L-1)\|\mathbf{x}\|_2^2$ holds for all $2K_a$ -sparse vectors \mathbf{x} . So the RIP of $\mathring{\mathbb{A}}$ implies that

$$\begin{aligned} \|\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)\|_F^2 &\geq (1 - \delta_{2K_a})L(L-1)\|\mathbf{g}^\circ \odot (\mathbf{b}_2 - \mathbf{b}_1)\|_2^2 \\ &\geq (1 - \delta_{2K_a})L(L-1)g_{\min}^2\|\mathbf{b}_2 - \mathbf{b}_1\|_2^2 \\ &\geq \frac{1}{2}(1 - \delta_{2K_a})L^2g_{\min}^2\|\mathbf{b}_2 - \mathbf{b}_1\|_2^2 \end{aligned} \quad (\text{F.14})$$

An upper bound on $\sigma_{\max}^2(\Sigma(\mathbf{b}_r)) = \|\Sigma(\mathbf{b}_r)\|_{op}^2$ can be found as follows. Note that for any binary $2K_a$ -sparse vector \mathbf{b} , it holds that

$$\begin{aligned} \sigma_{\max}(\Sigma(\mathbf{b})) &= \|\Sigma(\mathbf{b})\|_{op} \\ &= \left\| \sum_{k=1}^{K_{\text{tot}}} g_k^\circ b_k \mathbf{a}_k \mathbf{a}_k^H + N_0 \mathbf{I} \right\|_{op} \\ &\leq g_{\max} \left\| \sum_{k \in \text{supp}(\mathbf{b})} \mathbf{a}_k \mathbf{a}_k^H \right\|_{op} + N_0 \\ &= g_{\max} \left\| \sum_{k \in \text{supp}(\mathbf{b})} (\mathbf{a}_k \mathbf{a}_k^H - \mathbf{I}) + 2K_a \mathbf{I} \right\|_{op} + N_0 \\ &\leq g_{\max} \left\| \sum_{k \in \text{supp}(\mathbf{b})} (\mathbf{a}_k \mathbf{a}_k^H - \mathbf{I}) \right\|_{op} + g_{\max} 2K_a + N_0 \end{aligned} \quad (\text{F.15})$$

Now $\sum_{k \in \text{supp}(\mathbf{b})} (\mathbf{a}_k \mathbf{a}_k^H - \mathbf{I})$ is a sum of $2K_a$ random matrices $\mathbf{a}_k \mathbf{a}_k^H$, with \mathbf{a}_k drawn i.i.d. from the sphere of radius \sqrt{L} , and therefore sub-Gaussian. A generic large deviation result for such matrices, e.g., the complex version of [128, Theorem 4.6.1], shows that

$$\left\| \sum_{k \in \text{supp}(\mathbf{b})} (\mathbf{a}_k \mathbf{a}_k^H - \mathbf{I}) \right\|_{op} \leq \left(\sqrt{K_a} + C \left(\sqrt{L} + t \right) \right)^2 \quad (\text{F.16})$$

holds with probability at least $1 - 2\exp(-t^2)$ for some universal constant $C > 0$. Let $t = \sqrt{\beta} \max(\sqrt{K_a}, \sqrt{L})$ for some $\beta > 0$. Then (F.15) gives that

$$\sigma_{\max}(\Sigma(\gamma)) \leq (1 + \beta C') g_{\max} \max\{K_a, L\} + N_0 \quad (\text{F.17})$$

holds with probability at least $1 - \exp(-\beta \max\{K_a, L\})$ for some universal constants $C' > 0$. So (F.13) can be further bounded using (F.14) and (F.17) as

$$\frac{\|\Sigma(\mathbf{b}_2) - \Sigma(\mathbf{b}_1)\|_F^2}{\sigma_{\max}^2(\Sigma(\mathbf{b}_r))} \geq \frac{(1 - \delta_{2K_a})g_{\min}^2 \|\mathbf{b}_2 - \mathbf{b}_1\|_2^2}{2 \left((1 + C'\beta)g_{\max} \max\left\{\frac{K_a}{L}, 1\right\} + \frac{N_0}{L} \right)^2} \quad (\text{F.18})$$

Together with (F.10) and (F.13) this implies that, if the pilot matrix satisfies the RIP of order $2K_a$ with constant $\delta_{2K_a} < 1$, then $\psi(\cdot)$ is strongly convex along the line between any two K_a -sparse vectors with constant

$$m^* \geq \frac{(1 - \delta_{2K_a})g_{\min}^2}{2 \left((1 + C'\beta)g_{\max} \max\left\{\frac{K_a}{L}, 1\right\} + \frac{N_0}{L} \right)^2} \quad (\text{F.19})$$

with probability exceeding $1 - \exp(-\beta \max\{K_a, L\})$. Since the bound is independent of the chosen vectors and the number of $2K_a$ sparse binary vectors is bounded by $\binom{K_{\text{tot}}}{2K_a} \leq (eK_{\text{tot}}/K_a)^{2K_a} \leq (eK_{\text{tot}}/K_a)^{2\max\{K_a, L\}}$, (F.19) holds in the set of *all* $2K_a$ -sparse vectors with probability exceeding

$$1 - \exp\left(-2 \max\{K_a, L\} \left(\frac{\beta}{2} - \log\left(\frac{eK_{\text{tot}}}{2K_a}\right)\right)\right) \quad (\text{F.20})$$

This probability exceeds $1 - \epsilon$ if

$$\beta \geq 2 \log\left(\frac{eK_{\text{tot}}}{2K_a}\right) + \frac{\log(2/\epsilon)}{\max\{K_a, L\}} \quad (\text{F.21})$$

We get that

$$m^* \geq \frac{(1 - \delta_{2K_a})g_{\min}^2}{2 \left(C' \left(2 \log\left(\frac{eK_{\text{tot}}}{2K_a}\right) + \frac{\log(2/\epsilon)}{\max\{K_a, L\}} \right) g_{\max} \max\left\{\frac{K_a}{L}, 1\right\} + \frac{N_0}{L} \right)^2} \quad (\text{F.22})$$

holds with probability exceeding $1 - \epsilon$.

Let $k_d = \|\mathbf{b}_2 - \mathbf{b}_1\|_0 \leq 2K_a$ denote the number of positions in which \mathbf{b}_2 and \mathbf{b}_1 differ, i.e. their Hamming distance. Then the Renyi divergence (F.8) can be lower bound as:

$$\mathcal{D}_t(p_{\mathbf{b}}, p_{\mathbf{b}^\circ}) \geq t \frac{m^*}{4} k_d \quad (\text{F.23})$$

Putting everything together, we can complete the union bound. Note that there are $\binom{K_a}{k_d} \binom{K_{\text{tot}} - K_a}{k_d} \leq (3eK_{\text{tot}}K_a)^{k_d}$ ways to choose a support which differs from the true support in k_d positions. Now, denote by \mathcal{C} the event that the pilot matrix \mathbf{A} is such that the RIP condition (F.14) holds, and the bound (F.22). Using (F.3), Theorem 15 and (F.23)

we get that

$$\begin{aligned}
 \mathbb{P}(\mathbf{b}^* \neq \mathbf{b}^\circ, \mathcal{C}) &\leq \sum_{\mathbf{b} \in \Theta_{K_a} \setminus \{\mathbf{b}^\circ\}} \exp\left(-\frac{M}{4} \mathcal{D}_{1/2}(p_{\mathbf{b}}, p_{\mathbf{b}^\circ})\right) \\
 &\leq \sum_{k_d=1}^{2K_a} (3eK_{\text{tot}}K_a)^{k_d} \exp\left(-M \frac{m^*}{4} k_d\right) \\
 &= \sum_{k_d=1}^{2K_a} \exp\left(-k_d \left(M \frac{m^*}{4} - \log(3eK_{\text{tot}}K_a)\right)\right)
 \end{aligned} \tag{F.24}$$

So let

$$M \geq \frac{4}{m^*} \log\left(3eK_{\text{tot}}K_a \frac{1+\epsilon}{\epsilon}\right) \tag{F.25}$$

which is precisely condition (3.13), then

$$\begin{aligned}
 \mathbb{P}(\mathbf{b}^* \neq \mathbf{b}^\circ, \mathcal{C}) &\leq \sum_{k_d=1}^{2K_a} \left(\frac{\epsilon}{1+\epsilon}\right)^{k_d} \\
 &\leq \epsilon
 \end{aligned} \tag{F.26}$$

Finally

$$\begin{aligned}
 \mathbb{P}(\mathbf{b}^* \neq \mathbf{b}^\circ) &\leq \mathbb{P}(\mathbf{b}^* \neq \mathbf{b}^\circ, \mathcal{C}) + \mathbb{P}(\bar{\mathcal{C}}) \\
 &\leq \epsilon + \epsilon + \exp(-C_\delta L)
 \end{aligned} \tag{F.27}$$

This concludes the proof of Theorem 11.

G. Proof of the RIP, Theorem 12

Let us first define some basic properties.

Definition 3 (Sub-Exponential Norm). *Let X be a real scalar random variable. Define the sub-exponential norm of X as*

$$\|X\|_{\psi_1} := \inf \left\{ t > 0 : \mathbb{E} \left[\exp \left(\frac{|X|}{t} \right) \right] \leq 2 \right\}. \quad (\text{G.1})$$

A well known property of sub-exponential variables is that

$$\mathbb{P}(|X| > t) \leq 2 \exp(-ct/\|X\|_{\psi_1}) \quad \forall t > 0 \quad (\text{G.2})$$

for some universal constant $c > 0$.

Definition 4 (Sub-Exponential Random Vector). *Let \mathbf{X} be a random vector in \mathbb{R}^n . \mathbf{X} is said to be sub-exponential if all its marginals are scalar sub-exponential random variables, i.e. if*

$$\sup_{\mathbf{x} \in S^{n-1}} \|\langle \mathbf{X}, \mathbf{x} \rangle\|_{\psi_1} < \infty \quad (\text{G.3})$$

then we define $\|\mathbf{X}\|_{\psi_1} := \sup_{\mathbf{x} \in S^{n-1}} \|\langle \mathbf{X}, \mathbf{x} \rangle\|_{\psi_1}$, where S^{n-1} is the unit sphere in \mathbb{R}^n .

Basic properties of sub-exponential random variables and vectors can be found e.g. in Ch. 2 and 3 of [128].

Definition 5 (Convex Concentration Property (2.2 in [129])). *Let \mathbf{X} be a random vector in \mathbb{R}^n . \mathbf{X} has the convex concentration property with constant K if for every 1-Lipschitz convex function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$, we have $\mathbb{E}[\phi(\mathbf{X})] < \infty$ and for every $t > 0$,*

$$\mathbb{P}(|\phi(\mathbf{X}) - \mathbb{E}[\phi(\mathbf{X})]| \geq t) \leq 2 \exp(-t^2/K^2) \quad (\text{G.4})$$

For the RIP of $\mathring{\mathbf{A}}/\sqrt{m}$ we first establish the following results for generic matrices $\mathbf{R} \in \mathbb{R}^{m \times N}$ with independent normalized columns.

Theorem 16. *Let $\mathbf{R} \in \mathbb{R}^{m \times N}$ be a matrix with independent columns $\mathbf{R}_{:,i}$, normalized such that $\|\mathbf{R}_{:,i}\|_2^2 = m$, with ψ_1 -norm at most ψ . Also assume that $N \geq m$. Then the RIP*

constant of \mathbf{R}/\sqrt{m} satisfies $\delta_{2s}(\mathbf{R}/\sqrt{m}) < \delta$ with probability $\geq 1 - \exp(-C'\sqrt{c_{\delta,\xi}m})$ for

$$2s = c_{\delta,\xi} \frac{m}{\log^2(eN/c_{\delta,\xi}m)}, \quad (\text{G.5})$$

where $c_{\delta,\xi} \leq \min\{1, \frac{\delta^2}{(3C\xi^2)^2}\}$ for any $\xi > \psi + 1$ and $C, C' > 0$ are universal constants. \square

Proof. We make use of the following generic RIP result from [107, Theorem 3.3] for matrices with i.i.d. sub-exponential columns:

Theorem 17. Let $m \geq 1$ and s, N be integers such that $1 \leq s \leq \min(N, m)$. Let $\mathbf{R}_{:,1}, \dots, \mathbf{R}_{:,N} \in \mathbb{R}^m$ be independent sub-exponential random vectors normalized such that $\mathbb{E}[\|\mathbf{R}_{:,i}\|^2] = m$ and let $\psi = \max_{i \leq N} \|\mathbf{R}_{:,i}\|_{\psi_1}$. Let $\theta' \in (0, 1)$, $K, K' \geq 1$ and set $\xi = \psi K + K'$. Then for the matrix \mathbf{R} with columns $\mathbf{R}_{:,i}$

$$\delta_s\left(\frac{\mathbf{R}}{\sqrt{m}}\right) \leq C\xi^2 \sqrt{\frac{s}{m}} \log\left(\frac{eN}{s\sqrt{s/m}}\right) + \theta' \quad (\text{G.6})$$

holds with probability larger than

$$1 - \exp\left(-cK\sqrt{s} \log\left(\frac{eN}{s\sqrt{s/m}}\right)\right) \quad (\text{G.7})$$

$$- \mathbb{P}\left(\max_{i \leq N} \|\mathbf{R}_{:,i}\|_2 \geq K'\sqrt{m}\right) - \mathbb{P}\left(\max_{i \leq N} \left|\frac{\|\mathbf{R}_{:,i}\|_2^2}{m} - 1\right| \geq \theta'\right), \quad (\text{G.8})$$

where $C, c > 0$ are universal constants. \square

In order to prove Theorem 16 we shall apply Theorem 17. Let us abbreviate $\delta_s = \delta_s\left(\frac{\mathbf{R}}{\sqrt{m}}\right)$. Since $\|\mathbf{R}_{:,i}\|_2^2 = m$, the last two terms in (G.8) vanish for all $K' > 1$ and $\theta' > 0$. Therefore, we consider the bound

$$\delta_s \leq C\xi^2 \sqrt{\frac{s}{m}} \log\left(\frac{eN}{s\sqrt{s/m}}\right) =: D, \quad (\text{G.9})$$

that holds with probability larger than

$$\mathbb{P}(\delta_s \leq D) \geq 1 - \exp\left(-cK\sqrt{s} \log\left(\frac{eN}{s\sqrt{s/m}}\right)\right) \quad (\text{G.10})$$

Let $s = cm/\log^2(e\frac{N}{cm})$ for any $0 < c \leq 1$. Note that the conditions $c \leq 1$ and $N \geq m$ guarantee that $\log(e\frac{N}{cm}) \geq 1$. Plugging into (G.9) we see that the RIP-constant satisfies

$$\delta_s \leq C\xi^2 \sqrt{c} \frac{\log(e(\frac{N}{cm})^{3/2} \log^3(e\frac{N}{cm}))}{\log(e\frac{N}{cm})} \quad (\text{G.11})$$

$$\leq C\xi^2 \sqrt{c} \left(\frac{3}{2} + \frac{3 \log \log e\frac{N}{cm}}{\log e\frac{N}{cm}} \right) \quad (\text{G.12})$$

$$\leq C\xi^2 \sqrt{c} \left(\frac{3}{2} + \frac{3}{e} \right) \quad (\text{G.13})$$

$$\leq 3C\xi^2 \sqrt{c} \quad (\text{G.14})$$

where in the first line we made use of $m \leq N$ and in the last line we used $\log \log x / \log x \leq 1/e$. This bound fails with probability:

$$\mathbb{P}(\delta_s > D) \leq \exp \left(-\hat{c}K \sqrt{s} \log \left(e \frac{N\sqrt{m}}{s^{3/2}} \right) \right) \quad (\text{G.15})$$

$$\leq \exp \left(-\hat{c}K \sqrt{s} \log \left(e \frac{N}{m} \right) \right) \quad (\text{G.16})$$

$$= \exp(-\hat{c}K \sqrt{c} \sqrt{m}) \quad (\text{G.17})$$

where in the second line we used $s \leq m$. The statement of Theorem 16 follows by choosing c small enough such that $\delta_s \leq \delta$. \square

We want to apply Theorem 16, which holds for real values matrices \mathbf{R} , to the matrix

$$\mathbb{A}^R := \sqrt{2}[\text{Re}(\mathring{\mathbb{A}}); \text{Im}(\mathring{\mathbb{A}})] \in \mathbb{R}^{2L(L-1) \times K_{\text{tot}}}, \quad (\text{G.18})$$

i.e., the real matrix obtained by stacking real and imaginary part of $\mathring{\mathbb{A}}$, with $m = 2L(L-1)$ and $N = K_{\text{tot}}$. For this we need to show that

- i The columns of \mathbb{A}^R are normalized to $\sqrt{2L(L-1)}$;
- ii The columns of \mathbb{A}^R are sub-exponential with ψ_1 norm independent of the dimension.

Consider the k -th column $\mathbb{A}_{:,k}^R$ of \mathbb{A}^R . We have

$$\begin{aligned}
 \|\text{Re}(\mathbb{A}_{:,k}^R); \text{Im}(\mathbb{A}_{:,k}^R)\|_2^2 &= \|\text{Re}(\mathbb{A}_{:,k}^R)\|_2^2 + \|\text{Im}(\mathbb{A}_{:,k}^R)\|_2^2 \\
 &= \|\mathbb{A}_{:,k}^R\|_2^2 \\
 &= \sum_{i \neq j} |a_{k,i} a_{k,j}^H|^2 \\
 &= \left(\sum_{i=1}^L |a_{k,i}|^2 \right)^2 - \sum_{i=1}^L |a_{k,i}|^2 \\
 &= (\|\mathbf{a}_k\|_2^2)^2 - \|\mathbf{a}_k\|_2^2 \\
 &= L(L-1),
 \end{aligned} \tag{G.19}$$

where we have used the normalization of the pilot matrix \mathbf{A} . This shows i.

For ii we need to show that all marginal distributions of the columns \mathbb{A}^R are sub-exponential. Note that for any vector $\mathbf{u} \in \mathbb{R}^{2L(L-1)}$ the marginal $\langle \mathbb{A}_{:,k}^R, \mathbf{u} \rangle$ can be expressed as a quadratic form in $\mathbf{a}_k^R := \sqrt{2}[\text{Re}(\mathbf{a}_k); \text{Im}(\mathbf{a}_k)]$ as the following calculation shows. Let $\mathbf{U}, \tilde{\mathbf{U}} \in \mathbb{R}^{L \times L}$ be two matrices with zeros on the diagonal such that $\mathbf{u} = [\text{vec}_{\text{non-diag}}(\mathbf{U}); \text{vec}_{\text{non-diag}}(\tilde{\mathbf{U}})]$. Then it holds:

$$\begin{aligned}
 \langle \mathbb{A}_{:,k}^R, \mathbf{u} \rangle &= \sqrt{2} \sum_{i \neq j} \left(\text{Re}(a_{k,i} a_{k,j}^H) U_{ij} + \text{Im}(a_{k,i} a_{k,j}^H) \tilde{U}_{ij} \right) \\
 &= (\mathbf{a}_k^R)^\top \mathbf{Q}_{\mathbf{u}} \mathbf{a}_k^R
 \end{aligned} \tag{G.20}$$

with

$$\mathbf{Q}_{\mathbf{u}} = \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbf{U} & \tilde{\mathbf{U}} \\ -\tilde{\mathbf{U}} & \mathbf{U} \end{pmatrix} \tag{G.21}$$

and therefore $\|\mathbf{Q}_{\mathbf{u}}\|_F^2 = \|\mathbf{u}\|_2^2$.

This form of $\mathbf{Q}_{\mathbf{u}}$ follows from the identities:

$$\text{Re}(a_{k,i} a_{k,j}^H) = \text{Re}(a_{k,i}) \text{Re}(a_{k,j}) + \text{Im}(a_{k,i}) \text{Im}(a_{k,j}) \tag{G.22}$$

$$\text{Im}(a_{k,i} a_{k,j}^H) = -\text{Re}(a_{k,i}) \text{Im}(a_{k,j}) + \text{Im}(a_{k,i}) \text{Re}(a_{k,j}) \tag{G.23}$$

We can now use the following concentration result for quadratic forms from [129] which states that a random vector which satisfies the convex concentration property also satisfies the following inequality, known as Hanson-Wright inequality [130]:

Theorem 18 (Theorem 2.5 in [129]). *Let \mathbf{X} be a mean zero random vector in \mathbb{R}^n , which satisfies the convex concentration property with constant B , then for any $n \times n$ matrix \mathbf{Y}*

and every $t > 0$,

$$\begin{aligned} & \mathbb{P}(|\mathbf{X}^\top \mathbf{Y} \mathbf{X} - \mathbb{E}[\mathbf{X}^\top \mathbf{Y} \mathbf{X}]| > t) \\ & \leq 2 \exp \left(-c \min \left(\frac{t^2}{2B^4 \|\mathbf{Y}\|_F^2}, \frac{t}{B^2 \|\mathbf{Y}\|_{op}} \right) \right) \end{aligned} \quad (\text{G.24})$$

□.

Note that a random variable with such a mixed tail behavior is especially sub-exponential. This can be seen by bounding its moments. Let Z be a random variable with

$$\mathbb{P}(|Z| > t) \leq 2 \exp \left(-c \min \left(\frac{t^2}{B^4 \|\mathbf{Y}\|_F^2}, \frac{t}{B^2 \|\mathbf{Y}\|_{op}} \right) \right) \quad (\text{G.25})$$

Since $\|\mathbf{Y}\|_{op} \leq \|\mathbf{Y}\|_F$, we have $\mathbb{P}(|Z| > t) \leq 2 \exp(-c \min(x(t)^2, x(t)))$ for $x(t) = \frac{t}{B^2 \|\mathbf{Y}\|_F}$. It follows

$$\begin{aligned} \mathbb{E}[|Z|^p] &= \int_0^\infty \mathbb{P}(|Z|^p > u) du = p \int_0^\infty \mathbb{P}(|Z| > t) t^{p-1} dt \\ &\leq 2p(B^2 \|\mathbf{Y}\|_{op})^p \left(\int_0^1 e^{-x^2} x^{p-1} dx + \int_1^\infty e^{-x} x^{p-1} dx \right) \\ &\leq 2p(B^2 \|\mathbf{Y}\|_{op})^p (\Gamma(p/2) + \Gamma(p)) \\ &\leq 4p(B^2 \|\mathbf{Y}\|_{op})^p \Gamma(p) \leq 4p(pB^2 \|\mathbf{Y}\|_{op})^p \end{aligned} \quad (\text{G.26})$$

where $\Gamma(\cdot)$ is the Gamma function. So

$$(\mathbb{E}[|Z|^p])^{\frac{1}{p}} \leq (4p)^{\frac{1}{p}} p B^2 \|\mathbf{Y}\|_{op} \leq c p B^2 \|\mathbf{Y}\|_{op} \quad (\text{G.27})$$

where $c = 4e^{1/e}$. (G.27) is equivalent to $\|Z\|_{\psi_1} \leq c B^2 \|\mathbf{Y}\|_{op}$ by elementary properties of sub-exponential random variables, e.g. [128, Proposition 2.7.1].

The convex concentration property was introduced in Definition 5. In our case the pilots $\mathbf{a}_k \in \mathbb{C}^L$ are distributed uniformly on the complex L -dimensional sphere of radius L , therefore the real versions $\mathbf{a}_k^R \in \mathbb{R}^{2L}$ are distributed uniformly on the sphere of radius $2L$. A classical result states that a spherical random variable $\mathbf{X} \sim \text{Unif}(\sqrt{n}S^{n-1})$ has the even stronger (non-convex) concentration property (e.g. [128, Theorem 5.1.4]):

Theorem 19 (Concentration on the Sphere). *Let $\mathbf{X} \sim \text{Unif}(\sqrt{n}S^{n-1})$ be uniformly distributed on the Euclidean sphere of radius \sqrt{n} . Then there is an universal constant $c > 0$, such that for every 1-Lipschitz function $f : \sqrt{n}S^{n-1} \rightarrow \mathbb{R}$*

$$\mathbb{P}(f(\mathbf{X}) - \mathbb{E}[f(\mathbf{X})]) \leq 2 \exp(-ct^2) \quad (\text{G.28})$$

□

So in particular the columns $\mathbb{A}_{:,k}^R = \mathbf{a}_k^R$ have the convex concentration property with some constant $c > 0$, independent of the dimension and it follows by (G.20) and Theorem 18 applied to $\mathbf{X} = \mathbf{a}_k^R$ and $\mathbf{Y} = \mathbf{Q}_{\mathbf{u}}$ that the marginals of $\langle \mathbb{A}_{:,k}^R, \mathbf{u} \rangle$ uniformly satisfy the tail bound of the Hanson-Wright inequality. As shown in (G.27), this implies that the columns of \mathbb{A}^R are sub-exponential with

$$\|\mathbb{A}_{:,k}^R\|_{\psi_1} = \max_{\mathbf{u} \in S^{2L(L-1)-1}} \|\langle \mathbb{A}_{:,k}^R, \mathbf{u} \rangle\|_{\psi_1} \leq C \quad (\text{G.29})$$

for some universal constant $C > 0$.

With this we can apply Theorems 16 and together with 18 and 19 it follows that \mathbb{A}^R , as defined in (G.18), has RIP of order $2s$ with RIP constant $\delta_{2s}(\mathbb{A}^R/\sqrt{2m}) < \delta$ as long as

$$2s \leq C_\delta \frac{m}{\log^2(eK_{\text{tot}}/m)}. \quad (\text{G.30})$$

Then, for the complex valued $\mathring{\mathbb{A}}$ it holds that

$$\left\| \frac{\mathring{\mathbb{A}}\mathbf{x}}{\sqrt{m}} \right\|_2 = \left\| \frac{\sqrt{2}[\text{Re}(\mathring{\mathbb{A}}); \text{Im}(\mathring{\mathbb{A}})]\mathbf{x}}{\sqrt{2m}} \right\|_2 = \left\| \frac{\mathbb{A}^R\mathbf{x}}{\sqrt{2m}} \right\|_2 \quad (\text{G.31})$$

for any $\mathbf{x} \in \mathbb{R}^{K_{\text{tot}}}$ and therefore the RIP of $\mathbb{A}^R/\sqrt{2m}$ implies the RIP of $\mathring{\mathbb{A}}/\sqrt{m}$ with the same constants, which concludes the proof of Theorem 12.

H. Proof of the Recovery Guarantee for NNLS, Theorem 13

Throughout this section let $\boldsymbol{\gamma}^*$ denote the NNLS estimate

$$\boldsymbol{\gamma}^* = \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}_+^{K_{\text{tot}}}} \|\mathbb{A}\boldsymbol{\gamma} - \mathbf{w}\|_2^2 \quad (\text{H.1})$$

as introduced in Section 3.2.2, where \mathbb{A} is the $L^2 \times K_{\text{tot}}$ matrix whose k -th column is given by $\text{vec}(\mathbf{a}_k \mathbf{a}_k^H)$ and

$$\mathbf{w} = \text{vec}(\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}} - N_0 \mathbf{I}_L), \quad (\text{H.2})$$

where $\widehat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ is assumed to be the empirical covariance matrix (3.6) of M iid samples from a Gaussian distribution $\mathcal{CN}(0, \boldsymbol{\Sigma}_{\mathbf{y}})$ with covariance matrix

$$\boldsymbol{\Sigma}_{\mathbf{y}} = \sum_{i=1}^{K_{\text{tot}}} \gamma_i^\circ \mathbf{a}_i \mathbf{a}_i^H + N_0 \mathbf{I}_L \quad (\text{H.3})$$

where $\boldsymbol{\gamma}^\circ = (\gamma_1^\circ, \dots, \gamma_{K_{\text{tot}}}^\circ) \in \mathbb{R}_+^{K_{\text{tot}}}$ is the true (unknown) activity pattern. So \mathbf{w} can be expressed as

$$\mathbf{w} = \mathbb{A}\boldsymbol{\gamma}^\circ + \mathbf{d} \quad (\text{H.4})$$

for $\mathbf{d} := \text{vec}(\boldsymbol{\Sigma}_{\mathbf{y}} - \widehat{\boldsymbol{\Sigma}}_{\mathbf{y}})$. Let us introduce some notation.

Definition 6 (Robust NSP (4.21 in [65])). $\mathbb{A} \in \mathbb{C}^{L^2 \times K_{\text{tot}}}$ is said to satisfy the robust ℓ_q nullspace property (NSP) of order s with parameters $0 < \rho < 1$ and $\tau > 0$ if

$$\|\mathbf{d}_S\|_q \leq \frac{\rho}{s^{1-1/q}} \|\mathbf{v}_{\bar{S}}\|_1 + \tau \|\mathbb{A}\mathbf{v}\|_2 \quad \forall \mathbf{v} \in \mathbb{R}^{K_{\text{tot}}} \quad (\text{H.5})$$

holds for all subsets $S \subset [K_{\text{tot}}]$ with $|S| \leq s$. The set \bar{S} denotes here the complement of S in $[K_{\text{tot}}]$.

Furthermore let the ℓ_1 -error of the best s -sparse approximation to $\boldsymbol{\gamma}^\circ$ be denoted as:

$$\sigma_s(\boldsymbol{\gamma}^\circ)_1 = \min_{\|\boldsymbol{\gamma}\|_0 \leq s} \|\boldsymbol{\gamma}^\circ - \boldsymbol{\gamma}\|_1 \quad (\text{H.6})$$

If γ° is assumed to actually be s -sparse, then we obviously have $\sigma_s(\gamma^\circ)_1 = 0$. The statement of Theorem 13 will be an immediate consequence of the following theorem:

Theorem 20. *If $\mathbb{A} \in \mathbb{C}^{L^2 \times K_{tot}}$ has the robust ℓ_2 NSP of order s with constants $\tau > 0$ and $\rho \in (0, 1)$ and there exists a $\mathbf{t} \in \mathbb{C}^{K_{tot}}$, such that $\mathbf{1} = \mathbb{A}^H \mathbf{t}$, where $\mathbf{1} := (1, \dots, 1)^\top$, then for $p \in [1, 2]$ the NNLS estimate γ^* in (H.1) satisfies*

$$\|\gamma^* - \gamma^\circ\|_p \leq \frac{2C\sigma_s(\gamma^\circ)_1}{s^{1-1/p}} + \frac{2D}{s^{\frac{1}{2}-\frac{1}{p}}} \left(\tau + \frac{\|\mathbf{t}\|_2}{s^{\frac{1}{2}}} \right) \|\mathbf{d}\|_2 \quad (\text{H.7})$$

with $C := \frac{(1+\rho)^2}{1-\rho}$, $D = \frac{(3+\rho)}{1-\rho}$ and $\mathbf{d} = \text{vec}(\Sigma_{\mathbf{y}} - \widehat{\Sigma}_{\mathbf{y}})$ □

Proof. This proof is adapted from [110] to our setting. First, we will need some implications which follow from the NSP [65, Theorem 4.25]. Assume that \mathbb{A} satisfies the robust NSP as stated in the theorem. Then, for any $p \in [1, 2]$ and for all $\mathbf{x}, \mathbf{z} \in \mathbb{R}^{K_{tot}}$,

$$\begin{aligned} \|\mathbf{x} - \mathbf{z}\|_p &\leq \frac{C}{s^{1-1/p}} (\|\mathbf{x}\|_1 - \|\mathbf{z}\|_1 + 2\sigma_s(\mathbf{x})_1) \\ &\quad + D\tau s^{1/p-1/2} \|\mathbb{A}(\mathbf{x} - \mathbf{z})\|_2 \end{aligned} \quad (\text{H.8})$$

holds, with C, D as defined in the statement of the theorem. If $\mathbf{x}, \mathbf{z} \geq 0$ are non-negative and there exists \mathbf{t} such that $\mathbf{1} = \mathbb{A}^H \mathbf{t}$ we use:

$$\begin{aligned} \|\mathbf{x}\|_1 - \|\mathbf{z}\|_1 &= \langle \mathbf{1}, \mathbf{x} - \mathbf{z} \rangle = \langle \mathbb{A}^H \mathbf{t}, \mathbf{x} - \mathbf{z} \rangle \\ &= \langle \mathbf{t}, \mathbb{A}(\mathbf{x} - \mathbf{z}) \rangle \leq \|\mathbf{t}\|_2 \|\mathbb{A}(\mathbf{x} - \mathbf{z})\|_2 \end{aligned} \quad (\text{H.9})$$

where we have used Cauchy-Schwarz inequality (note that $\langle \mathbf{t}, \mathbb{A}(\mathbf{x} - \mathbf{z}) \rangle$ is real). So inequality (H.8) implies:

$$\|\mathbf{x} - \mathbf{z}\|_p \quad (\text{H.10})$$

$$\leq \frac{2C\sigma_s(\mathbf{x})_1}{s^{1-1/p}} + \left(D\tau + \frac{C \cdot \|\mathbf{t}\|_2}{s^{1/2}} \right) \frac{\|\mathbb{A}(\mathbf{x} - \mathbf{z})\|_2}{s^{\frac{1}{2}-\frac{1}{p}}} \quad (\text{H.11})$$

Now, let's take $\mathbf{y} = \mathbb{A}\mathbf{x} + \mathbf{d}$. Since $\|\mathbb{A}(\mathbf{x} - \mathbf{z})\|_2 \leq \|\mathbb{A}\mathbf{z} - \mathbf{y}\|_2 + \|\mathbf{d}\|_2$ we get for all non-negative \mathbf{z} and \mathbf{x} :

$$\|\mathbf{x} - \mathbf{z}\|_p \quad (\text{H.12})$$

$$= \frac{2C\sigma_s(\mathbf{x})_1}{s^{1-1/p}} + \left(D\tau + \frac{C \cdot \|\mathbf{t}\|_2}{s^{1/2}} \right) \frac{\|\mathbb{A}\mathbf{z} - \mathbf{y}\|_2 + \|\mathbf{d}\|_2}{s^{\frac{1}{2}-\frac{1}{p}}} \quad (\text{H.13})$$

Now take $\mathbf{z} = \boldsymbol{\gamma}^*$ and $\mathbf{x} = \boldsymbol{\gamma}^\circ$, then $\mathbf{y} = \mathbf{w}$ (see (H.2)). Since $\boldsymbol{\gamma}^\circ \in \mathbb{R}_+^{K_{\text{tot}}}$ is itself a feasible point of the minimization we have $\min_{\boldsymbol{\gamma} \in \mathbb{R}_+^{K_{\text{tot}}}} \|\mathbb{A}\boldsymbol{\gamma} - \mathbf{b}\|_2 \leq \|\mathbf{d}\|_2$, yielding:

$$\|\boldsymbol{\gamma}^* - \boldsymbol{\gamma}^\circ\|_p \leq \frac{2C\sigma_s(\boldsymbol{\gamma}^\circ)_1}{s^{1-1/p}} + 2 \left(D\tau + \frac{C \cdot \|\mathbf{t}\|_2}{s^{1/2}} \right) \frac{\|\mathbf{d}\|_2}{s^{\frac{1}{2}-\frac{1}{p}}} \quad (\text{H.14})$$

It is easily checked that $C \leq D$ for $\rho \in (0, 1)$, which gives the result. \square

In our case we choose $\mathbf{t} = t \cdot \text{vec}(\mathbf{I}_L) \in \mathbb{R}^{L^2}$ with some $t > 0$. Let \mathbb{A}_k be the k -th column of \mathbb{A} . It holds that

$$\mathbb{A}_k^H \text{vec}(\mathbf{I}_L) = \text{trace}(\mathbf{a}_k \mathbf{a}_k^H) = \|\mathbf{a}_k\|_2^2. \quad (\text{H.15})$$

Using the normalization of the pilots $\|\mathbf{a}_k\|_2^2 = L$, we get that:

$$\mathbb{A}^H \mathbf{t} = tL \cdot \mathbf{1}$$

so $t = 1/L$, and therefore $\|\mathbf{t}\|_2^2 = 1/L$ gives the desired condition $\mathbb{A}^H \mathbf{t} = \mathbf{1}$. Before we can make use of Theorem 20 it remains to show that \mathbb{A} has the robust ℓ_2 -NSP with high probability. To this end, we will restrict to those measurements which are related to the isotropic part of \mathbb{A} . Let $\mathring{\mathbb{A}}$ be the centered version of \mathbb{A} defined in (3.28). Now it is easy to check (revert the vectorization) that this special structure gives us the inequality:

$$\|\mathbb{A}\mathbf{v}\|_2^2 = \|\mathring{\mathbb{A}}\mathbf{v}\|_2^2 + \|\mathbb{A}^{\text{diag}}\mathbf{v}\|_2^2 \geq \|\mathring{\mathbb{A}}\mathbf{v}\|_2^2 \quad (\text{H.16})$$

where $\mathbb{A}^{\text{diag}} \in \mathbb{C}^{L \times K_{\text{tot}}}$ is defined as the non-isotropic part of \mathbb{A} with its k -th column defined by $\mathbb{A}_{:,k}^{\text{diag}} = \text{vec}(\text{diag}(\mathbf{a}_k \mathbf{a}_k^H))$. This shows that if $\mathring{\mathbb{A}}$ has the ℓ_2 -NSP of order s with constants τ and ρ , then so does \mathbb{A} , since

$$\begin{aligned} \|\mathbf{v}_S\|_2 &= \frac{\rho}{\sqrt{s}} \|\mathbf{v}_{\bar{S}}\|_1 + \tau \|\mathring{\mathbb{A}}\mathbf{v}\|_2 \\ &\leq \frac{\rho}{\sqrt{s}} \|\mathbf{v}_{\bar{S}}\|_1 + \tau \|\mathbb{A}\mathbf{v}\|_2 \end{aligned} \quad (\text{H.17})$$

holds for all subsets $S \subset [K_{\text{tot}}]$ with $|S| \leq s$. It is well-known that the robust ℓ_2 -NSP of order s is implied by the RIP of order $2s$ with sufficiently small constants [65]. The following theorem specifies how RIP is related to the ℓ_2 -NSP

Theorem 21. *If $\mathring{\mathbb{A}}$ has RIP of order $2s$ with a constant bound as $\delta_{2s}(\mathring{\mathbb{A}}/\sqrt{m}) \leq \delta < 4/\sqrt{41} \approx 0.62$ then $\mathring{\mathbb{A}}/\sqrt{m}$ has the robust ℓ_2 -NSP of order s with parameters ρ and τ' with $\rho \leq \delta/(\sqrt{1-\delta^2} - \delta/4)$ and $\tau' \leq \sqrt{1+\delta}/(\sqrt{1-\delta^2} - \delta/4)$.*

Furthermore $\mathring{\mathbb{A}}$ has the robust ℓ_2 -NSP of order s with parameters ρ and $\tau = \tau'/\sqrt{m}$. \square

Proof. The first part is shown in [65, Theorem 6.13]. The last statements follows immediately from

$$\tau L \left\| \frac{1}{L} \mathring{\mathbb{A}} \mathbf{v} \right\|_2 = \tau \left\| \mathring{\mathbb{A}} \mathbf{v} \right\|_2 \quad (\text{H.18})$$

□

Theorem 12 establishes the RIP of $\mathring{\mathbb{A}}/\sqrt{L(L-1)}$ under the assumptions of 13. If we fix $\delta < 4/\sqrt{41}$, Theorem 21 implies the robust ℓ_2 -NSP of order s for $\mathring{\mathbb{A}}$ with explicit bounds on τ and ρ . For example, $\delta = 0.5$ gives $\rho < 0.68$ and $\tau < 3/L$. As shown in (H.17) the robust ℓ_2 -NSP of $\mathring{\mathbb{A}}$ implies the ℓ_2 -NSP of the uncentered version \mathbb{A} of the same order with the same constants. Finally, the application of Theorem 20 concludes the proof of Theorem 13.

I. Analysis of Error of the Sample Covariance Matrix

Let $\mathbf{\Sigma}_{\mathbf{y}} \in \mathbb{R}^{L \times L}$ be fixed and let $\{\mathbf{y}(t) : t \in [M]\}$ be M i.i.d. samples from $\mathcal{CN}(0, \mathbf{\Sigma}_{\mathbf{y}})$. We first consider the simple case where $\mathbf{\Sigma}_{\mathbf{y}}$ is diagonal, given by $\mathbf{\Sigma}_{\mathbf{y}} = \text{diag}(\boldsymbol{\beta})$ and let $\mathbf{\Delta} = \hat{\mathbf{\Sigma}}_{\mathbf{y}} - \mathbf{\Sigma}_{\mathbf{y}}$ be the deviation of the sample covariance matrix from its mean. Then $\|\mathbf{d}\|_2 = \|\mathbf{\Delta}\|_F$ and the (i, j) -th component of $\mathbf{\Delta}$ is given by

$$\Delta_{ij} = \frac{1}{M} \sum_{t \in [M]} y_i(t) y_j^*(t) - \beta_i \delta_{ij} \quad (\text{I.1})$$

$$= \frac{\sqrt{\beta_i \beta_j}}{M} \sum_{t \in [M]} \left(\frac{y_i(t)}{\sqrt{\beta_i}} \frac{y_j^*(t)}{\sqrt{\beta_j}} - \delta_{ij} \right) \quad (\text{I.2})$$

where $\delta_{ij} = 1_{\{i=j\}}$ denotes the discrete delta function. Let $Y_{ij}(t) := \frac{y_i(t)}{\sqrt{\beta_i}} \frac{y_j^*(t)}{\sqrt{\beta_j}} - \delta_{ij}$. Then

$$|\Delta_{ij}|^2 = \frac{\beta_i \beta_j}{M^2} \left| \sum_{t=1}^M Y_{ij}(t) \right|^2 \quad (\text{I.3})$$

Since all $Y_{ij}(t)$ are zero mean and are independent for fixed i, j . Therefore the variance of their sum $\mathbb{E} \left[\left| \sum_{t=1}^M Y_{ij}(t) \right|^2 \right]$ is the sum of their variances. In the following we show that $\mathbb{E}[|Y_{ij}|^2] = 1$ for all i, j . For $i \neq j$, we have that

$$\begin{aligned} \mathbb{E}[|Y_{ij}|^2] &= \frac{\mathbb{E}[|y_i(t) y_j(t)^*|^2]}{\beta_i \beta_j} \\ &\stackrel{(a)}{=} \frac{\mathbb{E}[|y_i(t)|^2]}{\beta_i} \frac{\mathbb{E}[|y_j(t)|^2]}{\beta_j} \\ &= 1, \end{aligned} \quad (\text{I.4})$$

where in (a) we used the independence of the different components of $\mathbf{y}(t)$. Also, for $i = j$, we have that

$$\begin{aligned}\mathbb{E}[|Y_{ij}|^2] &= \mathbb{E}\left[\left|\frac{|y_i(t)|^2}{\beta_i} - 1\right|^2\right] \\ &= \frac{\mathbb{E}[|y_i(t)|^4]}{\beta_i^2} - 2\frac{\mathbb{E}[|y_i(t)|^2]}{\beta_i} + 1 \\ &\stackrel{(a)}{=} 2 - 2 + 1 \\ &= 1,\end{aligned}\tag{I.5}$$

where in (a) we used the identity $\mathbb{E}[|y_i(t)|^4] = 2\mathbb{E}[|y_i(t)|^2]^2$ for complex Gaussian random variables. Overall, from (I.4) and (I.5), we can write $\mathbb{E}[|\Delta_{ij}|^2] = \frac{\beta_i\beta_j}{M}$. Thus, we have that

$$\begin{aligned}\mathbb{E}[\|\Delta\|_F^2] &= \sum_{ij} \mathbb{E}[|\Delta_{ij}|^2] = \frac{\sum_{i,j} \beta_i\beta_j}{M} \\ &= \frac{(\sum \beta_i)^2}{M} = \frac{\text{trace}(\Sigma_{\mathbf{y}})^2}{M}.\end{aligned}\tag{I.6}$$

To see how fast $\|\Delta\|_F$ concentrates around its mean, note that for fixed i, j the $Y_{ij}(t)$ are independent sub-exponential random variables with sub-exponential norm ≤ 1 (see e.g. [128, Lemma 2.7.7]). Therefore, by the elemental Bernstein inequality we can estimate that for any $\alpha > 0$

$$\begin{aligned}\mathbb{P}\left(\left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \alpha\right) &= \mathbb{P}\left(\left|\sum_{t=1}^M Y_{ij}(t)\right| > \sqrt{\alpha}\right) \\ &\leq 2\exp(-c\min\{\alpha/M, \sqrt{\alpha}\})\end{aligned}\tag{I.7}$$

for some universal constant $c > 0$. By a union bound we can see that

$$\begin{aligned}\mathbb{P}\left(\min_{i,j} \left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \alpha\right) \\ \leq \binom{L}{2} \mathbb{P}\left(\left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \alpha\right) \\ \leq 2\exp(2\log(eL) - c\min\{\alpha/M, \sqrt{\alpha}\})\end{aligned}\tag{I.8}$$

By choosing α properly we can get the following statement:

Theorem 22. *Let $\epsilon > 0$*

$$\|\Delta\|_F \leq \frac{\text{trace}(\Sigma_{\mathbf{y}})}{\sqrt{M}} \sqrt{\frac{\log\left(\frac{2(eL)^2}{\epsilon}\right)}{c}} \quad (\text{I.9})$$

holds with probability exceeding $1 - \epsilon$, if $cM > \log(2(eL)^2/\epsilon)$, where $c > 0$ is the constant in (I.8). \square

Proof. In (I.8) choose $\alpha = M\delta$ with $\delta = \log(2(eL)^2/\epsilon)/c$. Then $\min\{\alpha/M, \sqrt{\alpha}\} = \min\{\delta, \sqrt{\delta M}\}$. Under the condition on M stated in the Theorem, $\min\{\delta, \sqrt{\delta M}\} = \delta$. So

$$\begin{aligned} & \mathbb{P}\left(\min_{i,j} \left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \delta M\right) \\ & \leq 2 \exp(2 \log(eL) - 2 \log(eL) + \log(\epsilon/2)) \\ & = \epsilon. \end{aligned} \quad (\text{I.10})$$

and the statement of the Theorem follows from

$$\begin{aligned} & \mathbb{P}\left(\|\Delta\|_F > \frac{\text{trace}(\Sigma_{\mathbf{y}})}{\sqrt{M/\delta}}\right) \\ & = \mathbb{P}\left(\|\Delta\|_F^2 > \frac{\text{trace}(\Sigma_{\mathbf{y}})^2}{M/\delta}\right) \\ & = \mathbb{P}\left(\sum_{ij} \frac{\beta_i \beta_j}{M^2} \left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \frac{\text{trace}(\Sigma_{\mathbf{y}})^2}{M/\delta}\right) \\ & \leq \mathbb{P}\left(\min_{ij} \left|\sum_{t=1}^M Y_{ij}(t)\right|^2 > \delta M\right) \\ & \leq \epsilon \end{aligned} \quad (\text{I.11})$$

where in the second equality we used (I.3) and in the last inequality we used (I.10). \square

Now, assume that the covariance matrix $\Sigma_{\mathbf{y}}$ is not in a diagonal form and let $\Sigma_{\mathbf{y}} = \mathbf{U} \text{diag}(\boldsymbol{\beta}) \mathbf{U}^H$ be the singular value decomposition of $\Sigma_{\mathbf{y}}$. By multiplying all the vectors $\mathbf{y}(t)$ by the orthogonal matrix \mathbf{U}^H to whiten them and noting the fact that multiplying by \mathbf{U}^H does not change the Frobenius norm of a matrix, we can see that the bound in Theorem 22, which depends on $\Sigma_{\mathbf{y}}$ only through its trace, holds true in general also for non-diagonal covariance matrices. Finally, since in Theorem 13 $\Sigma_{\mathbf{y}} = \sum_{k=1}^{K_{\text{tot}}} \gamma_k \mathbf{a}_k \mathbf{a}_k^H + N_0 \mathbf{I}_L$ and the

pilot sequences satisfy $\|\mathbf{a}_k\|_2^2 = L$, it holds that

$$\text{trace}(\mathbf{\Sigma}_{\mathbf{y}}) = \sum_{k=1}^{K_{\text{tot}}} \gamma_k \text{trace}(\mathbf{a}_k \mathbf{a}_k^H) + N_0 \text{trace}(\mathbf{I}_L) = L(\|\boldsymbol{\gamma}\|_1 + N_0), \quad (\text{I.12})$$

which gives (3.34) and (3.35).

Remark 5. *It is worthwhile to mention that although (I.6) was derived under the Gaussianity of the observations $\{\mathbf{y}(t) : t \in [M]\}$, the result can be easily modified for general distribution of the components of $\mathbf{y}(t)$. More specifically, let us define*

$$\max_i \frac{\mathbb{E}[|y_i(t)|^4]}{\mathbb{E}[|y_i(t)|^2]^2} =: \varsigma < \infty. \quad (\text{I.13})$$

Then, using (I.4) and applying (I.13) to (I.5), we can obtain the following upper bound

$$\mathbb{E}[\|\boldsymbol{\Delta}\|_{\mathbf{F}}^2] \leq \max\{\varsigma - 1, 1\} \times \frac{\sum_{i,j} \beta_i \beta_j}{M} \quad (\text{I.14})$$

$$\leq \max\{\varsigma - 1, 1\} \times \frac{\text{trace}(\mathbf{\Sigma}_{\mathbf{y}})^2}{M}, \quad (\text{I.15})$$

which is equivalent to (I.6) up to the constant multiplicative factor $\max\{\varsigma - 1, 1\}$. \diamond

Bibliography

- [1] E. Dahlman, S. Parkvall, and J. Skold, *4G: LTE/LTE-Advanced for Mobile Broad-band*. Amsterdam ; New York: Academic Press, 2. edition ed., Oct. 2013.
- [2] E. Dahlman, S. Parkvall, and J. Skold, *5G NR: The Next Generation Wireless Access Technology*. London, United Kingdom ; San Diego, CA: Academic Press, Aug. 2018.
- [3] K. Senel and E. G. Larsson, “Grant-Free Massive MTC-Enabled Massive MIMO: A Compressive Sensing Approach,” *IEEE Trans. Commun.*, vol. 66, pp. 6164–6175, Dec. 2018.
- [4] L. Liu and W. Yu, “Massive Connectivity with Massive MIMO-Part I: Device Activity Detection and Channel Estimation,” *IEEE Trans. Signal Process.*, vol. 66, pp. 2933–2946, June 2018.
- [5] L. Liu and W. Yu, “Massive Connectivity with Massive MIMO-Part II: Achievable Rate Characterization,” *IEEE Trans. Signal Process.*, vol. 66, pp. 2947–2959, June 2018.
- [6] L. Liu, E. G. Larsson, W. Yu, P. Popovski, C. Stefanovic, and E. de Carvalho, “Sparse Signal Processing for Grant-Free Massive Connectivity: A Future Paradigm for Random Access Protocols in the Internet of Things,” *IEEE Signal Process. Mag.*, vol. 35, pp. 88–99, Sept. 2018.
- [7] Y. Wu, X. Gao, S. Zhou, W. Yang, Y. Polyanskiy, and G. Caire, “Massive Access for Future Wireless Communication Systems,” *ArXiv191012678 Cs Math*, Feb. 2020.
- [8] X. Chen, T.-Y. Chen, and D. Guo, “Capacity of Gaussian Many-Access Channels,” *IEEE Trans. Inf. Theory*, vol. 63, pp. 3516–3539, June 2017.
- [9] Y. Polyanskiy, “A perspective on massive random-access,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2523–2527, June 2017.
- [10] R. Ahlswede, “Multi-way communication channels,” in *Second International Symposium on Information Theory: Tsahkadsor, Armenia, USSR, Sept. 2 - 8, 1971*, 1973.

- [11] H. Liao, "Multiple access channels (Ph.D. Thesis abstr.)," *IEEE Trans. Inf. Theory*, vol. 19, pp. 253–253, Mar. 1973.
- [12] A. El Gamal and T. M. Cover, "Multiple User Information Theory," *Proc IEEE*, vol. 68, no. 1, pp. 1466–1483, 1980.
- [13] R. Gallager, "A perspective on multiaccess channels," *IEEE Trans. Inf. Theory*, vol. 31, pp. 124–142, Mar. 1985.
- [14] N. Abramson, "THE ALOHA SYSTEM: Another alternative for computer communications," in *Proceedings of the November 17-19, 1970, Fall Joint Computer Conference*, AFIPS '70 (Fall), (Houston, Texas), pp. 281–285, Association for Computing Machinery, Nov. 1970.
- [15] E. Casini, R. De Gaudenzi, and O. Del Rio Herrero, "Contention Resolution Diversity Slotted ALOHA (CRDSA): An Enhanced Random Access Scheme for Satellite Access Packet Networks," *IEEE Trans. Wirel. Commun.*, vol. 6, pp. 1408–1419, Apr. 2007.
- [16] G. Liva, "Graph-Based Analysis and Optimization of Contention Resolution Diversity Slotted ALOHA," *IEEE Trans. Commun.*, vol. 59, pp. 477–487, Feb. 2011.
- [17] E. Paolini, G. Liva, and M. Chiani, "Coded Slotted ALOHA: A Graph-Based Method for Uncoordinated Multiple Access," *IEEE Trans. Inf. Theory*, vol. 61, pp. 6815–6832, Dec. 2015.
- [18] J. Massey and P. Mathys, "The collision channel without feedback," *IEEE Trans. Inf. Theory*, vol. 31, pp. 192–204, Mar. 1985.
- [19] A. Ephremides and B. Hajek, "Information theory and communication networks: An unconsummated union," *IEEE Trans. Inf. Theory*, vol. 44, pp. 2416–2434, Oct. 1998.
- [20] M. Berlioli, A. Munari, G. Cocco, and G. Liva, *Modern Random Access Protocols*. Now Publishers, 2016.
- [21] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite block-length regime," *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.
- [22] O. Ordentlich and Y. Polyanskiy, "Low Complexity Schemes for the Random Access Gaussian Channel," *2017 IEEE Int. Symp. Inf. Theory ISIT*, pp. 2533–2537, 2017.
- [23] A. Vem, K. R. Narayanan, J. Cheng, and J.-F. Chamberland, "A user-independent serial interference cancellation based coding scheme for the unsourced random access Gaussian channel," in *2017 IEEE Information Theory Workshop (ITW)*, (Kaohsiung, Taiwan), pp. 121–125, IEEE, Nov. 2017.

-
- [24] V. K. Amalladinne, A. Vem, D. K. Soma, K. R. Narayanan, and J.-F. Chamberland, "A Coupled Compressive Sensing Scheme for Uncoordinated Multiple Access," *arXiv:1809.04745*, Sept. 2018.
- [25] E. Marshakov, G. Balitskiy, K. Andreev, and A. Frolov, "A Polar Code Based Unsourced Random Access for the Gaussian MAC," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, pp. 1–5, Sept. 2019.
- [26] A. K. Pradhan, V. K. Amalladinne, K. R. Narayanan, and J.-F. Chamberland, "Polar Coding and Random Spreading for Unsourced Multiple Access," *arXiv:1911.01009*, Nov. 2019.
- [27] A. Pradhan, V. Amalladinne, A. Vem, K. R. Narayanan, and J.-F. Chamberland, "A Joint Graph Based Coding Scheme for the Unsourced Random Access Gaussian Channel," *ArXiv190605410 Cs Math*, June 2019.
- [28] S. S. Kowshik, K. Andreev, A. Frolov, and Y. Polyanskiy, "Energy efficient coded random access for the wireless uplink," *IEEE Trans. Commun.*, pp. 1–1, 2020.
- [29] A. Joseph and A. Barron, "Least squares superposition codes of moderate dictionary size are reliable at rates up to capacity," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 2541–2557, 2012.
- [30] C. Rush, A. Greig, and R. Venkataramanan, "Capacity-achieving Sparse Superposition Codes via Approximate Message Passing Decoding," *IEEE Trans. Inf. Theory*, vol. 63, no. 3, pp. 2016–2020, 2017.
- [31] J. Barbier and F. Krzakala, "Approximate message-passing decoder and capacity-achieving sparse superposition codes," *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 1–32, 2017.
- [32] D. L. Donoho, A. Maleki, and A. Montanari, "Message Passing Algorithms for Compressed Sensing," *PNAS*, vol. 106, pp. 18914–18919, Nov. 2009.
- [33] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," *IEEE Int. Symp. Inf. Theory Proc.*, pp. 1–18, 2011.
- [34] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, 2011.
- [35] R. Venkataramanan, S. Tatikonda, and A. Barron, "Sparse Regression Codes," *CIT*, vol. 15, pp. 1–195, June 2019.

- [36] R. Berthier, A. Montanari, and P.-M. Nguyen, “State Evolution for Approximate Message Passing with Non-Separable Functions,” *arXiv:1708.03950*, Aug. 2017.
- [37] T. Tanaka, “A statistical-mechanics approach to large-system analysis of CDMA multiuser detectors,” *IEEE Trans. Inf. Theory*, vol. 48, no. 11, pp. 2888–2910, 2002.
- [38] D. Guo and S. Verdú, “Randomly spread CDMA: Asymptotics via statistical physics,” *IEEE Trans. Inf. Theory*, vol. 51, pp. 1983–2010, June 2005.
- [39] D. Guo, D. Baron, and S. Shamai, “A single-letter characterization of optimal noisy compressed sensing,” in *2009 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 52–59, Sept. 2009.
- [40] S. Rangan, A. K. Fletcher, and V. K. Goyal, “Asymptotic Analysis of MAP Estimation via the Replica Method and Applications to Compressed Sensing,” *IEEE Trans. Inform. Theory*, vol. 58, pp. 1902–1923, Mar. 2012.
- [41] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová, “Probabilistic reconstruction in compressed sensing: Algorithms, phase diagrams, and threshold achieving matrices,” *J. Stat. Mech. Theory Exp.*, vol. 2012, no. 08, p. P08009, 2012.
- [42] J. Barbier, N. Macris, M. Dia, and F. Krzakala, “Mutual Information and Optimality of Approximate Message-Passing in Random Linear Estimation,” *arXiv:1701.05823*, Jan. 2017.
- [43] G. Reeves and H. D. Pfister, “The Replica-Symmetric Prediction for Compressed Sensing with Gaussian Matrices is Exact,” pp. 665–669, 2016.
- [44] J. Barbier and N. Macris, “The adaptive interpolation method: A simple scheme to prove replica formulas in Bayesian inference,” *Probab. Theory Relat. Fields*, vol. 174, pp. 1133–1185, Aug. 2019.
- [45] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, UK ; New York: Cambridge University Press, illustrated edition ed., May 2005.
- [46] T. L. Marzetta and H. Yang, *Fundamentals of Massive MIMO*. Cambridge University Press, Nov. 2016.
- [47] T. L. Marzetta, “Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas,” *IEEE Trans. Wirel. Commun.*, vol. 9, pp. 3590–3600, Nov. 2010.
- [48] S. Haghighatshoar, P. Jung, and G. Caire, “Improved Scaling Law for Activity Detection in Massive MIMO Systems,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 381–385, June 2018.

-
- [49] K. Senel and E. G. Larsson, "Device Activity and Embedded Information Bit Detection Using AMP in Massive MIMO," in *2017 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6, Dec. 2017.
 - [50] C. Wang, O. Y. Bursalioglu, H. Papadopoulos, and G. Caire, "On-the-Fly Large-Scale Channel-Gain Estimation for Massive Antenna-Array Base Stations," in *2018 IEEE International Conference on Communications (ICC)*, pp. 1–6, May 2018.
 - [51] M. E. Tipping, "Sparse bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Sept. 2001.
 - [52] A. C. Faul and M. E. Tipping, "Analysis of Sparse Bayesian Learning," in *Advances in Neural Information Processing Systems 14* (T. G. Dietterich, S. Becker, and Z. Ghahramani, eds.), pp. 383–389, MIT Press, 2002.
 - [53] M. E. Tipping, A. Faul, J. J. T. Avenue, and J. J. T. Avenue, "Fast Marginal Likelihood Maximisation for Sparse Bayesian Models," in *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, pp. 3–6, 2003.
 - [54] J. Palmer, B. D. Rao, and D. P. Wipf, "Perspectives on Sparse Bayesian Learning," in *Advances in Neural Information Processing Systems 16* (S. Thrun, L. K. Saul, and B. Schölkopf, eds.), pp. 249–256, MIT Press, 2004.
 - [55] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum Likelihood from Incomplete Data Via the *EM* Algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, pp. 1–22, Sept. 1977.
 - [56] T. Moon, "The expectation-maximization algorithm," *IEEE Signal Process. Mag.*, vol. 13, pp. 47–60, Nov. 1996.
 - [57] D. Wipf and B. Rao, "Sparse Bayesian Learning for Basis Selection," *IEEE Trans. Signal Process.*, vol. 52, pp. 2153–2164, Aug. 2004.
 - [58] S. Ji, Y. Xue, and L. Carin, "Bayesian Compressive Sensing," *IEEE Trans. Signal Process.*, vol. 56, pp. 2346–2356, June 2008.
 - [59] D. P. Wipf and B. D. Rao, "An Empirical Bayesian Strategy for Solving the Simultaneous Sparse Approximation Problem," *IEEE Trans. Signal Process.*, vol. 55, pp. 3704–3716, July 2007.
 - [60] G. Tang and A. Nehorai, "Performance Analysis for Sparse Support Recovery," *IEEE Trans. Inf. Theory*, vol. 56, pp. 1383–1399, Mar. 2010.

- [61] P. Pal and P. P. Vaidyanathan, "Parameter identifiability in Sparse Bayesian Learning," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (Florence, Italy), pp. 1851–1855, IEEE, May 2014.
- [62] P. Pal and P. P. Vaidyanathan, "Pushing the Limits of Sparse Support Recovery Using Correlation Information," *IEEE Trans. Signal Process.*, vol. 63, pp. 711–726, Feb. 2015.
- [63] A. Koochakzadeh, H. Qiao, and P. Pal, "On Fundamental Limits of Joint Sparse Support Recovery Using Certain Correlation Priors," *IEEE Trans. Signal Process.*, vol. 66, pp. 4612–4625, Sept. 2018.
- [64] O. Balkan, K. Kreutz-Delgado, and S. Makeig, "Localization of More Sources Than Sensors via Jointly-Sparse Bayesian Learning," *IEEE Signal Process. Lett.*, vol. 21, pp. 131–134, Feb. 2014.
- [65] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis, Birkhäuser Basel, 2013.
- [66] S. Khanna and C. R. Murthy, "On the Support Recovery of Jointly Sparse Gaussian Sources using Sparse Bayesian Learning," *ArXiv170304930 Cs Math*, Mar. 2017.
- [67] S. Khanna and C. R. Murthy, "Corrections to "On the Restricted Isometry of the Columnwise Khatri–Rao Product"," *IEEE Trans. Signal Process.*, vol. 67, pp. 2387–2388, May 2019.
- [68] Z. Chen, F. Sohrabi, Y.-F. Liu, and W. Yu, "Phase Transition Analysis for Covariance Based Massive Random Access with Massive MIMO," *ArXiv200304175 Cs Math*, Mar. 2020.
- [69] H. Abeida, Q. Zhang, J. Li, and N. Merabtine, "Iterative Sparse Asymptotic Minimum Variance Based Approaches for Array Processing," *IEEE Trans. Signal Process.*, vol. 61, pp. 933–944, Feb. 2013.
- [70] G.-O. Glentis, K. Zhao, A. Jakobsson, H. Abeida, and J. Li, "SAR imaging via efficient implementations of sparse ML approaches," *Signal Processing*, vol. 95, pp. 15–26, Feb. 2014.
- [71] X. Yang, G. Li, and Z. Zheng, "DOA Estimation of Noncircular Signal Based on Sparse Representation," *Wireless Pers Commun*, vol. 82, pp. 2363–2375, June 2015.

-
- [72] L. Ramesh and C. R. Murthy, "Sparse Support Recovery Via Covariance Estimation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6633–6637, Apr. 2018.
- [73] E. Arikan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Trans. Inform. Theory*, vol. 55, pp. 3051–3073, July 2009.
- [74] I. Tal and A. Vardy, "List Decoding of Polar Codes," *IEEE Trans. Inf. Theory*, vol. 61, pp. 2213–2226, May 2015.
- [75] A. Balatsoukas-Stimming, M. B. Parizi, and A. Burg, "LLR-based Successive Cancellation List Decoding of Polar Codes," *IEEE Trans. Signal Process.*, vol. 63, pp. 5165–5179, Oct. 2015.
- [76] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-Free Massive MIMO: Uniformly great service for everyone," in *2015 IEEE 16th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, (Stockholm, Sweden), pp. 201–205, IEEE, June 2015.
- [77] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: Possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal Process. Mag.*, vol. 22, pp. 41–53, July 2005.
- [78] A. Weiss, "On the accuracy of a cellular location system based on RSS measurements," *IEEE Trans. Veh. Technol.*, vol. 52, pp. 1508–1518, Nov. 2003.
- [79] K. Schaubach, N. Davis, and T. Rappaport, "A ray tracing method for predicting path loss and delay spread in microcellular environments," in *[1992 Proceedings] Vehicular Technology Society 42nd VTS Conference - Frontiers of Technology*, (Denver, CO, USA), pp. 932–935, IEEE, 1992.
- [80] R. Levie, Ç. Yapar, G. Kutyniok, and G. Caire, "RadioUNet: Fast Radio Map Estimation with Convolutional Neural Networks," *ArXiv191109002 Cs Eess Math Stat*, May 2020.
- [81] A. Fengler, P. Jung, and G. Caire, "SPARCs for Unsourced Random Access," *ArXiv190106234 Cs Math*, June 2020.
- [82] A. Fengler, S. Haghighatshoar, P. Jung, and G. Caire, "Non-Bayesian Activity Detection, Large-Scale Fading Coefficient Estimation, and Unsourced Random Access with a Massive MIMO Receiver," *ArXiv191011266 Cs Math*, Aug. 2020.

- [83] A. Fengler, P. Jung, and G. Caire, “SPARCs and AMP for Unsourced Random Access,” *IEEE Int. Symp. Inf. Theory Proc.*, pp. 2843–2847, July 2019.
- [84] A. Fengler, S. Haghghatshoar, P. Jung, and G. Caire, “Grant-Free Massive Random Access With a Massive MIMO Receiver,” *Asilomar Conf. Signals Syst. Comput.*, Nov. 2019.
- [85] A. Fengler, P. Jung, and G. Caire, “Unsourced Multiuser Sparse Regression Codes achieve the Symmetric MAC Capacity,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 3001–3006, June 2020.
- [86] A. Fengler and P. Jung, “On the Restricted Isometry Property of Centered Self Khatri-Rao Products,” *arXiv:1905.09245*, May 2019.
- [87] V. K. Amalladinne, J.-F. Chamberland, and K. R. Narayanan, “A Coded Compressed Sensing Scheme for Unsourced Multiple Access,” *IEEE Trans. Inf. Theory*, 2020.
- [88] D. Guo and S. Verdú, “Randomly Spread CDMA: Asymptotics via Statistical Physics,” *arXiv:cs/0503063*, Mar. 2005.
- [89] D. Guo, S. Shamai, and S. Verdú, “Mutual information and minimum mean-square error in Gaussian channels,” *IEEE Trans. Inf. Theory*, vol. 51, no. 4, pp. 1261–1282, 2005.
- [90] S. Verdú, “Mismatched Estimation and Relative Entropy,” *IEEE Trans. Inf. Theory*, vol. 56, pp. 3712–3720, Aug. 2010.
- [91] D. Guo, Y. Wu, S. Shamai, and S. Verdú, “Estimation in Gaussian noise: Properties of the minimum mean-square error,” *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2371–2385, 2011.
- [92] W. Rudin, *Principles of Mathematical Analysis*. McGraw-Hill, 1976.
- [93] W. Kautz and R. Singleton, “Nonrandom binary superimposed codes,” *IEEE Trans. Inf. Theory*, vol. 10, pp. 363–377, Oct. 1964.
- [94] A. Dyachkov and V. Rykov, “Survey of Superimposed Code Theory,” *Problems of Control and Information Theory*, vol. 12, pp. 229–242, Jan. 1983.
- [95] S. Györi, “Coding for a multiple access OR channel: A survey,” *Discrete Applied Mathematics*, vol. 156, pp. 1407–1430, May 2008.

-
- [96] S.-C. Chang and J. Wolf, "On the T-user M-frequency noiseless multiple-access channel with and without intensity information," *IEEE Trans. Inf. Theory*, vol. 27, pp. 41–48, Jan. 1981.
 - [97] A. J. Grant and C. Schlegel, "Collision-type multiple-user communications," *IEEE Trans. Inf. Theory*, vol. 43, pp. 1725–1736, Sept. 1997.
 - [98] A. Han Vinck and K. Keuning, "On the capacity of the asynchronous T-user M-frequency noiseless multiple-access channel without intensity information," *IEEE Trans. Inf. Theory*, vol. 42, pp. 2235–2238, Nov. 1996.
 - [99] I. Zadik, Y. Polyanskiy, and C. Thrampoulidis, "Improved bounds on Gaussian MAC and sparse regression via Gaussian inequalities," *IEEE Int. Symp. Inf. Theory Proc.*, pp. 430–434, 2019.
 - [100] K. Hsieh, C. Rush, and R. Venkataramanan, "Spatially Coupled Sparse Regression Codes: Design and State Evolution Analysis," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1016–1020, June 2018.
 - [101] A. Yedla, Y.-Y. Jian, P. S. Nguyen, and H. D. Pfister, "A simple proof of threshold saturation for coupled vector recursions," in *2012 IEEE Information Theory Workshop*, pp. 25–29, Sept. 2012.
 - [102] A. Yedla, Y.-Y. Jian, P. S. Nguyen, and H. D. Pfister, "A Simple Proof of Maxwell Saturation for Coupled Scalar Recursions," *IEEE Trans. Inform. Theory*, vol. 60, pp. 6943–6965, Nov. 2014.
 - [103] S. Kudekar, T. Richardson, and R. Urbanke, "Threshold saturation via spatial coupling: Why convolutional LDPC ensembles perform so well over the BEC," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 803–834, 2011.
 - [104] G. Caire, R. R. Müller, and T. Tanaka, "Iterative multiuser joint decoding: Optimal power allocation and low-complexity implementation," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 1950–1973, 2004.
 - [105] V. K. Amalladinne, A. K. Pradhan, C. Rush, J.-F. Chamberland, and K. R. Narayanan, "On Approximate Message Passing for Unsourced Access with Coded Compressed Sensing," *ArXiv200103705 Cs Math*, Jan. 2020.
 - [106] J. Sherman and W. J. Morrison, "Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix," *Ann. Math. Statist.*, vol. 21, pp. 124–127, Mar. 1950.

- [107] R. Adamczak, A. E. Litvak, A. Pajor, and N. Tomczak-Jaegermann, “Restricted Isometry Property of Matrices with Independent Columns and Neighborly Polytopes by Random Sampling,” *Constr. Approx.*, vol. 34, pp. 61–88, Aug. 2011.
- [108] O. Guédon, A. E. Litvak, A. Pajor, and N. Tomczak-Jaegermann, “Restricted isometry property for random matrices with heavy-tailed columns,” *Comptes Rendus Mathématique*, vol. 352, pp. 431–434, May 2014.
- [109] M. Slawski and M. Hein, “Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization,” *Electron. J. Stat.*, vol. 7, pp. 3004–3056, 2013.
- [110] R. Kueng and P. Jung, “Robust Nonnegative Sparse Recovery and the Nullspace Property of 0/1 Measurements,” *IEEE Trans. Inf. Theory*, vol. 64, pp. 689–703, 2018.
- [111] X. Song, S. Haghighatshoar, and G. Caire, “A Scalable and Statistically Robust Beam Alignment Technique for Millimeter-Wave Systems,” *IEEE Trans. Wirel. Commun.*, vol. 17, pp. 4792–4805, July 2018.
- [112] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Springer Texts in Electrical Engineering, New York: Springer-Verlag, second ed., 1994.
- [113] 802.11, “IEEE Standard for Information technology—Telecommunications and information exchange between systems Local and metropolitan area networks—Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” *IEEE Std 80211-2016 Revis. IEEE Std 80211-2012*, pp. 1–3534, Dec. 2016.
- [114] S. Sesia, I. Toufik, and M. Baker, *LTE, The UMTS Long Term Evolution: From Theory to Practice*. Wiley Publishing, 2009.
- [115] L. Ramesh, C. R. Murthy, and H. Tyagi, “Sample-Measurement Tradeoff in Support Recovery Under a Subgaussian Prior,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 2709–2713, July 2019.
- [116] L. Ramesh, C. R. Murthy, and H. Tyagi, “Sample-Measurement Tradeoff in Support Recovery under a Subgaussian Prior,” *ArXiv1912.11247 Cs Math*, Sept. 2020.
- [117] J. Kim, W. Chang, B. Jung, D. Baron, and J. C. Ye, “Belief propagation for joint sparse recovery,” *arXiv:1102.3289*, Feb. 2011.

-
- [118] Z. Chen, F. Sotrab, and W. Yu, “Sparse Activity Detection for Massive Connectivity,” *IEEE Trans. Signal Process.*, vol. 66, pp. 1890–1904, Apr. 2018.
 - [119] A. Javanmard and A. Montanari, “State evolution for general approximate message passing algorithms, with applications to spatial coupling,” *Inf Inference*, vol. 2, pp. 115–144, Dec. 2013.
 - [120] A. Greig and R. Venkataramanan, “Techniques for improving the finite length performance of sparse superposition codes,” *IEEE Trans. Commun.*, vol. 66, no. 3, pp. 905–917, 2018.
 - [121] E. G. Larsson and R. Moosavi, “Piggybacking an Additional Lonely Bit on Linearly Coded Payload Data,” *IEEE Wirel. Commun. Lett.*, vol. 1, pp. 292–295, Aug. 2012.
 - [122] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, “Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios,” *IEEE Wirel. Commun.*, vol. 23, pp. 60–67, Oct. 2016.
 - [123] D. Bankov, E. Khorov, and A. Lyakhov, “On the Limits of LoRaWAN Channel Access,” in *2016 International Conference on Engineering and Telecommunication (EnT)*, pp. 10–14, Nov. 2016.
 - [124] T. S. Rappaport, *WIRELESS COMMUNICATION PRINCIPLES AND PRACTICE*. Theodore S. Rappaport, Jan. 2014.
 - [125] K. Andreev, E. Marshakov, and A. Frolov, “A Polar Code Based TIN-SIC Scheme for the Unsourced Random Access in the Quasi-Static Fading MAC,” *ArXiv200506899 Cs Math*, May 2020.
 - [126] A. Decurninge, I. Land, and M. Guillaud, “Tensor-Based Modulation for Unsourced Massive Random Access,” *ArXiv200606797 Cs Math*, June 2020.
 - [127] D. Jiang and Y. Cui, “ML Estimation and MAP Estimation for Device Activities in Grant-Free Random Access with Interference,” *ArXiv200202595 Cs Eess Math*, Feb. 2020.
 - [128] R. Vershynin, *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2018.
 - [129] R. Adamczak, “A note on the Hanson-Wright inequality for random vectors with dependencies,” *Electron Commun Probab*, vol. 20, no. 72, p. 13 pp., 2015.

- [130] M. Rudelson and R. Vershynin, “Hanson-Wright inequality and sub-gaussian concentration,” *Electron. Commun. Probab.*, vol. 18, no. 82, 2013.