

Chapter 1

Hierarchical Compressed Sensing



Jens Eisert, Axel Flinth, Benedikt Groß, Ingo Roth, and Gerhard Wunder

1.1 Introduction

The field of compressed sensing studies the recovery of structured signals from linear measurements [12, 19]. Originally focusing on the structure of sparsity of vectors, the framework was quickly extended to the structure of low-rankness of matrices. These structures are simultaneously restrictive and rich. They are restrictive so that they allow for signal recovery using far fewer linear measurements than the ambient dimensions suggest and rich in that they naturally appear in a plethora of applications. That being said, in many practically relevant applications, the signals feature a more restrictive structure than mere sparsity or low-rankness. A particularly important broad class arising in a wealth of contexts is *hierarchically structured signals*. Such structures are in the focus of this book chapter.

J. Eisert

Dahlem Center for Complex Quantum Systems and Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany

e-mail: jense@physics.fu-berlin.de

A. Flinth

Institute for Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

e-mail: flinth@chalmers.se

B. Groß · G. Wunder (✉)

Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany

e-mail: benedikt.gross@fu-berlin.de; g.wunder@fu-berlin.de

I. Roth

Quantum Research Centre, Technology Innovation Institute, Abu Dhabi, UAE

Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany

e-mail: i.roth@fu-berlin.de

The perhaps simplest examples are constituted by *hierarchically sparse vectors*. A two-level hierarchically sparse vector is a vector consisting of multiple blocks with a restricted support as follows: only a small number of the blocks have non-vanishing entries and the blocks are themselves sparse. An illustrative example can be given via imagining a telecommunication base station responsible for handling a large set of potential users. If in each instance, only a few users actively transmit, and the messages that are transmitted are sparsely representable, the vector compiling all messages in its blocks is hierarchically sparse. The hierarchically sparse vectors will serve as the main illustrative example of the entire chapter. It is straightforward to generalize this notion for vectors with a hierarchy of nested blocks with sparsity assumptions restricting the number of non-vanishing blocks on each level.

Another hierarchical structure of interest is given by replacing the sparsity constraint on the vector-valued blocks by a low-rank assumption of matrix-valued blocks. A motivating example here can be found in quantum tomography, where quantum states can be modelled as low-rank Hermitian matrices. Hierarchical structures of quantum states arise here in the tasks of performing quantum tomography with a partially uncalibrated measurement device or de-mixing sparse sums of quantum states.

An intriguing feature of hierarchically structured signals is that their recovery task is amenable to efficient thresholding algorithms. In general, thresholding algorithms such as the iterative hard-thresholding pursuit are built on the insight that, in contrast to the original recovery problem, the projection onto the set of structured signals is efficient and in fact often particularly simple. This allows one to employ algorithmic strategies such as projective gradient descent.

For hierarchically sparse signals, it turns out that the calculation of the projection has the same computational complexity as the thresholding onto sparse signals. Furthermore, the hierarchical structure allows for the parallelization of the projections for the blocks on each level yielding potential for further reducing the time complexity by exploiting the restrictive structure. Based on this insight, we formally introduce variants of the *iterative hard-thresholding (IHT) algorithm* and the *hard-thresholding pursuit (HTP)* for hierarchically sparse signals.

For the IHT and HTP algorithm, recovery guarantees for measurement maps that act close to isometrically, on sparse vectors, exist. Due to their similarity, the recovery algorithms for hierarchically sparse signals inherit the recovery guarantees from the original IHT and HTP provided the measurement map exhibits a *restricted isometry property (RIP)* that is adapted to the hierarchically structured signal set. We refer to the modified RIP restricted to hierarchically sparse signals as the *hierarchically restricted isometry property (HiRIP)*.

In this chapter, we derive a series of theoretical results concerning the HiRIP. Requiring only HiRIP instead of RIP for the measurement opens up the possibility of exploiting multiple benefits. First, standard measurement ensembles such as random Gaussian matrices can achieve HiRIP with a reduced sampling complexity compared to RIP. The achievable logarithmic improvement mirrors the reduced complexity of the restricted signal set compared to standard sparse

vectors. Second, we introduce an ensemble of operators that do have the HiRIP, but not RIP in any parameter regime. We give this flexible class of operators the name *hierarchical measurements* since they are naturally adapted to hierarchical structures. Hierarchical measurements combine different measurement maps on each level of the hierarchy and, as we show, inherit HiRIP from standard RIP and coherence properties of their constituent maps. An important instance of hierarchical measurement is Kronecker products of measurements such that each factor acts on the blocks of a certain hierarchy level.

Finally, we illustrate how the framework of hierarchical compressed sensing can be applied in applications in machine-type communications and quantum technologies providing motivating examples and evaluations of practical performances.

Let us end with an outline of the remainder of the chapter. In Sects. 1.2 and 1.3, respectively, we formally introduce hierarchically sparse vectors and present the algorithms used to recover them. Section 1.4 is devoted to theoretical results concerning the hierarchical restricted isometry property (HiRIP) and step by step develops a flexible toolkit to establish the HiRIP for large classes of measurement ensembles. In Sect. 1.5, we move on to discussing the sparse, low-rank signal model, including how the recovery algorithms can be adapted. In Sect. 1.6, we provide a more specific discussion of selected applications. We close with a conclusion as well as a small outlook in Sect. 1.7.

1.2 Hierarchically Sparse Vectors

We consider structured sparse signals that are vectors over the field \mathbb{K} , referring to either the reals \mathbb{R} or the complex numbers \mathbb{C} , and are hierarchically structured into blocks. The support is restricted by sparsity assumptions on one or multiple hierarchy levels. The simplest instance of hierarchically sparse signals is two-level hierarchically sparse vector with constant block sizes and sparsities [20, 41–43].

Definition 1.1 (Two-Level Hierarchically Sparse Vectors) Let $N, n, s, \sigma \in \mathbb{N}$. A vector $x \in \mathbb{K}^{Nn}$ is called (s, σ) -*hierarchically sparse*, if it consists of blocks $x_i \in \mathbb{K}^n$, $x^\top = (x_1^\top, \dots, x_i^\top, \dots, x_N^\top)^\top$, where at most s blocks x_i are non-zero, and each of the non-zero blocks is at most σ -sparse.

For brevity, we write (s, σ) -*sparse*, dropping the *hierarchically* in the following. We refer to the set of (s, σ) -sparse vectors in \mathbb{K}^{Nn} as $\mathcal{S}_{s, \sigma}^{N, n}$ or simply \mathcal{S} if the parameters are clear from the context. We also call the support $\text{supp}(x) \subset [N] \times [n]$ of an (s, σ) -sparse vector a (s, σ) -*sparse support*, where $[n] := \{1, \dots, n\}$. The definition of a two-level hierarchically sparse vectors can be generalized in several directions: We can allow different block sizes and block sparsities. Furthermore, each block is allowed to be a hierarchically sparse vector itself. This gives rise to a more general recursive definition of hierarchically sparse vectors with arbitrary many levels. The defining data of such a general hierarchically sparse vector can

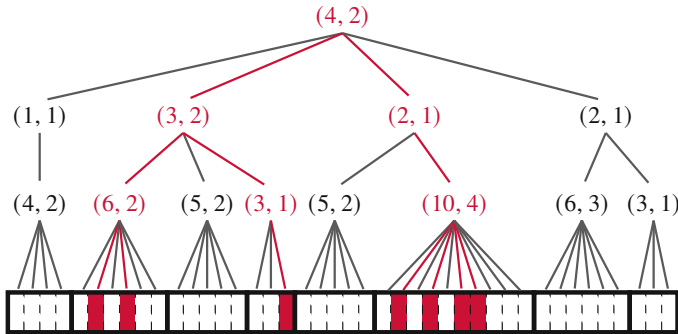


Fig. 1.1 This figure shows an example of a hierarchically sparse vector. The grouping of the entries is encoded in a rooted tree. The children of a vertex constitute a block at their level. The pair of values at each vertex indicates the block size (the number of children) and the sparsity, i.e. the number of children with non-vanishing entries. The leaves of the tree are identified with the entries of the vector. The support of the vector drawn below and corresponding vertices with non-vanishing entries are highlighted in red. ©2020 IEEE. Reprinted, with permission, from Ref. [35]

be collected in a rooted tree consisting of nodes, labelled by block sizes and sparsities, see Fig. 1.1. We refer to Ref. [35] for a formal definition of general hierarchically sparse vectors. Other special cases of hierarchically sparse vectors have been considered in the literature. Prominent examples are *block sparse* [13, 14] and *level sparse* [3, 28] vectors.

Another setting where the hierarchical sparsity naturally emerges is so-called *bi-sparsity*, see e.g. Ref. [18]. In said reference, a Hermitian matrix $X \in \mathbb{K}^{n \times n}$ is called *bisparsity* if there exists a set $S \subseteq [N]$ with $|S| \leq s$ so that X_{ij} is non-zero, only if both i and j are in S . Clearly, any bisparsity matrix can be interpreted as an (s, s) -sparse vector. More generally, a matrix $Y \in \mathbb{K}^{N \times n}$ with Y_{ij} non-zero for $i \in S$ and $j \in \Sigma$ for sets with cardinalities $|S| = s, |\Sigma| \leq \sigma$ can be regarded as (s, σ) -bisparsity and in the same manner identified with an (s, σ) -sparse vector. Bisparsity is of course more restrictive than hierarchical sparsity, but the projection operator onto the set of bisparsity matrices is—in stark contrast to its hierarchical sparsity counterpart—NP-hard to compute. Hierarchical sparsity can thus be seen as a relaxation of bisparsity which allows for more efficient recovery procedures. We refer to Ref. [18] for a more comprehensive discussion on these matters, as well as other ways to relax the bisparsity structures. We encounter this relaxation in conjunction with blind deconvolution in Sect. 1.6.2 and a non-commutative analog of it in our discussion of blind quantum tomography in Sect. 1.6.3.

For simplicity and notational clarity, we content ourselves to present the framework for two-level hierarchically sparse vectors. It is straightforward to generalize the algorithmic strategies and most analytical results of this chapter to the general definition of hierarchically sparse vectors outlined above, see Ref. [35] for details.

1.3 Hierarchical Thresholding and Recovery Algorithms

We study the linear inverse problem of recovering an (s, σ) -hierarchically sparse vector $x \in \mathbb{K}^{Nn}$ from noisy linear measurements of the form

$$y = Mx + e,$$

where $M \in \mathbb{K}^{m \times Nn}$ is the linear measurement operator and $e \in \mathbb{K}^m$ encodes additive noise. The recovery task can be cast as the constraint optimization problem

$$\underset{x \in \mathbb{K}^{Nn}}{\text{minimize}} \quad \frac{1}{2} \|y - Mx\|^2 \quad \text{subject to } x \text{ is } (s, \sigma)\text{-sparse.}, \quad (1.1)$$

where $\|y\| = [\sum_i |y_i|^2]^{1/2}$ denotes the ℓ_2 -norm.

So-called hard-thresholding algorithms solve the analogous optimization problem to (1.1) for standard s -sparse recovery by making use of the projection of a vector $z \in \mathbb{K}^n$ onto the set of s -sparse vectors. The projection onto s -sparse vectors,

$$\mathbb{T}_s(z) := \underset{x \in \mathbb{K}^n}{\text{argmin}} \|x - z\| \quad \text{subject to } x \text{ } s\text{-sparse,}$$

can be computed efficiently via hard thresholding, i.e. by setting to zero all but the s largest entries in absolute value. Note that since the set of s -sparse vectors in \mathbb{K}^n is not a convex set, the projection is non-unique. But for the arguments made here every solution works equally well. Using a quick-select algorithm [24], the average computational complexity of the thresholding operation is in $O(n)$ with worst-case complexity $O(n^2)$.

Following the blueprint of model-based compressed sensing [4], we can derive variants of standard hard-thresholding algorithms for the more restrictive sparsity structure under consideration here by modifying the thresholding operator accordingly. The projection of a vector $z \in \mathbb{K}^{Nn}$ onto the set \mathcal{S} of (s, σ) -hierarchically sparse vectors,

$$\mathbb{T}_{s,\sigma}(z) = \underset{x \in \mathcal{S}}{\text{argmin}} \frac{1}{2} \|x - z\|^2,$$

can be computed via *hierarchical hard thresholding*: first, the standard hard thresholding operation \mathbb{T}_σ is applied to each block. Then, all but the s blocks with largest ℓ_2 -norm are set to zero. The procedure is summarized as Algorithm 1 and illustrated in Fig. 1.2. We find that the average computational complexity of the hierarchical thresholding operation scales as $O(Nn)$, i.e. linear in the overall vector space dimension as for the standard hard thresholding. Furthermore, the hard thresholding and ℓ_2 -norm calculation of the different blocks can be parallelized, reducing the time complexity to $O(\max(N, n))$. The hierarchical thresholding operation can be

Algorithm 1: Hierarchical hard thresholding

input : $z \in \mathbb{K}^{Nn}$, sparsity levels (s, σ)
1 for $i \in [N]$ **do**
2 | $x_i = \mathbb{T}_\sigma(z_i)$;
3 | $v_i = \|x_i\|$;
4 end
5 $I = \text{supp}(\mathbb{T}_s((v_1, \dots, v_N)))$;
6 for $i \in [N] \setminus I$ **do**
7 | $x_i = 0$
8 end
output : (s, σ) -hierarchically sparse vector $x = (x_1^\top, \dots, x_N^\top)^\top$

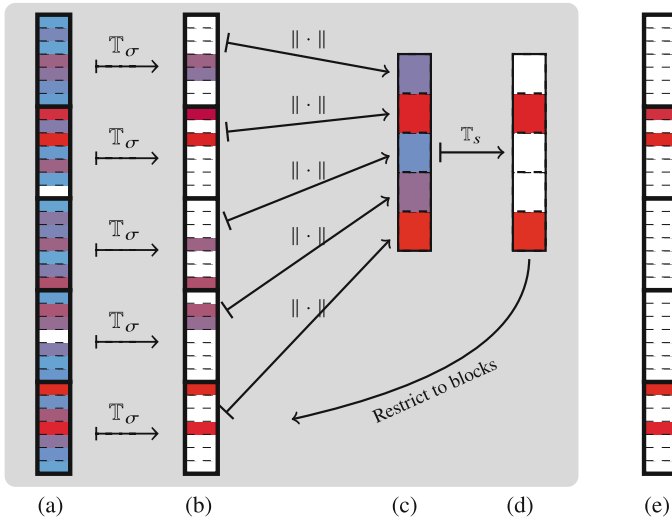


Fig. 1.2 In this figure, the evaluation of the hierarchical thresholding operator $\mathbb{T}_{s,\sigma}$ is illustrated. Starting with a given dense vector (a), each block is thresholded to its best σ -sparse approximation (b). To determine the s dominant blocks, the ℓ_2 -norm is calculated for each block. The resulting vector (c) of length N is again thresholded to its best s -sparse approximation (d). The resulting blocks indicated by the s -sparse approximation (d) are selected from the σ -sparse approximation (b). The remaining (s, σ) -sparse support (e) is the output of $\mathbb{T}_{s,\sigma}$. ©2020 IEEE. Reprinted, with permission, from Ref. [35]

extended recursively to general hierarchically sparse signals described in Sect. 1.2 without increasing the overall computational complexity.

Equipped with an efficient thresholding operation, we can formulate recovery algorithms for hierarchically sparse signals following standard strategies. A particularly simple algorithm is the *iterative hard thresholding algorithm* [5] which performs a projected gradient descent. The resulting *hierarchical iterative hard-thresholding algorithm* (HiHT, Algorithm 2 [53]) alternates gradient descent steps of the objective function (1.1) with the hard-thresholding operation $\mathbb{T}_{s,\sigma}$.

Algorithm 2: HiIHT algorithm

input : data $y \in \mathbb{K}^m$, measurement operator $M \in \mathbb{K}^{m \times Nn}$, sparsity levels (s, σ)
initialize: $x^{(0)} = 0$

1 repeat

2 | $\bar{x}^{(t)} = x^{(t-1)} + \tau^{(t)} M^* (y - Mx^{(t-1)});$

3 | $x^{(t)} = \mathbb{T}_{s, \sigma}(\bar{x}^{(t)});$

4 until *stopping criterion is met at $t = t^*$*

output : (s, σ) -sparse vector $x^{(t^*)}$

Algorithm 3: HiHTP algorithm

input : data $y \in \mathbb{K}^m$, measurement operator $M \in \mathbb{K}^{m \times Nn}$, sparsity levels (s, σ)
initialize: $x^{(0)} = 0$

1 repeat

2 | $\bar{x}^{(t)} = x^{(t-1)} + \tau^{(t)} M^* (y - Mx^{(t-1)});$

3 | $I^{(t)} = \text{supp}(\mathbb{T}_{s, \sigma}(\bar{x}^{(t)}));$

4 | $x^{(t)} = \underset{x}{\text{argmin}} \frac{1}{2} \|y - Mx\|^2 \quad \text{subject to} \quad \text{supp}(x) \subseteq I^{(t)};$

5 until *stopping criterion is met at $t = t^*$*

output : (s, σ) -sparse vector $x^{(t^*)}$

Here, $\tau^{(t)}$ is a suitably chosen step size. The original IHT algorithms use constant steps $\tau^{(t)} = 1$ for all t . Alternative strategies include backtracking as in the normalized iterative hard thresholding (NIHT) algorithm [6].

Faster convergence can be achieved with an adaption of the hard-thresholding pursuit [17] to hierarchical sparsity, the HiHTP [35, 36]. Compared to the HiIHT, the HiHTP algorithm uses the result of the thresholded gradient step as a proxy to guess the support of the correct solution in each step. Subsequently, a linear least-squares problem is solved on the support guess. The solution can be computed via pseudo-inverse or an approximate method. Notably, with this modification, if the algorithm finds the correct solution, it does this in a finite number of steps to the precision of the least-squares problem solver. The HiHTP algorithm is given as Algorithm 3.

The computational complexity of both algorithms, HiIHT and HiHTP, is typically dominated by the matrix vector multiplication with the measurement matrix M and M^* , scaling in general as $O(mNn)$. If a fast matrix vector multiplication is available for the measurement matrix, this scaling can be significantly improved.

The additional least-square solution in the HiHTP algorithm contributes $O(s\sigma m^2)$ operations. For this reason, HiIHT can be faster per iteration than the HiHTP in certain parameter regimes. Note that the computational complexity, featuring the overall vector space dimension Nn and the total sparsity $s\sigma$, is identical to the complexity of the original IHT and HTP algorithms.

Modifications using hierarchically sparse thresholding can also be applied to other compressed sensing algorithms such as the *CoSAMP* [32], the *Subspace Pursuit* [10], or *Orthogonal Matching Pursuit*, see e.g. Refs. [30, 46] and the references therein. Proximal operators of the convex relaxations of the problem

(1.1) can be calculated using soft-thresholding operations yielding a hierarchical version of the LASSO algorithms [20, 41–43]. Due to their similarity, the HiIHT and HiHTP algorithms inherit their convergence proofs and recovery guarantees with slight modifications from their non-hierarchical counterparts. To this end, we make use of the variant of the *restricted isometry property (RIP)* [8] adapted to hierarchically sparse signals.

Definition 1.2 (Hierarchical Restricted Isometry Property (HiRIP)) Given a linear operator $M : \mathbb{K}^{Nn} \rightarrow \mathbb{K}^m$, we denote by $\delta_{s,\sigma}$ the smallest constant such that

$$(1 - \delta_{s,\sigma})\|x\|^2 \leq \|Mx\|^2 \leq (1 + \delta_{s,\sigma})\|x\|^2$$

holds for all (s, σ) -hierarchically sparse vectors $x \in \mathbb{K}^{Nn}$.

We will also refer to the standard s -sparse RIP constant δ_s , defined analogously with the bounds holding for all s -sparse vectors. The standard RIP constant dominates the HiRIP constant as $\delta_{s\sigma} \geq \delta_{s,\sigma}$ since $\mathcal{S}_{s,\sigma}$ is a subset of the set of $s \cdot \sigma$ -sparse vectors. But as we will see below, using the HiRIP allows for a considerably more fine-grained analysis, yielding improvements in the sampling complexity.

In terms of a HiRIP condition, we can guarantee a robust and stable convergence to the correct solution for the hierarchical hard-thresholding algorithms. To this end, given $x \in \mathbb{K}^{Nn}$ and a support set $\Omega \subset [N] \times [n]$, we denote by $x|_{\Omega}$ the projection of x onto the subspace of \mathbb{K}^{Nn} indicated by Ω .

Theorem 1.1 (Recovery Guarantee for HiIHT and HiHTP [35, 53]) *Suppose the measurement operator $M : \mathbb{K}^{Nn} \rightarrow \mathbb{K}^m$ satisfies the HiRIP condition*

$$\delta_{3s,2\sigma} < \delta_*,$$

where δ_* is a threshold, equal to $1/\sqrt{3}$ for the HiHTP algorithm and equal to $\sqrt{2} - 1$ for the HiIHT algorithm. Then, for $x \in \mathbb{K}^{Nn}$, $e \in \mathbb{K}^m$, and $\Omega \subset [N] \times [n]$ an (s, σ) -hierarchically sparse support set, the sequence $(x^k)_k$ defined by HiIHT (Algorithm 2) or HiHTP (Algorithm 3), respectively, with $y = Mx|_{\Omega} + e$ satisfies, for any $k \geq 0$,

$$\|x^k - x|_{\Omega}\| \leq \rho^k \|x^0 - x|_{\Omega}\| + \tau \|e\|,$$

where the constants ρ and τ depend on which algorithm is used: for HiIHT,

$$\rho^{\text{HiIHT}} = \sqrt{3}\delta_{3s,2\sigma}, \quad \tau^{\text{HiIHT}} = \frac{2.18}{1 - \rho^{\text{HiIHT}}},$$

whereas for HiHTP,

$$\rho^{HiHTP} = \left(\frac{2\delta_{3s,2\sigma}}{1 - \delta_{(2s,2\sigma)}^2} \right)^{1/2}, \quad \tau^{HiHTP} = \frac{5.15}{1 - \rho^{HiHTP}}.$$

The theorem's proof follows closely along the lines of the standard proofs for HTP and IHT as found, e.g. in Refs. [7, 17, 19]. A detailed proof can be found in Refs. [35, 53], respectively.

1.4 Hierarchically Restricted Isometric Measurements

The results of the last section make it clear that the HiRIP property has the same role for hierarchically sparse recovery as the RIP takes on for sparse recovery. If we can prove that an operator A , for appropriate hi-sparsity levels (s, σ) , has the HiRIP, it is guaranteed that HiHTP can recover x from the measurements Ax . In this chapter, we will establish the HiRIP for several families of measurement operators, using more and more specialized techniques.

1.4.1 Gaussian Operators

Let us first discuss the HiRIP properties of the arguably most well-known random construction of a measurement operator: the Gaussian random matrix. A random matrix $A \in \mathbb{K}^{m \times n}$ is thereby said to be *Gaussian* if the entries are i.i.d. distributed according to the standard normal distribution $\mathcal{N}(0, 1)$.

It has become a folklore result (see e.g. Ref. [19, Ch.9]) that if A is Gaussian, the renormalized matrix $\frac{1}{\sqrt{m}}A$ has the s -RIP with high probability if

$$m \gtrsim s \log \left(\frac{n}{s} \right),$$

where the notation $\gtrsim f(x)$ means greater than $C \cdot f(x)$, with C an unspecified universal numerical constant. It is therefore natural to ask how large m needs to be in order for $\frac{1}{\sqrt{m}}A$ to have the (s, σ) -HiRIP. Since (s, σ) -sparsity is more restrictive than $s\sigma$ -sparsity, we surely will not need more than $\text{const.} \cdot s\sigma \cdot \log \left(\frac{s\sigma}{nN} \right)$ measurements. But is the threshold lower for the HiRIP? And if so, how much?

In fact, the framework of *model-based compressed sensing* [4] gives us a standard route to answer this question for the Gaussian ensemble. Let us sketch this route in some detail. First, one realizes that for any normalized fixed $x \in \mathbb{K}^N$, the random vector $\frac{1}{\sqrt{m}}Ax$ is also Gaussian and as such obeys the following *measure concentration inequality*:

$$\mathbb{P} \left(\left| \left\| \frac{1}{\sqrt{m}} Ax \right\|^2 - 1 \right| > \delta \right) \leq 2 \exp \left(-cm\delta^2 \right),$$

where c is a numerical constant. For a fixed vector $x \in \mathbb{K}^n$, $\frac{1}{\sqrt{m}}A$ preserves its norm with high probability.

Second, we generalize the almost isometric behaviour to hold for all vectors supported on a certain k -dimensional subspace V . To this end, we first establish that it suffices that the measurement operator acts almost isometrically on a so-called ρ -net for the intersection of the Euclidean unit ball with V . A ρ -net for a set M is a set N with the property that for any $q \in M$, there exists a $p \in N$ with $\|q - p\|_2 < \rho$. It is not hard to construct a ρ -net for the normalized elements of V with cardinality [19]

$$|N| \leq C_{\text{net}} \left(1 + \frac{2}{\rho} \right)^k.$$

By choosing ρ suitably and applying a union bound over the ρ -net, we obtain for any support S with $|S| = k$

$$\mathbb{P} \left(\left| \left\| \frac{1}{\sqrt{m}} Ax \right\|^2 - 1 \right| > \delta \quad \forall x : \text{supp}(x) = S \right) \leq C\lambda^k \exp \left(-\tilde{c}m\delta^2 \right), \quad (1.2)$$

where C , λ , and \tilde{c} are universal numerical constants. With (1.2) at our disposal, it is only one step to establish an isometry property for $\frac{1}{\sqrt{m}}A \in \mathbb{R}^{m \times N \cdot n}$ for an entire union of subspaces such as structured sparse vectors. For instance, in order to get the (s, σ) -HiRIP, we need to take a union bound over all (s, σ) -sparse supports S . There are $\binom{N}{s} \binom{n}{\sigma}^s$ such supports. Therefore

$$\mathbb{P} \left(\left| \left\| \frac{1}{\sqrt{m}} Ax \right\|^2 - 1 \right| > \delta \quad \forall (s, \sigma)\text{-sparse } x \right) \leq C \binom{N}{s} \binom{n}{\sigma}^s \lambda^{s\sigma} \exp \left(-\tilde{c}m\delta^2 \right).$$

This probability is dominated by ϵ , if

$$m \geq \tilde{c}^{-1} \delta^{-2} \log \left(C \binom{N}{s} \binom{n}{\sigma}^s \lambda^{s\sigma} \epsilon^{-1} \right).$$

Using the Stirling approximation $\binom{p}{k} \sim \left(\frac{p}{k} \right)^k$, we obtain the more readable condition

$$m \gtrsim \delta^{-2} \left(s \log \left(\frac{N}{s} \right) + s\sigma \log \left(\frac{n}{\sigma} \right) + \log \left(\frac{1}{\epsilon} \right) \right).$$

Let us state this as a theorem.

Theorem 1.2 (HiRIP for Gaussian Matrices) *Let $A \in \mathbb{K}^{m, n \cdot N}$ be random Gaussian. Then there is a universal numerical constant $C > 0$ so that if*




Fig. 1.3 Left: the number of recovered signals from 100 noiseless Gaussian samples over the number of measurements m for HTP, HiLasso, and HiHTP. The signals consist of $N = 30$ blocks of size $n = 100$ with $s = 4$ blocks having $\sigma = 20$ non-vanishing real entries. Right: the number of recovered blocks over the number of measurements m for HTP and HiHTP. The dashed and dotted lines indicate the average number of correctly recovered zero and non-zero blocks, respectively. The solid lines show the total average number of recovered blocks. The signals consist of $N = 30$ blocks with $s = 4$ blocks having non-vanishing real entries. A signal or block is considered recovered if it deviates from the true signal by less than 10^{-5} in ℓ_2 -norm. ©2020 IEEE. Reprinted, with permission, from Ref. [35]

$$m \geq \frac{C}{\delta^2} \left(s \log \left(\frac{N}{s} \right) + s\sigma \log \left(\frac{n}{\sigma} \right) + \log \left(\frac{1}{\epsilon} \right) \right), \quad (1.3)$$

$\frac{1}{\sqrt{m}}A$ has an (s, σ) -HiRIP constant $\delta_{s,\sigma}(A) \leq \delta$ with probability as least $1 - \epsilon$.

The difference of the condition (1.3) compared to the one needed to establish the standard RIP

$$m \geq \frac{C}{\delta^2} \left(s\sigma \log \left(\frac{Nn}{s\sigma} \right) + \log \left(\frac{1}{\epsilon} \right) \right) \quad (1.4)$$

is subtle. After all, both thresholds can be written as $s\sigma$ multiplied with logarithmic terms in the dimension of surrounding space. However, for certain parameter regimes, the difference is significant. Indeed, in the scenario that $N \gg n$, (1.3) can be much smaller than (1.4). This establishes that for Gaussian random matrices, hierarchical thresholding algorithms are theoretically expected to have an improved sampling complexity compared to their standard counterparts. Also in the non-asymptotic regime, one can observe an improved sample requirement in numerical simulation, see Fig. 1.3.

Note that the above discussion can be applied without problems to *sub-Gaussian* matrices. A matrix is sub-Gaussian if the entries $a_{i,j}$ are i.i.d. distributed according to a distribution that obeys $\mathbb{P}(|a_{i,j}| > t) \leq \alpha \exp(-\beta t^2)$ for some $\alpha, \beta > 0$.

1.4.2 Coherence Measures

The discussion in the last section very much relies on the random nature of the measurement operator. This is a common feature of compressed sensing-related theories—in order to obtain an optimal scaling, one practically has no choice other than to use a random construction. A possible route to still establish (non-optimal) RIP results for non-random matrices is to take a detour via so-called *coherence measures*. The simplest result is as follows [19, Prop 6.2]: if we define the *mutual coherence* of a matrix with normalized columns a_i as

$$\mu(A) = \sup_{i,j} |\langle a_i, a_j \rangle|,$$

the RIP constants obey

$$\delta_s(A) \leq (s - 1)\mu(A). \quad (1.5)$$

To establish analogous results for the HiRIP constants, we need to use coherence measures adapted to the block structure. Such measures have been introduced in Ref. [43] for the analysis of the HiLasso algorithm. To work with these coherence measures, it is convenient to introduce further notation to refer to the blocks of a vector individually. To this end, we use the Kronecker product of matrices in the convention

$$A \otimes B = \begin{pmatrix} a_{1,1}B & \dots & a_{1,N}B \\ \vdots & \ddots & \vdots \\ a_{m,1}B & \dots & a_{m,N}B \end{pmatrix},$$

where $a_{i,j}$ denotes the entries of A . The Kronecker product trivially also provides a Kronecker product on vectors $\mathbb{K}^N \times \mathbb{K}^n \rightarrow \mathbb{K}^{Nn}$ understood as $n \times 1$ and $N \times 1$ matrices, respectively. Using the basis $\{e_i\}_{i \in [N]}$, $(e_i)_j = \delta_{i,j}$ of \mathbb{K}^N , we can rewrite a blocked vector $x \in \mathbb{K}^{Nn}$ with blocks $x_i \in \mathbb{K}^n$, $i \in [N]$, as the sum of products $x = (x_1^\top, x_2^\top, \dots, x_N^\top)^\top = \sum_{i \in [N]} e_i \otimes x_i$. The Kronecker product exemplifies the canonical vector space isomorphism of \mathbb{K}^{Nn} with the tensor product space $\mathbb{K}^N \otimes \mathbb{K}^n$. Analogously, we identify the measurement matrices $A \in \mathbb{K}^{m \times N \cdot n}$ with linear operators $A : \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^m$. We refer to $A_i \in \mathbb{K}^{m \times n}$, $i \in [N]$, defined through $A_i(v) = A(e_i \otimes v)$, $v \in \mathbb{K}^n$, as the *block operators* of A . Now we introduce the specialized coherence measures.

Definition 1.3 (Sub-coherence and Block Coherence) Let $A : \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^m$ with block operators $A_i \in \mathbb{K}^{m \times n}$ and let $\{a_{i,j}\}_{j \in [n]}$ be the columns of the i th block operator. We define

1. The *sub-coherence* $\nu(A)$ of A as the maximal mutual coherence of the block operators, i.e.

$$v(A) = \sup_i \mu(A_i) = \sup_i \sup_{j \neq k} | \langle a_{i,j}, a_{i,k} \rangle |.$$

2. The *sparse block coherence* $\mu_{\text{block}}^{\sigma\sigma}(A)$ of A as

$$\mu_{\text{block}}^{\sigma\sigma}(A) = \sup_{i \neq j} \rho^{\sigma\sigma}(A_i^* A_j),$$

where $\rho^{\sigma\sigma}(B)$ denotes the σ -sparse singular value of a matrix $B \in \mathbb{K}^{N \times N}$,

$$\rho^{\sigma\sigma}(B) = \sup_{\substack{u, v \text{ } \sigma\text{-sparse} \\ \|u\| = \|v\| = 1}} | \langle u, Bv \rangle |.$$

Intuitively, $v(A)$ measures the coherence within each block, whereas $\mu_{\text{block}}^{\sigma\sigma}(A)$ measures the coherence between the blocks. Note that we have used a different normalization in the definition of the sparse block coherence compared to Ref. [43]. We can establish the following bounds on the HiRIP constants in terms of the coherence measures.

Theorem 1.3 (HiRIP Through Coherence Bound) *Let $A : \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^m$ be an operator with block operators A_i and $s \in [N]$, $\sigma \in [n]$. It holds that*

1. $\sup_i \delta_\sigma(A_i) \leq \delta_{1,\sigma}(A)$ and $\mu_{\text{block}}^{\sigma\sigma}(A) \leq 2\delta_{2,\sigma}(A)$.
2. $\delta_{s,\sigma}(A) \leq \sup_i \delta_\sigma(A_i) + (s-1)\mu_{\text{block}}^{\sigma\sigma}(A)$.

In addition, if all columns of the block operators A_i are normalized, then

$$\delta_{s,\sigma}(A) \leq (\sigma-1)v(A) + (s-1)\mu_{\text{block}}^{\sigma\sigma}(A).$$

Proof

1. Let $j \neq k$ and $x, y \in \mathbb{K}^n$ be σ -sparse normalized vectors. First, we have

$$| \|A_j x\|^2 - \|x\|^2 | = | \|A(e_j \otimes x)\|^2 - \|e_j \otimes x\|^2 | \leq \delta_{1,\sigma}(A),$$

since $e_j \otimes x$ is $(1, \sigma)$ -sparse. This proves the first claim. For the second claim, we use the polarization identity to find

$$\langle A_j x, A_k y \rangle = \frac{1}{4} \sum_{\ell=0}^3 i^\ell \|A_j x + i^\ell A_k y\|^2 = \frac{1}{4} \sum_{\ell=0}^3 i^\ell \|A(e_j \otimes x + i^\ell e_k \otimes y)\|^2.$$

Since $e_j \otimes x$ and $e_k \otimes y$ have disjoint block supports, $e_j \otimes x + i^\ell e_k \otimes y$ are $(2, \sigma)$ -sparse for all ℓ . Hence,

$$\left| \frac{1}{4} \sum_{\ell=0}^3 i^\ell \left\| A(e_j \otimes x + i^\ell e_k \otimes y) \right\|^2 - \frac{1}{4} \sum_{\ell=0}^3 i^\ell \left\| e_j \otimes x + i^\ell e_k \otimes y \right\|^2 \right| \leq \delta_{2,\sigma} \cdot \frac{1}{4} \sum_{\ell=0}^3 \left\| e_j \otimes x + i^\ell e_k \otimes y \right\|^2.$$

Now we use that $\|e_j \otimes x + i^\ell e_k \otimes y\|^2 = 2$ for all ℓ . This proves both that the final bound above equals $\frac{1}{2}\delta_{s,\sigma}$ and that $\sum_{\ell=0}^3 i^\ell \|e_j \otimes x + i^\ell e_k \otimes y\|^2 = 0$, yielding the claim.

2. Let $x = \sum_i e_i \otimes x_i$ be an (s, σ) -sparse and normalized signal. There exists an $S \subseteq [N]$ with $|S| = s$ so that $x_i = 0$ for $i \notin S$. We have

$$\|Ax\|^2 = \sum_{i=1}^N \|A_i x_i\|^2 + \sum_{i \neq j} \langle A_i x_i, A_j x_j \rangle.$$

Each x_i is σ -sparse and, thus, $|\|A_i x_i\|^2 - \|x_i\|^2| \leq \delta_\sigma(A_i) \|x_i\|^2$. Taking the sum over i yields

$$\left| \|x\|^2 - \sum_{i=1}^N \|A_i x_i\|^2 \right| \leq \sup_i \delta_\sigma(A_i) \|x\|^2.$$

We still need to deal with the cross-block terms. Let the support of x_i be denoted S_i , the orthogonal projection onto the space supported on S_i with P_{S_i} , and define V as the subspace with the same support as x . Consider the operator $C : V \rightarrow V$,

$$y = \sum_i e_i \otimes y_i \mapsto \sum_{i \in S} e_i \otimes P_{S_i} \left(\sum_{k \in S \setminus \{i\}} A_i^* A_k y_k \right).$$

We have

$$\sum_{i \neq j} \langle A_i x_i, A_j x_j \rangle = \langle x, Cx \rangle,$$

and C is Hermitian. The latter implies that $|\langle x, Cx \rangle| \leq \lambda \|x\|^2$, where λ is the magnitude of the largest eigenvalue of C . To estimate λ , let $v = \sum_i e_i \otimes v_i$ be a normalized eigenvector for C and i such that $\|v_i\|$ is maximal. We have $\lambda v_i = P_{S_i} \sum_{k \in S} A_i^* A_k v_k$, and consequently

$$\lambda \|v_i\|^2 = \langle v_i, P_{S_i} \sum_{k \in S} A_i^* A_k v_k \rangle = \sum_{\substack{k \in S \\ k \neq i}} \langle A_i v_i, A_k v_k \rangle$$

$$\leq \sum_{\substack{k \in S \\ k \neq i}} \mu_{\text{block}}^{\sigma\sigma}(A) \|v_i\| \|v_k\| \leq (s-1) \mu_{\text{block}}^{\sigma\sigma}(A) \|v_i\|^2.$$

In the second step, we have used that $v_i = P_{S_i} v_i$, since $v \in V$. In the penultimate step, we have used that v_i and v_k all are σ -sparse and that each index k in the sum is different from i . In the final step, we have used the optimality of i . This proves that $\lambda \leq (s-1) \mu_{\text{block}}^{\sigma\sigma}(A)$ and therefore the claim.

Finally, the addition of the theorem follows from the claim with (1.5). \square

The above result can be applied to construct a large family of operators that have suitably small HiRIP constants without exhibiting RIP in this regime. Consider N pairwise orthogonal, p -dimensional subspaces of \mathbb{K}^m , and $E_i : \mathbb{K}^p \rightarrow \mathbb{K}^m$ isometric embeddings onto them. Let further $C \in \mathbb{K}^{p \times n}$ be a fixed matrix with $\delta_\sigma(C) = \delta < 1$. We consider the operator

$$A : \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^m, \quad x \mapsto \sum_{i=1}^N E_i C x_i.$$

The block operators of A are given by $E_i C$, $i \in [N]$, and each of them is compressively encoding \mathbb{K}^n into one of the mutually orthogonal subspaces. Due to the fact that the E_i are isometric, $\delta_\sigma(A_i) = \delta_\sigma(C)$ for each i . The pairwise orthogonality of the subspaces implies that $A_i^* A_j = 0$ for $i \neq j$, so that $\mu_{\text{block}}^{\sigma\sigma}(A) = 0$. Theorem 1.3 then implies that $\delta_{s,\sigma}(A) \leq \delta(C)$ for any s .

The above construction will generically not result in an operator with small $\delta_{s\sigma}(A)$. To this end, suppose that $s\sigma \leq n$ and $p \leq n - s\sigma$. Then, there exists an $s\sigma$ -sparse $w \in \mathbb{K}^n$ with $Cw = 0$. Now the vector $x = (w, 0, \dots, 0)$ is $s\sigma$ -sparse, but $|\|Ax\|^2 - \|x\|^2| = |\|0\|^2 - \|x\|^2| = \|x\|^2$. We conclude that $\delta_{s\sigma}(A) \geq 1$.

A disadvantage of this construction is that necessarily $m \geq Np \geq Ns\sigma$. This is a considerably worse scaling than we found for Gaussian random matrices, which exhibit the HiRIP for $m \gtrsim s\sigma$ up to log-factors. The scaling in N as opposed to the sparsity parameter on the block-level s arises from the encoding into mutually orthogonal subspaces. The idea of ‘mixing’ block operators can, however, be driven a lot further to avoid this overhead as we will see in the next section.

1.4.3 Hierarchical Measurement Operators

As we saw above, a measurement operator on $\mathbb{K}^N \otimes \mathbb{K}^n$ can always be thought of as a mixture of block operators, say

$$B(x) = \sum_{i=1}^N B_i x_i.$$

The inequalities in Theorem 1.3, part 1 imply that in order for B to have a small HiRIP constant, we need each block operator to be well-conditioned, and in addition that the blocks are incoherent. What can we do when they are not?

Assume that instead of just observing Bx , we are allowed to sample a few different linear combinations of the vectors $B_i x_i$,

$$y = \left(\sum_{i=1}^N a_{i,j} B_i x_i \right)_{j \in [M]} = \sum_{i=1}^N a_i \otimes B_i x_i,$$

with $a_i = (a_{j,i})_{j \in [M]} \in \mathbb{K}^M$. Can this make recovery easier? Let us define such measurement operators that act hierarchically on the block structure of the vectors as *hierarchical measurement operators*.

Definition 1.4 (Hierarchical Measurement Operators) Let $A \in \mathbb{K}^{M,N}$ and $B_i \in \mathbb{K}^{m,n}$, $i = 1, \dots, N$, be given and denote the i th column of A by a_i . We call the operator

$$\mathcal{H}: \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^M \otimes \mathbb{K}^m, \quad x \mapsto \sum_{i=1}^N a_i \otimes B_i x_i,$$

the *hierarchical measurement operator defined by A and $(B_i)_{i \in [N]}$* .

The structure and naming of hierarchical operators makes it easy to believe that they are an excellent fit for hierarchically sparse recovery. They are, however, by no means only of academic interest. We will discuss this more thoroughly in Sect. 1.6. For now, the practical interest might already become apparent by noting that an important special case of hierarchical measurement operators is the following: in the case of $B_i = B$ being equal, the hierarchical operator is the same as the *Kronecker product* $A \otimes B$ of the matrices A and B . How do the hierarchical isometry constants of \mathcal{H} relate to the ones of A and the B_i s? In order to discuss this question, we begin by proving the following lemma.

Lemma 1.1 (RIP Implies Nuclear Norm Isometry) Let $X \in \mathbb{K}^{N \times N}$ have the property that for some sets S and \bar{S} of cardinality s , $X_{i,j} = 0$ if either $i \notin S$ or $j \notin \bar{S}$.

1. If X is positive-definite Hermitian, which in particular implies $S = \bar{S}$,

$$|\langle A^* A, X \rangle - \|X\|_*| \leq \delta_s(A) \|X\|_*.$$

2. If S and \bar{S} are disjoint,

$$|\langle A^* A, X \rangle| \leq \delta_{2s}(A) \|X\|_*.$$

Here, $\|X\|_*$ denotes the nuclear norm, also known as the trace norm, of X , i.e. the sum of its singular values.

Proof Consider a singular value decomposition of X , $X = \sum_{i=1}^N \sigma_i v_i u_i^*$. We have $\langle A^* A, X \rangle = \sum_{i=1}^N \sigma_i \langle Au_i, Av_i \rangle$. Due to the assumption, for all i with $\sigma_i \neq 0$, $\text{supp}(v_i) \subset S$ and $\text{supp}(u_i) \subset \bar{S}$.

1. For X positive-definite, the σ_i are the eigenvalues of X , and $u_i = v_i$. Since each u_i is s -sparse, it holds that

$$|\langle A^* A, X \rangle - \|X\|_*| \leq \sum_{i=1}^N \sigma_i |\langle Au_i, Au_i \rangle - 1| \leq \sum_{i=1}^N \sigma_i \cdot \delta_s(A) = \delta_s(A) \|X\|_*.$$

2. Ref. [19, Prop. 6.3] states that since the supports of u_i and v_i are disjoint, we have $|\langle Au_i, Av_i \rangle| \leq \delta_{2s}(A)$. This in turn implies

$$|\langle A^* A, X \rangle| \leq \sum_{i=1}^N \sigma_i |\langle Au_i, Av_i \rangle| \leq \sum_{i=1}^N \sigma_i \delta_{2s}(A) = \delta_{2s}(A) \|X\|_*.$$

□

We now prove that \mathcal{H} inherits the HiRIP from the RIP of its constituent matrices, in that $\delta_{s,\sigma}(\mathcal{H})$ can be bounded in terms of $\delta_s(A)$ and the constants $\delta_\sigma(B_i)$.

Theorem 1.4 (Hierarchically Inherited HiRIP) *Let \mathcal{H} be the hierarchical operator defined by A and $(B_i)_{i \in [N]}$. We have for s, σ arbitrary*

$$\delta_{s,\sigma}(\mathcal{H}) \leq \delta_s(A) + \sup_i \delta_\sigma(B_i) + \delta_s(A) \cdot \sup_i \delta_\sigma(B_i).$$

Proof Let x be normalized and (s, σ) -sparse and S such that $a_i = 0$ for $i \notin S$. We have

$$\|\mathcal{H}(x)\| = \sum_{i,j=1}^N \langle a_i \otimes (B_i x_i), a_j \otimes (B_j x_j) \rangle = \sum_{i,j=1}^N \langle a_i, a_j \rangle \langle B_i x_i, B_j x_j \rangle = \langle A^* A, G \rangle,$$

where $G \in \mathbb{K}^{N \times N}$ denotes the matrix with non-vanishing entries $G_{i,j} = \langle B_i x_i, B_j x_j \rangle$ for $i \in S$ and $j \in S$. By Lemma 1.1, part 1,

$$|\langle A^* A, G \rangle - \|G\|_*| \leq \delta_s(A) \|G\|_*. \quad (1.6)$$

It remains to estimate $\|G\|_*$. In order to do this, consider the operator $M : \mathbb{K}^{|S|} \rightarrow \mathbb{K}^m$, $c \mapsto \sum_{i \in S} c_i B_i x_i$. By construction, $G = M^* M$, and therefore, $\|G\|_* = \|M\|^2 = \sum_{i \in S} \|B_i x_i\|^2$, where $\|\cdot\|$ here refers to the Frobenius norm. Consequently,

$$\left| \|G\|_* - \|x\|^2 \right| \leq \sum_{i \in S} \left| \|B_i x_i\|^2 - \|x_i\|^2 \right| \leq \sum_{i \in S} \delta_\sigma(B_i) \|x_i\|^2. \quad (1.7)$$

Combining (1.6) and (1.7), we obtain

$$\begin{aligned} \left| \langle A^* A, G \rangle - \|x\|^2 \right| &\leq \left| \langle A^* A, G \rangle - \|G\|_* \right| + \left| \|G\|_* - \|x\|^2 \right| \\ &\leq \delta_s(A) \left(1 + \sup_i \delta_\sigma(B_i) \right) \|x\|^2 + \sup_i \delta_\sigma(B_i) \|x\|^2, \end{aligned}$$

which proves the claim. \square

The theorem shows that hierarchical operators are a rich class of operators which much more often have the HiRIP than the RIP. To make this precise, we take a look at the special case of Kronecker products $A \otimes B$. Theorem 1.4 implies that if $\delta_s(A)$ and $\delta_\sigma(B)$ are small, $\delta_{s,\sigma}(A \otimes B)$ is also small. This is in stark contrast to the RIP of Kronecker products. Indeed, Ref. [25] derived that

$$\delta_s(A \otimes B) \geq \max(\delta_s(A), \delta_s(B)).$$

That is, in order for $A \otimes B$ to have the s -RIP (nota bene, not the $s\sigma$ -RIP), *both* A and B must have it. This obstacle leads to demanding performance bounds in applications [40].

The total number of measurements measured by a hierarchical operator is equal to mM . Together with the classical results on the RIP of Gaussian operators, the theorem implies that by choosing A and B Gaussian, we can hence build hierarchical operators having the (s, σ) -HiRIP using only

$$\text{const} \cdot s\sigma \log\left(\frac{n}{\sigma}\right) \log\left(\frac{N}{s}\right)$$

many measurements. This scaling is up to log-factors identical to the result Eq. (1.3) we established for fully Gaussian matrices. This is noteworthy, since while fully Gaussian matrix consists of $MN \cdot mn$ independent parameters, a Kronecker product $A \otimes B$ only has $MN + mn$. This constitutes a considerable de-randomization of the measurements, which can be e.g. exploited to reduce the storage complexity or to speed up calculations. We refer to Refs. [34, 35] for an extended discussion and an alternative direct proof of HiRIP for Kronecker product measurements.

Theorem 1.4 tells us that operators with small RIP constants can be combined to obtain an operator with a small HiRIP constant. We now take a look at the contrary question: To what extend are small RIP constants of the constituent operators required to bound the HiRIP constants of the hierarchical measurement operator?

In order to get a simple formulation of our first result, let us first note that there is an ambiguity in the definition of a hierarchical measurement operator. We can always simultaneously rescale a_i and B_i since $a_i \otimes B_i = (\lambda a_i) \otimes (\lambda^{-1} B_i)$. We

may thus w.l.o.g. assume that $\|a_i\| = 1$ for all i . Under this assumption, a small (s, σ) -HiRIP constant of \mathcal{H} indeed implies small σ -RIP constants of all B_i .

Proposition 1.1 (σ -RIP Bound from (s, σ) -HiRIP) *Let \mathcal{H} be a hierarchical measurement operator given by A and $(B_i)_{i \in [N]}$. Assume that the columns of A fulfil $\|a_i\| = 1$ for all i . Then, it holds that*

$$\sup_i \delta_\sigma(B_i) \leq \delta_{1, \sigma}(\mathcal{H}).$$

Proof The i th block operator \mathcal{H}_i of \mathcal{H} is given by $a_i \otimes B_i \in \mathbb{K}^{M \times N}$. The normalization implies that $\|(a_i \otimes B_i)x\|^2 = \|a_i\|^2 \cdot \|B_i x\|^2 = \|B_i x\|^2$ for each $x \in \mathbb{K}^N$. Thus, $\delta_\sigma(B_i) = \delta_\sigma(\mathcal{H}_i)$, and the result follows from Theorem 1.3, part 1. \square

The above result in essence states that for \mathcal{H} to have the (s, σ) -HiRIP, it is necessary that all B_i have the corresponding σ -RIP. Intriguingly, for the RIP requirement of A , the situation is very different. Indeed, if the B_i are mapping into incoherent subspaces, A does not need to have the RIP. The precise result is as follows.

Theorem 1.5 (HiRIP with Block Incoherence) *For a family $(B_i)_{i \in [N]}$, define the operator*

$$\mathcal{B} : \mathbb{K}^N \otimes \mathbb{K}^n \rightarrow \mathbb{K}^m, \quad x \mapsto \sum_{i=1}^N B_i x_i.$$

Let $A \in \mathbb{K}^{M \times N}$ and natural numbers s, σ , and t be given. The hierarchical operator \mathcal{H} given by A and $(B_i)_{i \in [N]}$ fulfils

$$\delta_{t, s, \sigma}(\mathcal{H}) \leq \sup_i \delta_\sigma(B_i) + \delta_s(A) \cdot \sup_i \delta_\sigma(B_i) + t\sqrt{s} \cdot \delta_{2s}(A) \cdot \mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B}).$$

Proof Let $x = \sum_i e_i \otimes x_i$ be a (ts, σ) -sparse, normalized vector and $S \subset [N]$ be such that $x_i = 0$ for $i \notin S$. We may subdivide S into t disjoint sets S_1, \dots, S_t with cardinality s each. For each pair $(k, \ell) \in [t] \times [t]$, we define a matrix $G^{k, \ell} \in \mathbb{K}^{N \times N}$ with non-vanishing entries $G_{i, j}^{k, \ell} = \langle B_i x_i, B_j x_j \rangle$ for $i \in S_k$ and $j \in S_\ell$. We may use the same reasoning as in the proof of Theorem 1.4 to argue that

$$\|\mathcal{H}(x)\|^2 = \sum_{k=1}^t \langle A^* A, G^{k, k} \rangle + \sum_{k \neq \ell} \langle A^* A, G^{k, \ell} \rangle.$$

Now, each matrix $G^{k, \ell}$ fulfils the assumption of Lemma 1.1, part 1 for $k = \ell$ and Lemma 1.1, part 2 for $k \neq \ell$. Hence,

$$\left| \|\mathcal{H}(x)\|^2 - \sum_{k=1}^t \|G^{k,k}\|_*^2 \right| \leq \delta_s(A) \cdot \sum_{k=1}^N \|G^{k,k}\|_* + \delta_{2s}(A) \cdot \sum_{k \neq \ell} \|G^{k,\ell}\|_*.$$

Still in analogy to the proof of Theorem 1.4, we find that $|\|G^{k,k}\|_* - \|x_k\|^2| \leq \sup_i \delta_\sigma(B_i) \|x_k\|^2$, and consequently

$$\left| \mathcal{H}(x) - \|x\|^2 \right| \leq \delta_s(A) \left(1 + \sup_i \delta_\sigma(B_i) \right) + \delta_{2s}(A) \cdot \sum_{k \neq \ell} \|G^{k,\ell}\|_*.$$

It remains to bound the terms with $k \neq \ell$. First, let us note that, since $G^{k,\ell}$ has rank at most s , $\|G^{k,\ell}\|_* \leq \sqrt{s} \|G^{k,\ell}\|$. We now use the definition of the intra-block coherence to argue that

$$\|G^{k,\ell}\| = \sqrt{\sum_{i \in \mathcal{S}_k, j \in \mathcal{S}_\ell} |\langle B_i x_i, B_j x_j \rangle|^2} \leq \mu_{\text{block}}^{(2\sigma, 2\sigma)} \sqrt{\sum_{i \in \mathcal{S}_k, j \in \mathcal{S}_\ell} \|x_i\|^2 \cdot \|x_j\|^2}.$$

Finally with

$$\sum_{k \neq \ell} \sqrt{\sum_{i \in \mathcal{S}_k} \|x_i\|^2} \cdot \sqrt{\sum_{j \in \mathcal{S}_\ell} \|x_j\|^2} \leq \left(\sum_k \sqrt{\sum_{i \in \mathcal{S}_k} \|x_i\|^2} \right)^2 \leq t \|x\|^2,$$

where we have used the Cauchy–Schwarz inequality in the final step, the claim follows. \square

Note that the above result shows that A does not need to have the ts -RIP in order for the hierarchical operator to exhibit the corresponding HiRIP. We may in particular choose $t = N/s$ and obtain an operator that acts isometrically on any vector with sparse blocks. In terms of sample complexity, the above result is still a bit opaque. By making a particular choice of t and using the methods of Gaussian random matrices discussed in Sect. 1.4.1, one can derive the following result (see Ref. [16] for a proof).

Proposition 1.2 (Sample Complexity) *Let $(B_i)_i$ and \mathcal{B} be as in Theorem 1.4.3. Assume that*

$$\left(t \mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B}) \right)^2 \leq \frac{N}{\log(N)},$$

and choose $A \in \mathbb{K}^{M \times N}$ as a Gaussian matrix. Let $\delta, \epsilon > 0$. Provided that

$$M \geq C \left(\left(t \mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B}) \right)^2 \cdot \frac{1}{\delta^2} \log \left(\frac{N (1 + \sup_i \delta_\sigma(B_i))^2}{\left(t \mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B}) \right)^2} \right) + \log \left(\frac{1}{\epsilon} \right) \right),$$

where C is a universal numerical constant, the hierarchical measurement operator \mathcal{H} defined by A and $(B_i)_{i \in N}$ obeys

$$\delta_{t, \sigma}(\mathcal{H}) \leq \delta + \sup_i \delta_\sigma(B_i)$$

with a probability at least $1 - \epsilon$.

This proposition shows that if $\mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B})$ is small enough, the number of ‘Gaussian linear combinations’ we take with A does not have to grow linearly in t in order to establish a (t, σ) -RIP—instead, only $(t \mu_{\text{block}}^{(2\sigma, 2\sigma)}(\mathcal{B}))^2$ is needed.

The square dependence here on $(t \mu_{\text{block}}^{(2\sigma, 2\sigma)})$ is of course inferior compared to the linear dependence of the sparsity we can achieve with the help of Theorem 1.4. It is unclear whether this is merely an artefact of the proof.

These results end our discussion of the hierarchical operators and with that our theoretical results on hierarchical restricted isometry properties.

1.5 Sparse De-mixing of Low-Rank Matrices

Generally, hierarchically sparse vectors arise from recursively assuming nested groupings of the vector entries to be sparsely non-vanishing. Another generalization of hierarchically structured vectors arises when we replace the sparsity assumption with another structure assumption such as a low rank when suitably reshape. One of the simplest of such examples is the de-mixing of a sparse sum of low-rank matrices from linear measurements. For $i \in [N]$, let $\mathcal{A}_i : \mathbb{K}^{n \times n} \rightarrow \mathbb{K}^m$ be linear maps and $\rho_i \in \mathbb{K}^{n \times n}$ be matrices of rank at most r . The problem of *de-mixing low-rank matrices* is to reconstruct the matrices ρ_i given data of the form

$$y = \sum_{i=1}^N \mathcal{A}_i(\rho_i).$$

A further structure assumption might be that out of the N matrices ρ_i actually only a number of s are non-vanishing, giving rise to the problem of de-mixing a sparse sum. We can straightforwardly cast the problem as the reconstruction problem of a hierarchically structured vector. To this end, we set $X = \sum_{i=1}^N e_i \otimes \rho_i$. We can regard X as a ‘vector’ in $\mathbb{K}^{Nn \times n}$ of matrix-valued blocks of rank- r and at most s vanishing blocks.

Algorithm 4: SDT algorithm

input : Data y , measurement \mathcal{A} , sparsity s and rank r of signal
initialize: $X^0 = 0$.
1 repeat
2 | Calculate step-widths μ^l
3 | $X^{l+1} = \bar{\mathbb{T}}_{s,r} \left(X^l + \text{diag}(\mu^l) P_{\mathcal{T}_{X^l}} (\mathcal{A}^* (y - \mathcal{A}(X^l))) \right)$
4 until *stopping criterion is met at* $l = l^*$
output : Recovered signal X^{l^*}

Compared to (s, σ) -sparse vectors, we have replaced the non-vanishing σ -sparse blocks by low-rank matrices. The de-mixing problem of a sparse sum of low-rank matrices then is the task to reconstruct such a hierarchically (block) sparse, (block-wise) low-rank vector X from linear measurements.

The principle strategy of hierarchical hard thresholding of Sect. 1.3 carries over to hierarchically sparse, low-rank vectors. The projection onto the set of rank- r matrices is given by the hard thresholding of the singular values. Let $\rho \in \mathbb{K}^{n \times n}$ have singular value decomposition $U \text{diag}(\Sigma) V^*$ with a vector of singular values $\Sigma \in \mathbb{K}^n$. We define

$$P_r(\rho) = U \text{diag}(\mathbb{T}_r(\Sigma)) V^* .$$

Basically, replacing the application of \mathbb{T}_σ in the hierarchically thresholding Algorithm 1 yields a projection onto hierarchically sparse, low-rank vectors which we will refer to as $\bar{\mathbb{T}}_{s,r}$.

Modifying the projective gradient descent of the HiIHT algorithm with this projection yields the so-called *sparse de-mixing thresholding* (SDT) algorithms, Algorithm 4 [38]. In contrast to the structure of a union of subspaces of sparse vectors, the set of rank r matrices constitutes an embedded differential manifold in the linear vector space of all matrices. The geometrical structure can be exploited in iterative hard-thresholding algorithms by projecting the gradient of the embedding space in the descent step onto the tangent space of the manifold at the current iterate [1, 47, 49]. At point ρ , the tangent space of the manifold of rank- r matrices is the linear span of the set of matrices that have the same row or column space as ρ [1]. For a hierarchically sparse, low-rank vector $X = \sum_{i=1}^N e_i \otimes \rho_i$, we use the projection onto the tangent space for each block. We denote by P_{V_i} and P_{U_i} the projection onto the row and column space of ρ_i , respectively. For ρ_i vanishing, we set the projections to be the identity. We define $P_{\mathcal{T}_X} : \mathbb{K}^{Nn \times n} \rightarrow \mathbb{K}^{Nn \times n}$ as $G = \sum_{i=1}^N e_i \otimes g_i \mapsto \sum_{i=1}^N e_i \otimes [g_i - (\text{Id} - P_{U_i})g_i(\text{Id} - P_{V_i})]$. The particularity of the SDT algorithm is that we allow for a different step size for each matrix block. We refer to Ref. [38] for more details on the algorithm and Ref. [50] for an implementation. The SDT algorithm without the sparse thresholding operation to determine the block support coincides with the algorithm proposed in Ref. [45].




Fig. 1.4 The figure (taken from Ref. [38]) displays the recovery rate for the SDT in different variants for different values of m for random Gaussian measurements. *DT* refers to the SDT algorithms without the sparsity constraint (SD) and *informed DT* to the SDT algorithm restricted to the correct support. A signal is considered successfully recovered, if the algorithm output deviates from the true signal by less than 10^{-3} in Frobenius norm. Each point is averaged over 50 iterations and signal instances with $r = 1$, $n = 16$, $N = 10$, and $s = 3$. One observes nearly coinciding recovery performances for the informed DT and the SDT algorithm. In comparison, the DT algorithm requires significantly more samples for recovery

Following the blueprint of model-based compressed sensing, one can also establish a recovery guarantee based on a RIP condition custom-tailored to the hierarchical structure at hand. For random Gaussian measurement ensembles, this gives rise to a sampling complexity of

$$\delta^{-2} \lceil s \log(N/s) + (2n + 1)rs \log \frac{1}{\delta} \rceil$$

to guarantee the correct recovery of $X \in \mathbb{K}^{Nn \times n}$ with at most s non-vanishing blocks of rank r [38, Theorem 6]. Many results derived in Sect. 1.4 that establish the HiRIP for hierarchically sparse vectors for different measurement ensembles can be generalized to hierarchically sparse, low-rank vectors. This allows one to guarantee recovery by the SDT algorithm for a large class of measurement ensembles.

Compared to an algorithm that does not exploit the sparsity of the de-mixing problem, the SDT algorithm can exhibit a significant improvement in the sampling complexity in relevant parameter regimes, Fig. 1.4.

Hierarchically sparse, low-rank vectors certainly constitute another important class of hierarchically structured signals as it encodes the de-mixing problem of a sparse sum of low-rank matrices. The theme of hierarchically combining low-rank and sparse structure assumptions in nested grouping of entries gives rise to a plethora of structures all of which can be efficiently reconstructed using recursive combinations of the hierarchical thresholding method introduced above.

1.6 Selected Applications

1.6.1 Channel Estimation in Mobile Communication

In mobile communication, a lot of users are simultaneously communicating with a base station through electromagnetic waves. Let us model the messages a user wants to transmit with a sequence $c \in \mathbb{K}^n$. To send this message, the user must first translate the message to a wave. A popular scheme for this is so-called *OFDM (Orthogonal Frequency-Division Multiplexing)*. This scheme can be imagined as each c_k giving rise to a complex exponential, a so-called tone, $b(\omega) = [1, e^{-i\omega t_1}, \dots, e^{-i\omega t_{n-1}}] \in \mathbb{K}^{1 \times n}$, where ω is the frequency and t_1, \dots, t_{n-1} are some discretization times. In OFDM, a fixed grid of the form $\omega_k = 2\pi k \bar{\omega}$, $k \in [n]$, is used, where $\bar{\omega}$ is the normalized frequency. Mathematically, this corresponds to applying the discrete Fourier transform to c .

As the electromagnetic waves travel from the user to the base station, they scatter on random features, e.g. buildings and trees, in the environment. This scattering causes random phase and amplitude shifts, modelled by so-called complex gains ρ_p . It also means that a single transmission results in several incoming wave-fronts, each with a different angle of arrival. This situation can be utilized if the base station has several antennas arranged in an array: when the wave-front arrives at each antenna, the wave-front travels slightly different distances before arriving at each antenna, i.e. if a ‘1’ arrives at antenna 0, antenna k will receive an ‘ $a_k(\theta)$ ’, where θ denotes the angle of the wave-front. Here, $a = [a_0, \dots, a_{n-1}] : [-\pi, \pi] \rightarrow \mathbb{K}^{1 \times n}$ is a function, often referred to as the *antenna manifold* in the communication literature. For the popular *uniform linear array (ULA)*, in which the antennas are placed at a uniform separation d along a straight line, the antenna manifold is after a change of variables $u = d \sin(\theta)$ given by

$$a(u) = [1, e^{2\pi diu}, e^{4\pi diu}, \dots, e^{2(n-1)\pi diu}] .$$

The parameter u actually takes on values in the entirety of $[-d, d]$, but let us for now assume that it lies on some grid $\{-\frac{d}{2N}, \dots, \frac{d}{2N}\}$.

Combining these two models, we see that for a specific user, all transmitted signals result in a collective measurement of the form $\sum_{\ell=1}^L \rho_p a(u_p)^* \langle b(\omega_p)^*, c \rangle$ where (ω_p, u_p) is given by the delay and angle of the k th wave-front. The communication is thus characterized by the *channel matrix* [9]

$$H = \sum_{p=1}^L \rho_p a(u_p)^* b(\omega_p) \in \mathbb{K}^{N \times n} .$$

Once we know H , the base station can easily decode any number of sent messages. Note that as long as the environment and the position of the user do not change drastically, H is expected to stay roughly constant.




Fig. 1.5 Comparison of HiIHT and IHT performances for channel matrix reconstruction. ©2019 IEEE. Reprinted, with permission, from Ref. [53]

Now suppose that we are only given a low-dimensional sample of H . To be concrete, define sub-sampling operators $P_u \in \mathbb{K}^{M \times N}$, $P_\omega \in \mathbb{K}^{m \times n}$ in angle and delay, and assume that we only observe $P_u H P_\omega^\top$. Can we still recover the entire matrix? To do this, we may utilize that, according to the above discussion, it has a sparse representation in the delay-angle domain. Indeed, defining $A = [e^{2k\pi i u_j}]_{k,j \in [N]} \in \mathbb{K}^{N \times N}$ and $B = [e^{-it_k \omega_\ell}]_{k,\ell \in [n]} \in \mathbb{K}^{n \times n}$, we get

$$P_u H P_\omega^\top = P_u A \left(\sum_{p \in [L]} \rho_p e_{u_{j_p}} \otimes e_{\omega_{\ell_p}} \right) B^* P_\omega^\top = (P_u A \otimes P_\omega B) X,$$

with $X = \sum_{p=1}^L \rho_p e_{u_{j_p}} \otimes e_{\omega_{\ell_p}}$. Note that X is not only sparse but also hierarchically sparse: only a few angle blocks are active, and for each such angle, only a few delays ω_k are utilized and vice versa. In fact, it is a reasonable assumption that the angles for the L paths are distinct, leading to a $(1, L)$ -sparse ground truth. We further observe that sampled H is a Kronecker product measurement of X , where the terms of the Kronecker product are sub-sampled Fourier matrices. Thus, the results of Sect. 1.4.3 imply that the recovery indeed is possible and provide an explicit sampling complexity.

In Fig. 1.5, the performance of HiIHT and IHT is compared for $m = n = 256$ and $N = 1024$. We generate data synthetically and inject the measurement with Gaussian noise of an SNR of 10dB. The recovery quality is measured in terms of the mean per-entry square error $\frac{1}{nN} \|H - \hat{H}\|^2$ between the actual channel matrix H and the estimate \hat{H} . This error is plotted against the sub-sampling factor M/N

for different values of L . We see that HiIHT handles a small sub-sampling factor considerably better than IHT. Indeed, only accessing one percent of the available antennas is enough to achieve reasonable performance with HiIHT, whereas IHT fails when less than about 10 percent of the antennas are utilized.

The communication setting presented here can be extended in several directions: first, we may drop the assumptions on the delays and angles to be on a grid—in the off-the-grid case, the vector X is arguably still approximately sparse. Second, we can model the case of multiple users by adding a third level to the hierarchical signal. On this level, sparsity naturally emerges assuming a sporadic user activity. We refer to Ref. [53] for details.

1.6.2 Secure Massive Access

With the rise of new communication technologies such as the Internet of Things (IoT) and Tactile Internet (TI), the amount of devices virtually explodes, and with it the amount of sensitive information gathered from various sensors and transmitted over the air. This development poses significant challenges on the security of communication channels and demands for new physical layers of security. In particular, it calls for fast and scalable low-overhead security schemes suitable for the frequent burst of spontaneous communication between low-complexity devices with a base station. Here, we use the hierarchical measurement framework to design a secure massive access procedure based on blind deconvolution, see also the discussion on bisparsity structures in Sect. 1.2. More details can be found in Ref. [52].

A base station sends out known pilots to enable all *user equipments* (UEs) to measure the channel between the station and the UE. The channel is here modelled as a filter in \mathbb{K}^N , where N is the length of the delay period. For each transmitting UE $p \in [N_d]$ and receiving base station antenna $q \in [N_r]$, there is one filter

$$h_{p,q} = (h_{p,q,1}, \dots, h_{p,q,i}, \dots, h_{p,q,N}) \in \mathbb{K}^N.$$

The concrete appearance of the filters is again determined by delays caused by reflections on random physical features in the environment. Therefore, it is reasonable to assume that each $h_{p,q}$ is sparse, and, for fixed UE q , all channels $h_{p,q}$ for $p = 1, \dots, N_t$ share the same sparsity pattern.

As in the previous section, the UE transmits their sequences $c_p \in \mathbb{K}^E$ by first linearly encoding them into signals $x_p = B_p c_p$ using a codebook $B_p \in \mathbb{K}^{N \times E}$ and then sending them over the channel. In an IoT scenario, the messages typically are very short, so that it can be assumed that they can be encrypted as sparse sequences c_p . During transmissions, these are convolved with the channel vectors, so that each of the base station's antennas receives a superposition of the UEs' signals,

$$y_q = \sum_{p=1}^{N_r} h_{p,q} \circledast (B_p c_p) + z_q$$

with $q = 1, \dots, N_t$ and \circledast denoting the circular convolution. We may now *lift* [29] the bilinear operation $(c_p, h_{p,q}) \rightarrow h_{p,q} \circledast B_p c_p$ to a linear operation $\text{conv}_p : \mathbb{K}^{E \times N} \rightarrow \mathbb{K}^N$ on the matrix $b_p h_{p,q}^\top \in \mathbb{K}^{E \times N}$ resulting in

$$y_q = \sum_{p=1}^{N_r} \text{conv}_p(b_p h_{p,q}^\top) + z_q. \quad (1.8)$$

We observe that the channel estimation task at the base station becomes the problem of simultaneously performing a blind deconvolution and de-mixing, naturally formalized as the linear reconstruction of a signal

$$X_q = (b_1 h_{1,q}^\top, \dots, b_{N_d} h_{N_d,q}^\top) \in (\mathbb{K}^{E \times N})^{N_d} \sim \mathbb{K}^{N_d \cdot E \cdot N}.$$

The signal further exhibits the following structure: our assumptions of σ -sparse channels and s -sparse messages imply that the matrices $b_p h_{p,q}^\top$ are all (s, σ) -bisparsity. As discussed in Sect. 1.2, we may relax this to simple hierarchical (s, σ) -sparsity. Additionally assuming a sparse user activity at a given time, i.e. $b_p \neq 0$ only for μ users, the vector X_q is a three-level (s, σ, μ) -sparse vector. Note that the operator conv_p has a structure that is not covered by our theoretical results. Still, we may try to recover it using the HiHTP algorithm.

We conduct simulations with $N_t = 1$ receive antenna and $N_r = 10$ total UEs. We set $N = 1024$ and $N = E = 128$. For each of the N_r users, a σ -sparse channel $h_k \in \mathbb{R}^E$ is drawn with the locations of the non-zeros distributed uniformly and entries drawn from the standard normal distribution. The signals are computed as $x_k = B c_k$, where $B \in \mathbb{R}^{N \times E}$ is a Gaussian random matrix and $c_k \in \mathbb{R}^E$ is s -sparse with values in $\{-1, 1\}$ if the user is active and 0 if the user is not active. This results in the data $y_1 \in \mathbb{R}^N$ as defined in (1.8).

We vary the number of active users μ , as well as the sparsities s and σ . Figures 1.6 and 1.7 show the rate of successful recovery for varying number of active users, averaged over 20 runs per setup. The x- and y-axis show the channel sparsity μ and the signal sparsity s , respectively. As can be seen, the HiHTP algorithm is indeed capable of recovering the ground truth, as long as the sparsity levels are low enough.

An interesting feature of the model is that it can be used to generate a secure communication scheme. To this end, we make use of the *reciprocity* of the channel: the channel $h_{p,q}^\uparrow$ for transmission from UE q to base station antenna p is equal to the channel $h_{p,q}^\downarrow$ for transmission in the other direction. This reciprocity condition is fulfilled for modern off-the-shelf WiFi devices [48]. Due to the reciprocity,




Fig. 1.6 Recovery rate for 2 of 10 active users. ©2018 IEEE. Reprinted, with permission, from Ref. [52]




Fig. 1.7 Recovery rate for 5 of 10 active users. ©2018 IEEE. Reprinted, with permission, from Ref. [52]

the channel itself can serve as a source of shared randomness for the secret key generation. The communication protocol consists of two phases:

Phase 1

1. The base station sends a predefined pilot signal to all UEs.
2. Each UE q measures the complex-valued channel gains $h_q^\downarrow = (h_{p,q}^\downarrow)_{p \in [N_r]}$.
3. Each UE encrypts his/her message m to a sequence $c_p = f(m, h_q^\downarrow)$, using some encryption scheme f and h_q^\downarrow as a random encryption key.

Phase 2

1. All the UEs q send their encrypted *sequences* c_q to the base station using the scheme discussed above. The encoding operators B_p are left public.
2. The base station receives the superposition of all the convolutions of the cipher text with the respective channels. With a hierarchical thresholding algorithm, the station inverts (1.8) and, thus, gains knowledge of the cipher texts c_p and channels $h_{p,q}^\uparrow$.
3. Due to reciprocity $h_{p,q}^\uparrow = h_{p,q}^\downarrow$, the base station thereby obtains the encryption keys h_q^\downarrow and decrypts the cipher texts.

The security of the scheme relies on the assumption that the channels of different users are independent of each other and cannot be inferred from another position. Unless a man in the middle has access to the antenna of a UE, the eavesdropper cannot use his/her channel coefficients to recover the message of another user.

We note that small variations between both channels, i.e. small violations of reciprocity, can be tolerated by adjusting the key generation process. One can for example quantize the channel gain sufficiently coarse to equalize the keys. Here, the hierarchical framework is applied to solve a blind deconvolution and de-mixing problem. Refs. [16, 51] present further examples of the hierarchical measurement framework applied to massive random access without a built-in security scheme.

1.6.3 Blind Quantum State Tomography

Quantum communication allows for the transmission of data under unprecedented levels of security [21]. Here, the security proofs are neither based on assumptions on the computational hardness of certain mathematical problems nor on the feasibility of practically reverting or predicting the randomness of physical processes: instead, there are proofs of security available based on the fundamental laws of nature themselves. Under mild assumptions, quantum key distribution can be proven secure under the most general attacks allowed by physics, within a paradigm of closed laboratories. Simultaneously, the advent of novel quantum computing devices promises solving certain tasks with a significantly improved computational complexity compared to classical computing devices. These tasks include NP problems at the heart of established and universally employed cryptographic schemes. It is beyond the scope of the present article to introduce the various applications of the quantum technologies [2]. Instead, we here focus on a particular context in which hierarchical compressed sensing naturally comes into play: this is the task of semi-device dependently identifying the state of a quantum device. Methods for such characterization and certification tasks are important diagnostic tools in the development of quantum technologies. We refer to Refs. [11, 27] for details.

The problem at hand here is the identification of quantum states prepared in some physical prescription. The recovery of unknown quantum states is called *quantum*

state tomography. A general quantum state is described by a trace-normalized, positive-definite complex matrix. Of particular interest are unit rank, so-called *pure* quantum states or more generally low-rank quantum states. Ideally, devices in quantum technologies operate or are envisioned to operate in pure quantum states of large dimensions. Quantum states of higher rank encode ‘classical’ statistical mixtures of pure states typically produced by noisy operations. We denote the set of rank- r quantum states by $\mathcal{D}_r^n \subset \mathbb{C}^{n \times n}$.

An important diagnostic task for quantum devices is, thus, to learn the low-rank quantum state of the device from linear measurements. Exploiting the rank constraint on the quantum states in the recovery task is crucial to devise quantum tomography protocols working in state spaces of sizeable dimension. This renders compressed sensing method of crucial importance for quantum tomography [15, 22, 23, 26, 33, 37, 39, 44].

That said, the apparatus with which one performs the measurements can especially for near-term devices not be reasonably assumed to be fully characterized: commonly, there are calibrating parameters that are not fully known. An important practical problem is, thus, the recovery of a low-rank quantum state ρ by means of measurement devices that are simultaneously themselves characterized by a handful of parameters, giving rise to sparse vectors ξ .

In a linear approximation of the measurement device calibration, this leads to the problem of *blind (self-calibrating) quantum state tomography*: let $\mathcal{A} : \mathbb{C}^{nd \times d} \rightarrow \mathbb{R}^m$ be a linear map describing the measurement and calibration model. Given data $y = \mathcal{A}(X) \in \mathbb{R}^m$ and the linear map \mathcal{A} , recover X under the assumption that

$$X \in \{\xi \otimes \rho \mid \xi \in \mathbb{K}^N \text{ } s\text{-sparse, } \rho \in \mathcal{D}_r^n\} \subset \mathbb{C}^{Nn \times n}. \quad (1.9)$$

The blind quantum state tomography problem can be regarded as a non-commutative analog of bisparsity recovery problems where the data is bilinear in two sparse vectors both to be recovered. Similarly to the vector case, already the projection onto the set of structured signal is an NP-hard problem. In fact, one can encode the *sparse PCA problem* [31] and thereby CLIQUE into the task of finding the closest element of the form $\xi \otimes \rho$ with $\xi \in \mathbb{K}^N$, $\rho \in \mathcal{D}_r^n$ to a given $X \in \mathbb{K}^{Nn \times n}$ in Frobenius norm, Ref. [38, Theorem 3]. For this reason, it is not possible to directly derive an efficient algorithm based on a hard-thresholding operation for the blind quantum tomography problem.

However, the problem of de-mixing a sparse sum of low-rank matrices introduced in Sect. 1.5 can be seen as a relaxation to the closest hierarchically structured signal class that still allows for an efficient projection. The analogy to the relation of bisparsity and hierarchical sparsity is imminent.

Consequently, the SDT algorithm is a natural candidate to efficiently tackle the blind tomography problem. Figure 1.8 shows numerical simulations of the performance of the SDT algorithms in the blind quantum tomography task for a random calibration model motivated by quantum technologies in comparison to a standard low-rank tomography algorithm. The relaxation to the hierarchical structured problem, however, comes at the cost of a sub-optimal scaling in complexity theory. While




Fig. 1.8 The figure (taken from Ref. [38]) displays the trace-norm reconstruction error for the SDT algorithm compared to a standard low-rank tomography algorithm for a different number of measurements m of sub-sampled random Pauli measurements. Each point depicts 30 random measurement and signal instances with $r = 1$, $d = 8$, $n = 10$, and $s = 3$. The dotted lines indicate the median. The inline figure shows the mean ℓ_2 -norm reconstruction error of the calibration coefficients for the SDT algorithm

a parameter counting of the original blind tomography problem hints at an optimal scaling of $O(\max\{s \log N, nr\})$, the sparse de-mixing problem introduces already in parameter count an additional factor of s to the second term $O(\max\{s \log N, snr\})$. Due to the sparsity assumption on the calibration parameters, the total number of calibration parameter N still only enters logarithmically. For this reason, the scheme remains highly scalable in practically relevant parameter regimes despite the relaxation. At the same time, using the framework of hierarchical compressed sensing outlined above provides a rich toolkit to equip the SDT with flexible guarantees for many ensembles of measurement and calibration models. Another algorithmic approach to bilinear structured problems such as the blind tomography problem is constraint alternating minimization. We refer to Ref. [38] for further details.

1.7 Conclusion and Outlook

In this chapter, we have introduced a framework for hierarchically compressed sensing with a focus mostly on the reconstruction of hierarchically sparse signals. In its core, standard approaches of compressed sensing naturally generalize to hierarchically structured signals, giving rise to recovery algorithms equipped with theoretical guarantees. Thereby, the successful recovery of hierarchically sparse signals via hard-thresholding algorithms can be established under a custom-tailored restricted isometry assumption. There are, however, a number of specific features that separate the hierarchical framework from its more generic counterparts.

At the heart of the approach is the fact that the projection operator onto the set of hierarchically structured signals is efficiently calculable via hierarchical hard thresholding. Unlike for, e.g., the bisparsity structure, it can be computed in linear time and is highly amenable to parallelization. This in turn renders the simple recovery algorithm interesting in realistic parameter regimes and under practical demands.

Furthermore, within the hierarchical framework, there is a large family of operators that obey the hierarchical, but not the standard restricted isometry property. This makes the framework potentially applicable in settings where standard compressed sensing is infeasible.

On a more theoretical level, the hierarchically sparse structure can be used as a relaxation of the complicated bisparsity structure. In particular, we have presented numerical evidence that instances of the sparse blind deconvolution problem can be solved using HiHTP. And we have invoked the same strategy for the quantum tomography problem and other related questions. While in this context theoretical guarantees are expected to be sub-optimal, the simplicity and flexibility of the hierarchical framework might still be of merit in order to analyze complicated measurement settings. We leave further exploring these matters to future research. A particularly interesting question is to analyze the HiRIP properties of the blind deconvolution operator.

Indeed, we have at the end of this chapter seen several exemplary applications where the hierarchical approach facilitates recovery. This brings us to the arguably most important feature of the framework: hierarchically structured signals naturally emerge in many applications. From our own background and past research, we can conclude this with some confidence. But of course, we very much suspect that there are many applications we are unaware of where the hierarchical framework is readily applicable. For the sake of clarity, we have mainly focused our exposition on the set of two-level hierarchically sparse vectors and merely hinted at the generalizations towards multiple levels potentially mixing low-rankness and sparsity and potentially even further structures that for themselves come with an efficient projection. We hope that we have conveyed that the approach, and even most of the results we presented, can be rather straightforwardly generalized to this rich family of hierarchical signal structures, leaving the playing field wide open.

Acknowledgments This work is a report of some of the findings of the DFG-funded project EI 519/9-1 within the priority programme ‘Compressed Sensing in Information Processing’ (CoSIP), jointly held by J. Eisert and G. Wunder. We specifically thank our coauthors, in particular, M. Barzegar, G. Caire, R. Fritschek, S. Haghighatshoar, D. Hangleiter, M. Kliesch, S. Stefanatos, R. Kueng, and J. Wilkens, with which we have explored this research theme over the years.

References

1. Absil, P.A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton (2009)

2. Acin, A., Bloch, I., Buhrman, H., Calarco, T., Eichler, C., Eisert, J., Esteve, D., Gisin, N., Glaser, S.J., Jelezko, F., Kuhr, S., Lewenstein, M., Riedel, M.F., Schmidt, P.O., Thew, R., Wallraff, A., Walmsley, I., Wilhelm, F.K.: The European quantum technologies roadmap. *New J. Phys.* **20**, 080201 (2018). <https://doi.org/10.1088/1367-2630/aad1ea>
3. Adcock, B., Hansen, A.C., Poon, C., Roman, B.: Breaking the coherence barrier: a new theory for compressed sensing. *Forum Math. Sigma* **5** (2017). <https://doi.org/10.1017/fms.2016.32>
4. Baraniuk, R.G., Cevher, V., Duarte, M.F., Hegde, C.: Model-based compressive sensing. *IEEE Trans. Inf. Theory* **56**, 1982–2001 (2010). <https://doi.org/10.1109/TIT.2010.2040894>
5. Blumensath, T., Davies, M.E.: Iterative thresholding for sparse approximations. *J. Four. An. App.* **14**, 629–654 (2008). <https://doi.org/10.1007/s00041-008-9035-z>
6. Blumensath, T., Davies, M.E.: Normalized iterative hard thresholding: guaranteed stability and performance. *IEEE J. Sel. Top. Sig. Proc.* **4**, 298–309 (2010). <https://doi.org/10.1109/JSTSP.2010.2042411>
7. Bouchot, J.L., Foucart, S., Hitczenko, P.: Hard thresholding pursuit algorithms: number of iterations. *App. Comp. Harm. An.* **41**, 412–435 (2016). <https://doi.org/10.1016/j.acha.2016.03.002>
8. Candes, E.J., Tao, T.: Decoding by linear programming. *IEEE Trans. Inf. Theory* **51**, 4203–4215 (2005). <https://doi.org/10.1109/TIT.2005.858979>
9. Chen, Z., Yang, C.: Pilot decontamination in wideband massive MIMO systems by exploiting channel sparsity. *IEEE Trans. Wirel. Commun.* **15**, 5087–5100 (2016). <https://doi.org/10.1109/TWC.2016.2553021>
10. Dai, W., Milenkovic, O.: Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inf. Theory* **55**, 2230–2249 (2009). <https://doi.org/10.1109/TIT.2009.2016006>
11. Eisert, J., Hangleiter, D., Walk, N., Roth, I., Markham, D., Parekh, R., Chabaud, U., Kashefi, E.: Quantum certification and benchmarking. *Nat. Rev. Phys.* **2**, 382–390 (2020). <https://doi.org/10.1038/s42254-020-0186-4>
12. Eldar, Y.C., Kutyniok, G.: *Compressed Sensing: Theory and Applications*. Cambridge University Press, Cambridge (2012)
13. Eldar, Y.C., Mishali, M.: Block sparsity and sampling over a union of subspaces. In: *Digital Signal Processing, 2009 16th International Conference on*, pp. 1–8 (2009). <https://doi.org/10.1109/ICDSP.2009.5201211>
14. Eldar, Y.C., Mishali, M.: Robust recovery of signals from a structured union of subspaces. *IEEE Trans. Inf. Theory* **55**, 5302–5316 (2009). <https://doi.org/10.1109/TIT.2009.2030471>
15. Flammia, S.T., Gross, D., Liu, Y.K., Eisert, J.: Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New J. Phys.* **14**, 095022 (2012). <https://doi.org/10.1088/1367-2630/14/9/095022>
16. Flinth, A., Groß, B., Roth, I., Eisert, J., Wunder, G.: *Hierarchical isometry properties of hierarchical measurements* (2021)
17. Foucart, S.: Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM J. Num. An.* **49**, 2543–2563 (2011). <https://doi.org/10.1137/100806278>
18. Foucart, S., Gribonval, R., Jacques, L., Rauhut, H.: Jointly low-rank and bisparsity recovery: questions and partial answers. Preprint (2019). ArXiv:1902.04731
19. Foucart, S., Rauhut, H.: *A Mathematical Introduction to Compressive Sensing*. Springer (2013)
20. Friedman, J., Hastie, T., Tibshirani, R.: A note on the group lasso and a sparse group lasso. Preprint (2010). ArXiv: 1001.0736
21. Gisin, N., Ribordy, G., Tittel, W., Zbinden, H.: Quantum cryptography. *Rev. Mod. Phys.* **74**, 145–195 (2002). <https://doi.org/10.1103/RevModPhys.74.145>
22. Gluza, M., Schweigler, T., Rauer, B., Krumnow, C., Schmiedmayer, J., Eisert, J.: Quantum read-out for cold atomic quantum simulators. *Phys. Commun.* **20**, 12 (2020). <https://doi.org/10.1038/s42005-019-0273-y>
23. Gross, D., Liu, Y.K., Flammia, S.T., Becker, S., Eisert, J.: Quantum state tomography via compressed sensing. *Phys. Rev. Lett.* **105**, 150401 (2010). <https://doi.org/10.1103/PhysRevLett.105.150401>

24. Hoare, C.A.R.: Algorithm 65: find. *Commun. ACM* **4**, 321–322 (1961). <https://doi.org/10.1145/366622.366647>
25. Jokar, S., Mehrmann, V.: Sparse solutions to underdetermined Kronecker product systems. *Linear Algebra Appl.* **431**, 2437–2447 (2009). <https://doi.org/10.1016/J.LAA.2009.08.005>
26. Kalev, A., Kosut, R.L., Deutsch, I.H.: Quantum tomography protocols with positivity are compressed sensing protocols. *NJP Quant. Inf.* **1**, 15018 (2015). <https://doi.org/10.1038/nnpjqi.2015.18>
27. Kliesch, M., Roth, I.: Theory of quantum system certification. *PRX Quantum* **2**, 010201 (2021). <https://doi.org/10.1103/PRXQuantum.2.010201>
28. Li, C., Adcock, B.: Compressed sensing with local structure: uniform recovery guarantees for the sparsity in levels class. *Appl. Comput. Harm. Anal.* **46**, 453–477 (2019). <https://doi.org/10.1016/j.acha.2017.05.006>
29. Ling, S., Strohmer, T.: Blind deconvolution meets blind demixing: algorithms and performance bounds. *IEEE Trans. Inf. Theory* **63**, 4497–4520 (2017)
30. Liu, H., Sun, F.: Hierarchical orthogonal matching pursuit for face recognition. In: *The First Asian Conference on Pattern Recognition*, pp. 278–282 (2011). <https://doi.org/10.1109/ACPR.2011.6166530>
31. Magdon-Ismail, M.: NP-hardness and inapproximability of sparse PCA. *Inf. Proc. Lett.* **126**, 35–38 (2017). <https://doi.org/10.1016/j.ipl.2017.05.008>
32. Needell, D., Tropp, J.A.: CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comp. Harm. Anal.* (2008). <https://doi.org/10.1016/j.acha.2008.07.002>
33. Riofrio, C.A., Gross, D., Flammia, S.T., Monz, T., Nigg, D., Blatt, R., Eisert, J.: Experimental quantum compressed sensing for a seven-qubit system. *Nat. Commun.* **8**, 15305 (2017). <https://doi.org/10.1038/ncomms15305>
34. Roth, I., Flinth, A., Kueng, R., Eisert, J., Wunder, G.: Hierarchical restricted isometry property for Kronecker product measurements. In: *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 632–638. IEEE (2018). <https://doi.org/10.1109/ALLERTON.2018.8635829>
35. Roth, I., Kliesch, M., Flinth, A., Wunder, G., Eisert, J.: Reliable recovery of hierarchically sparse signals for Gaussian and Kronecker product measurements. *IEEE Trans. Signal Process.* **68**, 4002–4016 (2020). <https://doi.org/10.1109/tsp.2020.3003453>
36. Roth, I., Kliesch, M., Wunder, G., Eisert, J.: Reliable recovery of hierarchically sparse signals. In: *Proceedings of the third “International Traveling Workshop on Interactions between Sparse models and Technology” (iTWIST’16)*, pp. 36–38 (2016)
37. Roth, I., Kueng, R., Kimmel, S., Liu, Y.K., Gross, D., Eisert, J., Kliesch, M.: Recovering quantum gates from few average gate fidelities. *Phys. Rev. Lett.* **121** (2018). <https://doi.org/10.1103/physrevlett.121.170502>
38. Roth, I., Wilkens, J., Hangleiter, D., Eisert, J.: Semi-device-dependent blind quantum tomography. Preprint (2020). ArXiv:2006.03069
39. Shabani, A., Kosut, R.L., Mohseni, M., Rabitz, H., Broome, M.A., Almeida, M.P., Fedrizzi, A., White, A.G.: Efficient measurement of quantum dynamics via compressive sensing. *Phys. Rev. Lett.* **106**, 100401 (2011). <https://doi.org/10.1103/PhysRevLett.106.100401>
40. Shabara, Y., Koksals, C.E., Ekici, E.: How long to estimate sparse MIMO channels. Preprint. arXiv:2101.07287 (2021)
41. Simon, N., Friedman, J., Hastie, T., Tibshirani, R.: A sparse-group Lasso. *J. Comp. Graph. Stat.* **22**, 231–245 (2013). <https://doi.org/10.1080/10618600.2012.681250>
42. Sprechmann, P., Ramirez, I., Sapiro, G., Eldar, Y.: Collaborative hierarchical sparse modeling. In: *2010 44th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6 (2010). <https://doi.org/10.1109/CISS.2010.5464845>
43. Sprechmann, P., Ramirez, I., Sapiro, G., Eldar, Y.C.: C-HiLasso: a collaborative hierarchical sparse modeling framework. *IEEE Trans. Sig. Proc.* **59**, 4183–4198 (2011). <https://doi.org/10.1109/TSP.2011.2157912>

44. Steffens, A., Riofrío, C.A., McCutcheon, W., Roth, I., Bell, B.A., McMillan, A., Tame, M.S., Rarity, J.G., Eisert, J.: Experimentally exploring compressed sensing quantum tomography. *Quantum Sci. and Technol.* **2**, 025005 (2017). <https://doi.org/10.1088/2058-9565/aa6ae2>
45. Strohmer, T., Wei, K.: Painless breakups-efficient demixing of low rank matrices. *J. Four. Ana. App.* **25**, 1–31 (2019). <https://doi.org/10.1007/s00041-017-9564-4>
46. Tropp, J.A.: Greed is good: algorithmic results for sparse approximation. *IEEE Trans. Inf. Theory* **50**, 2231–2242 (2004). <https://doi.org/10.1109/TIT.2004.834793>
47. Vandereycken, B.: Low-rank matrix completion by Riemannian optimization. *SIAM J. Opt.* **23**, 1214–1236 (2013). <https://doi.org/10.1137/110845768>
48. Vasisht D. Kumar, S., Katabi, D.: Decimeter-level localization with a single WiFi access point. In: NSDI, pp. 165–178 (2016)
49. Wei, K., Cai, J.F., Chan, T.F., Leung, S.: Guarantees of Riemannian optimization for low rank matrix recovery. *SIAM J. Math. Appl.* **37**, 1198–1222 (2016). <https://doi.org/10.1137/15M1050525>
50. Wilkens, J., Hangleiter, D., Roth, I.: (2020). Open source Gitlab repository at <https://gitlab.com/wilkensJ/blind-quantum-tomography>
51. Wunder, G., Flinth, A., Groß, B.: Measure concentration on the OFDM-based massive random access channel. In: 2021 IEEE Statistical Signal Processing Workshop (SSP), pp. 526–530 (2021)
52. Wunder, G., Roth, I., Fritschek, R., Groß, B., Eisert, J.: Secure massive IoT using hierarchical fast blind deconvolution. In: 2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW), pp. 119–124. IEEE (2018)
53. Wunder, G., Stefanatos, S., Flinth, A., Roth, I., Caire, G.: Low-overhead hierarchically-sparse channel estimation for multiuser wideband massive MIMO. *IEEE Trans. Wirel. Commun.* **18**, 2186–2199 (2019)